The Essentials of Monetary Policy

Nathaniel A. Throckmorton*

August 10, 2026

Table 1 lists some essential questions for monetary policy, the corresponding concepts, and why each question matters. For each question, important articles are cited and briefly summarized (with complete references at the end). For additional background, I encourage you to read the section "Central Banking and the Conduct of Monetary Policy" in Mishkin (2021). It provides an accessible introduction to the Federal Reserve, its monetary policy objectives and tools, and the practical considerations that shape how monetary policy is conducted.

1. **Why does the Fed raise interest rates when inflation rises?**

    - Taylor (1993) provides a simple framework for thinking about how the Federal Reserve should adjust interest rates as economic conditions change. A central idea is that when inflation rises above its target, the Fed should raise the nominal interest rate by more than the increase in inflation, thereby increasing the real interest rate and slowing economic activity and inflationary pressure. The Taylor rule also provides a useful benchmark for evaluating monetary policy without implying that the Fed should mechanically follow a formula.
    - Clarida et al. (2000) examine how the Federal Reserve's response to inflation changed over time. They find that monetary policy before Paul Volcker responded too weakly to increases in expected inflation, while the Volcker–Greenspan Fed raised interest rates much more aggressively. Their results illustrate the Taylor principle: to stabilize inflation, the nominal interest rate should generally rise by more than one-for-one with increases in inflation.

2. **What is the Fed trying to accomplish with monetary policy?**

    - Clarida et al. (1999) describe monetary policy as a stabilization problem in which the central bank uses interest rates to pursue its ultimate economic objectives. In their framework, policymakers seek to maintain low and stable inflation while also limiting fluctuations in economic activity around its potential. The paper helps distinguish the federal funds rate from the inflation and real economic outcomes that monetary policy is ultimately intended to stabilize.

*Associate Professor, Economics Department, William & Mary, P.O. Box 8795, Williamsburg, VA 23187. E-mail: nat@wm.edu. Websites: Personal, Google Scholar, IDEAS.

Table 1: Seven Essential Questions for Monetary Policy

#	**Question**	**Concept(s)**	**Why It Matters**
1	Why does the Fed raise interest rates when inflation rises?	Taylor principle and systematic monetary policy	Provides a simple framework for understanding how the Fed can adjust interest rates in response to inflation and economic activity.
2	What is the Fed trying to accomplish with monetary policy?	Dual mandate, economic stabilization, and optimal policy	Explains how the Fed balances its goals of stable prices and maximum employment when deciding where to set interest rates.
3	Why shouldn't the Fed simply do whatever seems best at each meeting?	Rules versus discretion, credibility, and time consistency	Explains why predictable and consistent monetary policy can improve economic outcomes by influencing expectations about future policy.
4	If the Fed does not literally follow a policy rule, what role should rules play?	Policy rules as benchmarks, judgment, and uncertainty	Shows how policy rules can provide useful benchmarks while still allowing policymakers to respond to unusual circumstances and uncertainty.
5	Where does inflation come from?	Demand and supply shocks and the Phillips curve	Helps distinguish between different sources of inflation and understand why the appropriate monetary policy response may depend on what is causing prices to rise.
6	How do we know whether the economy is strong or weak?	Economic activity, unemployment, labor-market slack, and the natural rate	Provides a framework for interpreting economic data and assessing whether conditions are likely to create inflationary or employment pressures.
7	How does the Fed explain its decisions, and why does communication matter?	Expectations, forward guidance, and central-bank communication	Explains how Fed statements and guidance about future policy can affect expectations, financial markets, and the broader economy.

- Federal Open Market Committee (2025), "2025 Statement on Longer-Run Goals and Monetary Policy Strategy", describes how the FOMC puts its dual mandate of maximum employment and price stability into practice. The Fed defines price stability as 2 percent inflation over the longer run, while recognizing that maximum employment cannot be summarized by a single fixed number. When the inflation and employment objectives conflict, the FOMC follows a "balanced approach," considering both how far the economy is from each objective and how long it is expected to take to return to them. The statement provides a useful real-world counterpart to Clarida et al. (1999).
- Chung et al. (2025) examine how different monetary policy strategies perform when the Fed faces uncertainty about how inflation behaves and how quickly it responds to economic conditions. The paper shows that the Fed's goal is not simply to return inflation to 2 percent as quickly as possible, but to pursue price stability while limiting unnecessary fluctuations in employment and economic activity.

3. **Why shouldn't the Fed simply do whatever seems best at each meeting?**

   - Bernanke and Mishkin (1997) explain how a central bank can combine clear policy objectives with flexibility in responding to changing economic conditions. Rather than viewing inflation targeting as a rigid policy rule, they describe it as "constrained discretion": policymakers retain discretion over individual decisions but operate within a framework of stated goals, transparency, and accountability. This provides a useful way to understand why the Fed should not simply do whatever seems best at each meeting, while also explaining why monetary policy cannot be reduced to mechanically following a formula.
   - Mishkin (2017) argues that giving central banks complete discretion can make monetary policy less predictable and weaken the credibility of their commitment to price stability. Instead, discretion should be made more "rule-like" by anchoring policy to clear objectives and communicating how policymakers are likely to respond as economic conditions change.

4. **If the Fed does not literally follow a policy rule, what role should rules play?**

   - Orphanides (2007) explains how simple policy rules, such as the Taylor rule, can provide useful benchmarks for evaluating and guiding monetary policy without being followed mechanically. Comparing actual policy decisions with the prescriptions of simple rules can help policymakers make their decisions more systematic and transparent while still allowing judgment in response to changing economic conditions.
   - Garga et al. (2025) explain how simple policy rules can provide useful benchmarks for monetary policy while uncertainty about the economy and the balance of risks may justify deviations from those benchmarks. The paper shows how policymakers can combine systematic policy with judgment by considering alternative economic scenarios and communicating how policy would respond if conditions develop differently than expected.

5. **Where does inflation come from?**

   - Shapiro et al. (2022) provides a simple framework for distinguishing whether inflation is being driven by changes in demand or supply by examining how prices and quantities move together. Applying this approach to the post-pandemic inflation surge shows that both supply constraints and strong demand contributed to rising prices, illustrating why understanding the source of inflation is essential for determining the appropriate monetary policy response.
   - Bernanke and Blanchard (2025) examine the causes of the post-pandemic inflation surge and show how both supply disruptions and unusually strong demand contributed to rising prices. They also show that the sources of inflation can change over time, with shocks to product markets driving much of the initial inflation surge and tight labor markets playing a greater role in its persistence.
   - Hajdini et al. (2025) review the lessons of the post-pandemic inflation surge, emphasizing how unusually strong demand interacted with supply constraints to generate inflation and how inflation expectations influenced its persistence. The paper shows that understanding where inflation comes from requires considering demand, supply, and expectations together rather than relying on any single explanation.

6. **How do we know whether the economy is strong or weak?**

   - Blanchard and Katz (1997) explain how the natural rate of unemployment provides a benchmark for assessing whether the labor market is tight or weak and whether labor-market conditions are likely to generate inflationary pressure. Because the natural rate is not directly observable and can change over time, policymakers must interpret unemployment alongside other economic evidence rather than relying on a single fixed threshold.
   - Abraham et al. (2020) show that no single statistic provides a complete picture of the strength of the labor market, so economists must consider evidence from unemployment, job vacancies, hiring, quits, and other indicators together. Their analysis illustrates how policymakers can combine multiple measures to assess whether the labor market is tight or weak and what those conditions may imply for inflation and monetary policy.
   - Williams (2023), "Measuring the Natural Rate of Interest: Past, Present, and Future", explains how the natural rate of interest, r^*, provides a benchmark for assessing whether monetary policy is stimulating or restraining economic activity. Because r^* cannot be directly observed and changes over time, policymakers must estimate it from economic data and account for substantial uncertainty when judging how tight or accommodative monetary policy actually is.
   - Holston et al. (2023) estimate the natural rate of interest, r^*, which provides a benchmark for assessing whether interest rates are stimulating or restraining economic activity. Their analysis illustrates both how economists use data to estimate this unobservable benchmark and why uncertainty about r^* makes assessing the stance of monetary policy difficult in real time.

7. **How does the Fed explain its decisions, and why does communication matter?**

   - Woodford (2005) explains that central bank communication is itself an important part of monetary policy because expectations about future interest rates affect longer-term interest rates and economic decisions today. By clearly communicating how policy is likely to respond to economic conditions, the Fed can influence expectations and make its interest-rate decisions more effective.
   - Blinder et al. (2008) review the evidence on how central bank communication affects financial markets, expectations, and the effectiveness of monetary policy. They show that clear communication can make policy decisions more predictable and help the public understand how the central bank is likely to respond as economic conditions change.
   - Tan and Wessel (2026), "What is a 'reaction function' in central banking? How does it differ from 'forward guidance'?", distinguish between forward guidance and the Fed's reaction function. Their discussion of Chair Kevin Warsh's approach shows how the Fed can reduce explicit guidance about future rate decisions while still communicating the principles that guide its decisions.

References

ABRAHAM, K. G., J. C. HALTIWANGER, AND L. E. RENDELL (2020): "How tight is the US labor market?" *Brookings Papers on Economic Activity*, 2020, 97–165.

BERNANKE, B. AND O. BLANCHARD (2025): "What caused the US pandemic-era inflation?" *American Economic Journal: Macroeconomics*, 17, 1–35.

BERNANKE, B. S. AND F. S. MISHKIN (1997): "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives*, 11, 97–116.

BLANCHARD, O. AND L. F. KATZ (1997): "What we know and do not know about the natural rate of unemployment," *Journal of Economic Perspectives*, 11, 51–72.

BLINDER, A. S., M. EHRMANN, M. FRATZSCHER, J. DE HAAN, AND D.-J. JANSEN (2008): "Central bank communication and monetary policy: A survey of theory and evidence," *Journal of economic literature*, 46, 910–945.

CHUNG, H., C. JONES, A. LEPETIT, AND F. M. MARTIN (2025): "Implications of Inflation Dynamics for Monetary Policy Strategies," Tech. Rep. 2025-072, Board of Governors of the Federal Reserve System.

CLARIDA, R., J. GALI, AND M. GERTLER (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

CLARIDA, R., J. GALÍ, AND M. GERTLER (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, 115, 147–180.

FEDERAL OPEN MARKET COMMITTEE (2025): "2025 Statement on Longer-Run Goals and Monetary Policy Strategy," Amended effective August 22, 2025.

GARGA, V., E. HERBST, A. MCKAY, G. NICOLÒ, AND M. PAUSTIAN (2025): "Monetary Policy, Uncertainty, and Communications," Tech. Rep. 2025-074, Board of Governors of the Federal Reserve System.

HAJDINI, I., A. SHAPIRO, A. L. SMITH, AND D. VILLAR (2025): "Inflation since the Pandemic: Lessons and Challenges," Tech. Rep. 2025-070, Board of Governors of the Federal Reserve System.

HOLSTON, K., T. LAUBACH, AND J. C. WILLIAMS (2023): "Measuring the Natural Rate of Interest after COVID-19," Tech. Rep. 1063, Federal Reserve Bank of New York.

MISHKIN, F. S. (2017): "Making Discretion in Monetary Policy More Rule-Like," Working Paper 24135, National Bureau of Economic Research.

——— (2021): *The Economics of Money, Banking, and Financial Markets*, Hoboken, NJ: Pearson, 13 ed.

ORPHANIDES, A. (2007): "Taylor Rules," Tech. Rep. 2007-18, Board of Governors of the Federal Reserve System.

SHAPIRO, A. H. ET AL. (2022): "How much do supply and demand drive inflation," *FRBSF Economic Letter*, 15.

TAN, Z. AND D. WESSEL (2026): "What Is a "Reaction Function" in Central Banking? How Does It Differ from "Forward Guidance"?" The Hutchins Center on Fiscal and Monetary Policy.

TAYLOR, J. B. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

WILLIAMS, J. C. (2023): "Measuring the Natural Rate of Interest: Past, Present, and Future," Remarks at the Thomas Laubach Research Conference, Board of Governors of the Federal Reserve System, Washington, DC.

WOODFORD, M. (2005): "Central bank communication and policy effectiveness," .

Carnegie-Rochester Conference Series on Public Policy 39 (1993) 195-214
North-Holland

Discretion versus policy rules in practice

John B. Taylor*
Stanford University, Stanford, CA 94305

Abstract

This paper examines how recent econometric policy evaluation research on monetary policy rules can be applied in a practical policymaking environment. According to this research, good policy rules typically call for changes in the federal funds rate in response to changes in the price level or changes in real income. An objective of the paper is to preserve the concept of such a policy rule in a policy environment where it is practically impossible to follow mechanically any particular algebraic formula that describes the policy rule. The discussion centers around a hypothetical but representative policy rule much like that advocated in recent research. This rule closely approximates Federal Reserve policy during the past several years. Two case studies—German unification and the 1990 oil-price shock—that had a bearing on the operation of monetary policy in recent years are used to illustrate how such a policy rule might work in practice.

The econometric evaluation of monetary and fiscal policy rules using new methods of "rational expectations" macroeconomics has been the subject of substantially increased research in recent years.[1] A number of factors

*Research was supported by a grant from the National Science Foundation at the National Bureau of Economic Research and by the Stanford Center for Economic Policy Research. I am grateful to Craig Furfine, Ben McCallum, Volker Wieland, and John Williams for helpful comments and assistance.

[1]The forthcoming volume by Bryant, Hooper, and Mann (1993) summarizes much of the empirical research with large multicountry models. A recent Federal Reserve System conference summarized in Taylor (1992) was largely devoted to the analysis of policy rules. A prototype empirical analysis was provided by Taylor (1979) with a full multicountry analysis described in Taylor (1993). Research by McCallum (1988) has also generated considerable interest in econometric evaluation of policy rules. Much of the material in this paper is drawn from Taylor (1993).

0167-2231/93/$06.00 © 1993 – Elsevier Science Publishers B.V. All rights reserved.

have motivated this research: the Lucas critique showing that traditional econometric policy evaluation was flawed, the recognition that rational expectations does not imply monetary policy ineffectiveness, the finding that credibility has empirically significant benefits, and the time inconsistency demonstration that policy rules are superior to discretion. Although one can find precursors of the new research on policy rules, the recent analysis has been made possible by new solution and estimation techniques for economy-wide equilibrium models, the development of empirical models of expectations-consistent wage and price dynamics, and the ability of multi-country empirical frameworks to handle international capital flows in efficient world markets.

The preferred policy rules that have emerged from this research have not generally involved fixed settings for the instruments of monetary policy, such as a constant growth rate for the money supply. The rules are responsive, calling for changes in the money supply, the monetary base, or the short-term interest rate in response to changes of the price level or real income. Some of the research has been quite precise about this response; the coefficients in the algebraic formulas for the policy rules provide exact instructions about how much the Fed should adjust its instruments each quarter in response to an increase in the price level or an increase in real GDP. While the exact coefficients differ from study to study, recently there has been some indication of a consensus about the functional forms and the signs of the coefficients in the policy rules.

Despite the emphasis on policy rules in recent macroeconomic research, the notion of a policy rule has not yet become a common way to think about policy in practice. Policymakers do not, and are not evidently about to, follow policy rules mechanically. Some of the reasons are purely technical. For example, the quarterly time period that has been used to evaluate policy in most econometric models is probably too short to average out blips in the price level due to factors such as temporary changes in commodity prices. On the other hand, a quarter is too long to hold the federal funds rate fixed between adjustments. For example, when the economy starts into recession, sharp and rapid interest-rate declines are appropriate. Many of these technical problems could be corrected, in principle, by modifications of these policy rules. A moving average of the price level over a number of quarters, for example, would be a way to smooth out temporary price fluctuations. Averaging real output—or nominal output—could also be considered. Going to a monthly model—and taking even longer-moving averages—would be a way to make the interest rate more responsive in the very short term. Such generalizations are an important task for future research.

However, these modifications would make the policy rule more complex and more difficult to understand. Even with many such modifications, it

is difficult to see how such algebraic policy rules could be sufficiently encompassing. For example, interpreting whether a rise in the price level is temporary or permanent is likely to require looking at several measures of prices (such as the consumer price index, the producer price index, or the employment cost index). Looking at expectations of inflation as measured by futures markets, the term structure of interest rates, surveys, or forecasts from other analysts is also likely to be helpful. Interpreting the level and the growth rate of the economy's potential output—which frequently is a factor in policy rules—involves predictions about productivity, labor-force participation, and changes in the natural rate of unemployment. While the analysis of these issues can be aided by quantitative methods, it is difficult to formulate them into a precise algebraic formula. Moreover, there will be episodes where monetary policy will need to be adjusted to deal with special factors. For example, the Federal Reserve provided additional reserves to the banking system after the stock-market break of October 19, 1987 and helped to prevent a contraction of liquidity and to restore confidence. The Fed would need more than a simple policy rule as a guide in such cases.

Does all this mean that we must give up on policy rules and return to discretion? In fact, arguments like the one in the previous paragraphs sound much like those used by advocates of discretion rather than rules. Even some of those who have advocated the use of rules in the past seem to have concluded that discretion is the only answer. For example, David Laidler (1991) argues, "We are left, then, with relying on discretionary policy in order to maintain price stability."

If there is anything about which modern macroeconomics is clear however—and on which there is substantial consensus—it is that policy rules have major advantages over discretion in improving economic performance. Hence, it is important to preserve the concept of a policy rule even in an environment where it is practically impossible to follow mechanically the algebraic formulas economists write down to describe their preferred policy rules.

The purpose of this paper is to begin to consider how the recent research on policy rules might apply in such an environment. Section 1 starts with some important semantic issues. Section 2 describes recent results on the design of policy rules that form the basis for this research. Sections 3 and 4 consider the use of such policy rules in practice. For concreteness, I center the discussion around a hypothetical but representative policy rule that is much like that advocated in recent research. This policy rule also describes recent Fed policy surprisingly accurately. I also discuss two case studies—German unification and the 1990 oil-price shock—that had bearing on the operation of monetary policy in recent years.

1. Semantic issues

There is considerable agreement among economists that a policy rule need not be interpreted narrowly as entailing fixed settings for the policy instruments. Although the classic rules versus discretion debate was usually carried on as if the only policy rule were the constant growth rate rule for the money supply, feedback rules in which the money supply responds to changes in unemployment or inflation are also policy rules. In the area of fiscal policy, the automatic stabilizers—transfer payments that automatically rise and tax revenues that automatically grow more slowly with a rise in the unemployment rate—can be interpreted as a "policy." In the area of exchange-rate policy, a fixed exchange-rate system is clearly a policy rule, but so are adjustable or crawling pegs.

Moreover, in my view, a policy rule need not be a mechanical formula, but here there is more disagreement among economists. A policy rule can be implemented and operated more informally by policymakers who recognize the general instrument responses that underlie the policy rule, but who also recognize that operating the rule requires judgment and cannot be done by computer. This broadens the definition of a policy rule significantly and permits the consideration of issues that would be excluded under the narrower definition. By this definition, a policy rule would include a nominal income rule in which the central bank takes actions to keep nominal income on target, but it would not include pure discretionary policy.

In broadening the definition beyond mechanical formulas, I do not mean to lose the concept of a policy rule entirely. Under pure discretion, the settings for the instruments of policy are determined from scratch each period with no attempt to follow a reasonably well-defined contingency plan for the future. A precise analytical distinction between policy rules and discretion can be drawn from the time-consistency literature. In three of the major contributions—Kydland and Prescott (1977), Barro and Gordon (1983), or Blanchard and Fischer (1989)—a policy rule is referred to as the "optimal," the "rules," or the "precommitted" solution, respectively, to a dynamic optimization problem. Discretionary policy is referred to as the "inconsistent," the "cheating," or the "shortsighted" solution, respectively. That literature demonstrates that the advantage of rules over discretion is like the advantage of a cooperative over a noncooperative solution in game theory. This is one of the reasons that researchers have focused on policy rules in recent normative policy research.

As argued above, the term "policy rule" need not necessarily mean either a fixed setting for the policy instruments or a mechanical formula. Saying so, however, does not change common usage. Among most policymakers, the term "policy rule" connotes either a fixed setting for the policy instruments

or a simplistic mechanical procedure. An alternative terminology would help focus attenion on the concept of a policy rule as defined here. For example, one alternative terminology was adopted in the 1990 *Economic Report of the President.* "Policy *rule*" was replaced by "*systematic* policy" or sometimes by "policy *system*" when a noun seemed more appropriate. For example, the 1990 *Economic Report of the President* said, "My Administration will...support a credible, *systematic* monetary policy program that sustains maximum economic growth while controlling and reducing inflation." (p. 4, italics added). The adjective "systematic" is defined in the *Oxford American Dictionary* as "methodical, according to a plan, and not casually or at random." Hence, this word connotes the important properties of a policy rule without bringing along the baggage of fixed settings or mechanical formulas.

With this broader definition of policy rules, comparing the performance of different rules becomes more challenging. Technically speaking, a policy rule is a contingency plan that lasts forever unless there is an explicit cancellation clause. While no policy rule will literally last forever, if a policy rule is to have any meaning, it must be in place for a reasonably long period of time. For a macroeconomic policy rule, several business cycles would certainly be sufficient, but for many purposes several years would do just as well. Policymakers need to make a commitment to stay with the rule if they are to gain the advantages of credibility associated with a rule. If economic analysis is to predict how the economy will perform with a policy rule, some durability of the rule is obviously required. In addition, econometric evaluation of policy rules is of little use if the policy rule is constantly changing.

A final semantic point relates to how different types of policy questions can be described using the language of policy rules. I find it useful to distinguish among three types of policy issues related to policy rules: (1) the design of a policy rule, (2) the transition to a new policy rule once it is designed, and (3) the day-to-day operation of a policy rule once it is in place. As I will describe below, certain policy actions that appear to be discretionary can be interpreted as transitions from one policy rule to another or even as part of the operation of an existing policy rule.

2. Policy design: the search for a good monetary policy rule

The policy design issues I consider in this paper focus entirely on monetary policy. The study of fiscal policy rules—automatic stabilizers or budget-balancing strategies—could be considered using the same approach. The design of fiscal policy rules is an important element of macroeconomic policy analysis despite problems with discretionary fiscal policy. Automatic stabilizers remain an important part of macroeconomic policy and help mitigate recessions. However, automatic stabilizers are affected by goals that go well

beyond macroeconomic policy. For example, changes in the progressivity of the tax system affect the responsiveness of the automatic stabilizers to economic fluctuations but are not made with stabilization policy in mind.

The forthcoming volume by Bryant, Hooper, and Mann (1993) compares what nine different multicountry econometric models say about the performance of different monetary policy rules. Seven of the nine models are estimated rational expectations models. The models were developed by the International Monetary Fund, the Federal Reserve Board, the Department of Finance in Canada, and several individual researchers.

All the policy rules evaluated in the Bryant comparison are interest-rate rules. The monetary authorities are assumed to adjust their interest rate in response either to (1) deviations of the money supply from some target, (2) deviations of the exchange rate from some target, or (3) weighted deviations of the inflation rate (or the price level) and real output from some target.

There are substantial differences from model to model, and there is no agreement on a particular policy rule with particular parameters. Yet there is some consensus. The policy rules that focus on the exchange rate or policies that focus on the money supply do not deliver as good a performance (measured in output and price variability) as policies that focus on the price level and real output directly. In other words, monetary policy rules in which the short-term interest rate instrument is raised by the monetary authorities if the price level and real income are above a target and is lowered if the price level and real income are below target, seem to work well. By how much the interest rate should change is still uncertain, but that a consensus is emerging about a functional form is very promising.

My own research on policy rules reported in Taylor (1993) is generally consistent with these results. Using my multicountry rational expectations model, I simulated economic performance of the G-7 countries under several different monetary policy rules. Economic performance was then examined under the different policy rules. The policy rules were ranked according to how successful they were in achieving price stability and output stability. The approach deals explicitly with several issues raised by the Lucas critique of traditional econometric policy evaluation methods. In fact, the three examples used in the original critique paper of Lucas—consumption demand, price determination, and investment demand—are part of this multicountry model. Endogenizing expectations using the rational expectations assumption, as Lucas did in his original paper, is precisely what automatically happens in this model. To be sure, the equations of the model could benefit from more theoretical research, but the approach does seem appropriate for estimating the long-term effects of different policy regimes.

The approach uses an empirically estimated distribution of shocks. Theoretical studies are useful for highlighting key parameters that affect the

answers. For example, in a standard nonrational expectations model, a fixed exchange-rate system will work better if country-specific shocks to the liquidity preference equations have a relatively large variance. In that case, a fixed exchange-rate system has the same advantages as interest-rate targeting. On the other hand, a flexible exchange-rate system will work better if country-specific shocks to the consumption or investment equations have a relatively large variance. To get any farther than this requires estimates of the size of the shocks.

For the flexible exchange-rate regime, I assumed that each central bank adjusts its short-term interest-rate target in response to changes in the price level and real output from a target. However, for the fixed exchange-rate system, the interest rates in the individual countries cannot be set independently of one other. For example, if the Fed raised the Federal funds rate above the Japanese call money rate, funds would flow quickly into the United States putting upward pressure on the dollar and threatening the fixed rate unless the Bank of Japan likewise raised the call money rate. In order to keep exchange rates from fluctuating, therefore, a common target for the "world" short-term interest rate must be chosen. Analogously with the flexible exchange-rate case, it was assumed that the world short-term interest rate rises if the world price level rises above the target.

My comparison of the flexible exchange-rate system with the fixed exchange-rate system shows that the fluctuations in real output are much larger in the United States, France, Germany, Italy, Japan, and the United Kingdom when exchange rates are fixed, compared with when they are flexible. The standard deviation of output nearly doubles in Germany and Japan under fixed exchange rates in comparison with flexible exchange rates. The fluctuations in real output in Canada are slightly less under fixed rates than under flexible rates, but there is a deterioration of price stability in Canada under fixed exchange rates. A change in the Canadian domestic policy rule under flexible exchange rates could easily match the output stability of the fixed exchange-rate case with more price stability. In this sense the flexible exchange-rate system dominates for all the countries I considered.

Inflation performance is also better with the flexible exchange-rate system than with the fixed-rate system. Price volatility—as measured by the standard deviation of the output deflator around its target—is greater in all countries under fixed exchange rates. Japan and Germany have more than twice as much price volatility under the system that fixes their exchange rate with the dollar.

In addition to finding that it is preferable for the central banks to set interest rates based on economic conditions in their own country (paying little attention to exchange rates), the results show that placing a positive weight on both the price level and real output in the interest-rate rule is

preferable in most countries. Placing some weight on real output works better than a simple price rule, but it is not clear whether the weight on output should be greater than or less than the weight on the price level. A general conclusion from these results is that placing some weight on real output in the interest-rate reaction function is likely to be better than a pure price rule.

Although there is not a consensus about the size of the coefficients of policy rules, it is useful to consider what a representative policy rule might look like. One policy rule that captures the spirit of the recent research and which is quite straightforward is:

$$r = p + .5y + .5(p - 2) + 2 \tag{1}$$

where

r is the federal funds rate,
p is the rate of inflation over the previous four quarters
y is the percent deviation of real GDP from a target.

That is,

y = 100(Y - Y*)/Y* where
Y is real GDP, and
Y* is trend real GDP (equals 2.2 percent per year from 1984.1 through 1992.3).

The policy rule in equation (1) has the feature that the federal funds rate rises if inflation increases above a target of 2 percent or if real GDP rises above trend GDP. If both the inflation rate and real GDP are on target, then the federal funds rate would equal 4 percent, or 2 percent in real terms. (Using the inflation rate over the previous four quarters on the right-hand side of equation (1) indicates that the interest-rate policy rule is written in "real" terms with the lagged inflation rate serving as a proxy for expected inflation.) The 2-percent "equilibrium" real rate is close to the assumed steady-state growth rate of 2.2 percent. This policy rule has the same coefficient on the deviation of real GDP from trend and the inflation rate.

The policy rule in equation 1 has the general properties of the rules that have emerged from recent research, and the coefficients are round numbers that make for easy discussion. What is perhaps surprising is that this rule fits the actual policy performance during the last few years remarkably well. Figure 1 shows the actual path for the federal funds rate and the path implied by the example policy rule during the 1987–1992 period. There is a significant deviation in 1987 when the Fed reacted to the crash in the stock market by

easing interest rates. In this sense the Fed policy has been conducted as if the Fed had been following a policy rule much like the one called for by recent research on policy rules.

For completeness, the paths of the two factors in the policy rule are illustrated in Figures 2 and 3. Note that according to this policy rule, the economy was above trend in the late 1980s and fell below trend during the 1990–91 recession. The gap between actual GDP and trend GDP has narrowed only slightly since the end of the 1990–91 recession. The inflation rate is shown in Figure 3. It certainly appears that the changes in inflation and real GDP influenced the path of the federal funds rate.

3. Discretion versus transitions between policy rules

Most macroeconomic research on policy rules has focused on the design of such rules, as summarized in the previous section. Questions about making a transition from one policy rule to a new policy rule have been given relatively little attention. This situation is not unique to macroeconomics. In general, economists have been better at determining what type of system works best than at determining how to make a transition to that system. In international trade theory, not much is known about the appropriate speed at which one should move to free trade. Also, economists have shown the benefits of a market economy, but there is relatively little research on the transition from one system to another. Because there has been relatively little research in this area and because the problems are harder, there is less formal framework than there is for the design of policy rules.

Examples of transitions

Suppose that it becomes clear that a policy in operation is not performing well and that a new policy system would work better. Suppose, for example, that the target inflation rate in the policy rule in the previous section is shown to be too high. Rather than aim for a 5-percent per year inflation rate, it is recognized that a target of 2-percent per year would be better for long-run economic performance. In this example, only the "intercept" term in the policy rule must be changed. This transition problem is, of course, none other than the problem of disinflation.

Similar examples can be given for fiscal policy rules. Analogous to a change in the intercept in the monetary policy rule would be a recognition that the budget deficit should be balanced at full employment. Analogous to a change in the response coefficient would be a recognition that an increase in the response of the automatic stabilizers to economic conditions would be desirable. The latter might entail a change in the unemployment

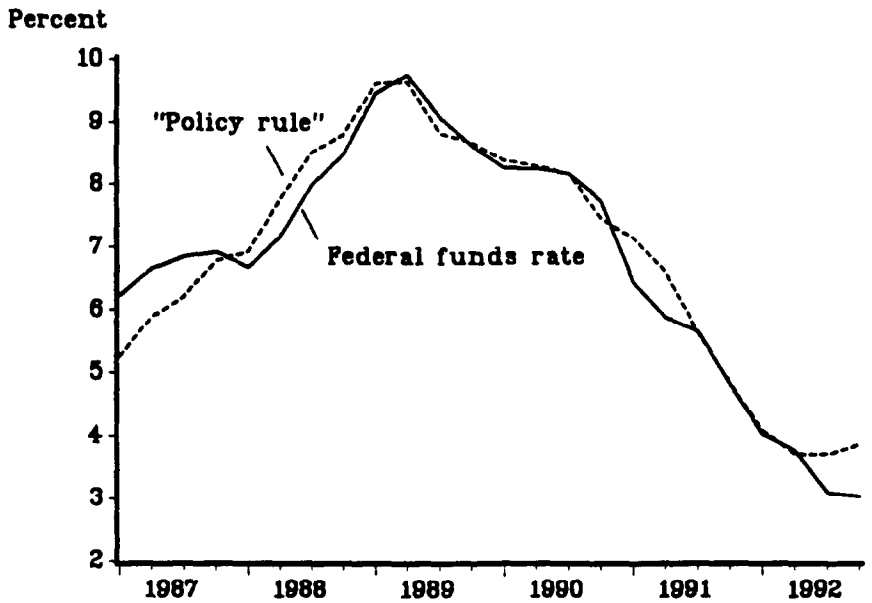


Figure 1. Federal funds rate and example policy rule.

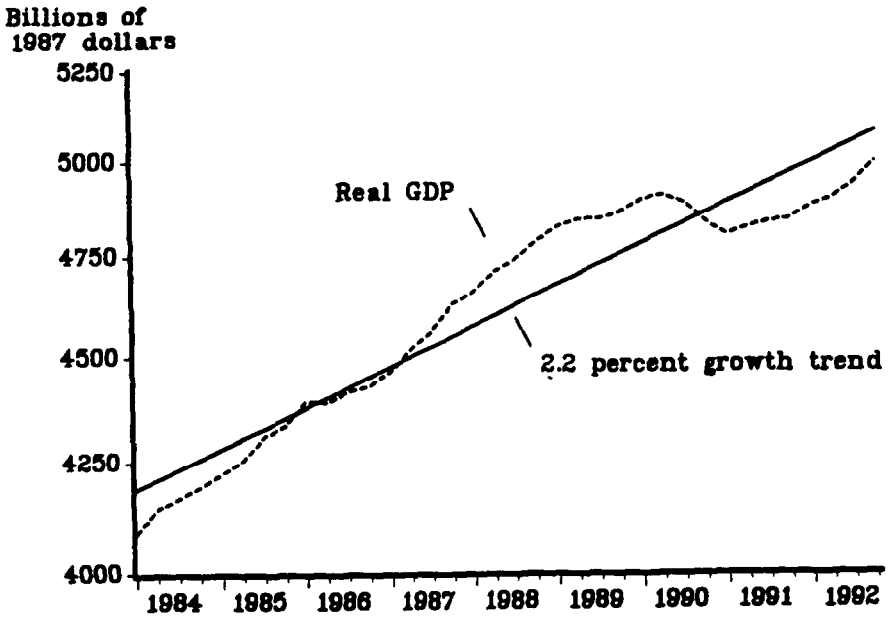


Figure 2. Real GDP and 2.2 percent growth trend.

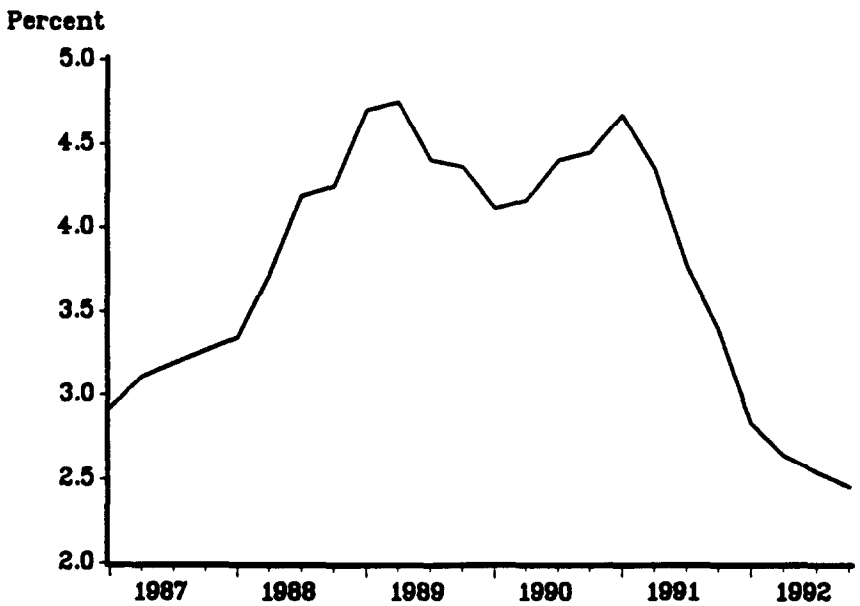


Figure 3. Inflation during previous 4 quarters (GDP deflator).

compensation system that determines at what unemployment rate long-term unemployment benefits are automatically paid.

Why do we need any special treatment of these transitions? First, the research that underlies the design of policy rules assumes that expectations are rational. This makes sense when a policy is in operation for a long time. People will have adjusted their behavior to the policy in place, and expectations of policy and other variables are most likely to be unbiased. However, in the period immediately after a new policy rule has been put in place, people are unlikely either to know about or understand the new policy or to believe that policymakers are serious about maintaining it. Simply assuming that people have rational expectations and know the policy rule is probably stretching things during this transition period. Instead, people may base their expectations partly on studying past policy in a Bayesian way, or by tring to anticipate the credibility of the new policy by studying the past records of policymakers, or by assessing whether the policy will work.

Because expectations only gradually converge during this transition period, the impact of the policy rule on the economy may be quite different than projected by an analysis that assumes rational expectations. This problem of learning about a new policy during a transition was worked out in the case of a change in the price level or inflation target in a very simple model in Taylor (1975). If the initial inflation rate is above the long-run inflation rate, as in the disinflation examples given above, then it is optimal to make the new policy as credible as possible. Announcing the policy and starting to use it is optimal. However, in the case where initial conditions have an inflation rate lower than is optimal, a welfare function that includes both inflation and unemployment can be increased only by gradually informing the public about the plans to move to a new policy. In this unusual case, the precise amount of information to release each period can be computed using optimal control theory.

A second reason for worrying about transitions is that there are natural rigidities in the economy that prevent people from changing their behavior instantly. People may have committed to projects, plans, or contracts under the assumption that the old policy was in place. Moreover, they may have assumed that other people they deal with have similar commitments. Long-term wage-setting commitments are primary examples, but there are many others, including long-term investment projects and loan contracts. Such rigidities suggest that the transition to a new policy rule should be gradual and announced publicly. This gives people a chance to unravel previous commitments without significant losses.

In my view there are many other examples of policy issues that can be usefully interpreted as transitions from one policy rule to another. In practice, however, there is little distinction between such transition issues and

what appears to be pure discretion. To highlight the distinction, I examined two transition problems more explicitly in Taylor (1993): (1) the transition to a monetary policy rule with a zero-inflation target, and (2) the transition toward a fiscal policy rule with a balanced full-employment government budget.

4. Discretion in the operation of policy rules

As stated in the introduction of this paper, operating monetary policy by mechanically following a policy rule like equation 1 is not practical. But how can the constructive results of research as summarized by such a policy rule be made operational? Using equation 1 as an example, I consider two possibilities. One is to try to make use of the specific form of the policy rule as one of the inputs to central bank decision-making. A second is to list the general principles that underlie the policy rule and to leave it up to the policymakers to decide the policy-setting without the guidance of the algebraic formula. Some combination of these two options could also be tried. After describing these two alternative approaches, I consider several case studies to illustrate how they might be used in practice.

Making use of a specific rule

Policymakers, such as the members of the FOMC, currently base their decisions on many factors: leading indicators, the shape of the yield curve, the forecasts of the Fed staff models, etc. There is no reason why a policy rule such as in equation 1 could not be added to the list, at least on an experimental basis. Each time the FOMC meets, the Fed staff could be asked to include in the briefing books information about how recent FOMC decisions compare with the policy rule. Forecasts for the next few quarters—a regular part of the staff briefing—could contain forecasts of the federal funds rate implied by the policy rule. There are many variants on this idea. For instance, there could be a range of entries corresponding to policy rules with different coefficients, or perhaps a policy rule where the growth rate of real GDP rather than its level appears. Bands for the federal funds rate path could span these variants.

At a minimum, experimenting with such a format would bring attention to the concept of a policy rule. "Learning by doing" with the rule, however, would likely bring changes and improvements in the rule and in the format for presenting and using the rule. If the policy rule comes so close to describing actual Federal Reserve behavior in recent years and if FOMC members believe that such performance was good and should be replicated in the future even under a different set of circumstances, then a policy rule could provide some guide to future decisions. This may be particularly relevant when the

membership of the FOMC changes. Such a policy rule could become a guide for future FOMCs.

Making use of general characteristics of policy rules

A second possible approach to making a policy rule operational does not try to use the details of any particular algebraic formulation. Instead, it requires a characterization of the fundamental properties of the rule. Patent laws provide a useful analogy. Patent laws establish the principle that inventors who obtain a patent have the rights to market their invention for a given number of years. The details are left to patent office officials and the court system. Where one draws the line between the fundamentals and the details will depend on many factors.

For example, some of the fundamental features of a monetary policy rule like equation 1 were summarized in the 1990 *Economic Report of the President* as,

> The Federal Reserve generally increases interest rates when inflationary pressures appear to be rising and lowers interest rates when inflationary pressures are abating and recession appears to be more of a threat.... Assessing just how much the policy instrument needs to be changed as circumstances evolve requires judgement. Thus, a policy approach that relies on the expertise of the FOMC members is appropriate and should be preserved. If the operating stance of policy is... measured by interest rates, appropriate settings vary with the interest sensitivity of aggregate demand. (p. 85)

Note that this characterization gives only the signs of the response coefficients of the policy rule. Rather than specifying the magnitudes of the coefficients, it states that the magnitudes should depend on the sensitivity of aggregate demand to interest rates. That is an implication of the design analysis, but it is considerably less specific than stating the magnitudes of the responses.

This characterization is not specific about the target for inflation or for real output. It states only that the federal funds rate should be adjusted when infation rises or falls and when output rises or falls. Certainly, more is needed if the characterization is to effectively convey the fundamental properties of a policy rule like equation 1.

Since the mid-1970s monetary targets have been used in many countries to state targets for inflation. If money velocity were stable, then, given an estimate of potential output growth, money targets would imply a target for the price level; given velocity and a real output target, the target price level would obviously fall out algebraically from the money supply target.

Even though the 1980s have shown that money velocity is not stable in the short run, the long-run stability of the velocity of some monetary measures allows one to state targets for the price level. For example, with an estimated secular growth of real output of $2\frac{1}{2}$ percent and a steady velocity, a money growth range of $2\frac{1}{2}$ percent to $6\frac{1}{2}$ percent—the Fed's targets for 1992—would imply that the price level target grows at 0 to 4 percent per year. Given biases such as index number problems in measuring prices, the 2-percent per year implicit target inflation rate is probably very close to price stability or "zero" inflation.

Case study one: the oil-price shock of 1990

Operating a monetary policy rule in the face of an oil-price shock is difficult and deserves particular study. It is even more difficult if the shock occurs during a transition to a new policy rule with lower inflation as perhaps was occurring in the early 1990s. I focus here on the events that followed the Iraqi invasion of Kuwait on August 2, 1990.

The oil-price shock occurred as the U.S. economy was growing slowly following the 1988-89 monetary tightening—increases in the federal funds rate that had been aimed at containing and reducing the rate of inflation (see Figure 1). If one characterizes Fed actions in terms of the policy rule described above, then the increase in the federal funds rate can be interpreted as occurring for two reasons. First, economic growth in 1987 and 1988 was very strong and inflation was rising; both factors would call for an increase in the federal funds rate according to a policy rule like that in equation 1. Moreover, the Fed had indicated that its intention was to move the economy toward price stability. In other words, the Fed had been attempting to gradually disinflate—to make a transition to greater price stability. In fact, the mean of the target growth rate ranges for the M2 money supply had been reduced from 7 percent in 1987 to 5 percent in 1990, and was reduced to $4\frac{1}{2}$ percent in 1991. The explicit intention of reducing the growth rate targets was to reduce the rate of inflation by an equivalent amount.

Iraq invaded Kuwait on August 2, 1990. Iraq and Kuwait had together been producing 4.3 million barrels of oil a day, and there was a threat to the supply of oil from Saudi Arabia. Not surprisingly the price of oil rose sharply from $21 per barrel at the end of July to $28 on August 6 and eventually to a peak of $46 in mid-October. The monthly average price rose from $17 in July to $36 in October. The effect that this increase in oil prices might have on the economy was of great concern, and major efforts were put in place to estimate the economic impacts. Task forces were assembled and many models—both traditional and forward-looking—were simulated to obtain estimates. The Council of Economic Advisers published a consensus estimate that a one-year temporary increase in oil prices of 50 percent could

temporarily raise the overall price level (GDP deflator) by about 1 percent and with a longer lag, cause real output to fall by about the same amount.

What should be the monetary and fiscal policy reaction to these changes? Suppose that a monetary policy rule like the one described above were in place. Taken literally, equation 1 would say that an increase in the central bank's interest-rate target—relative to what it otherwise would be—was in order; in the short run the price level would rise more than real output would fall. However, such an interest-rate increase would be inappropriate if the price level rise was temporary and would soon disappear. A more complex rule with a longer moving average of inflation or a two-sided moving average including expected inflation might be more appropriate, but the point here is to keep the rule simple.

In fact, analysis at the time suggested that the increase would be temporary. The futures market for oil was helpful in making this assessment. Although the spot price for oil doubled by mid-October, the one-year-ahead futures price changed very little. The December 1991 futures price rose only about $4 per barrel while the spot price rose by $25. Moreover, oil supply analyses suggested that increased oil production elsewhere could eventually make up most of the lost production in Iraq and Kuwait if the embargo continued. The main uncertainty was whether additional oil production facilities would be destroyed before the conflict ended. This uncertainty was dramatically resolved with the successful start of Desert Storm in mid-January 1991.

For these reasons an increase in interest rates to counteract the increase in the price level brought about by the oil shock would be inappropriate–despite the literal interpretation of equation 1. However, not adjusting interest rates in the face of a rising price level would require deviation from the policy rule.

In most countries, the oil-price rise was not viewed as requiring short-term changes in monetary policies. With central banks following monetary strategies that focused on adjustment of interest rates, this position is best interpreted as a policy response like the one discussed above for the United States: interest rates should follow the path that would have occurred without the oil-price shock. There was also a broad consensus that the credibility of economic policies that had been built up in the 1980s should be maintained, and that a clear message be sent that this was the intention of policymakers.

Fiscal policy was also a factor. The automatic stabilizers of fiscal policy provide some built-in response to any negative effects on real output and employment that an oil shock might have, and it was certainly the intention in the Untied States in the summer of 1990 to allow this response to work to mitigate the impact of the oil-price shock on the economy. Some international policy officials raised the possibility of over-riding the automatic stabilizers—offsetting them by increasing taxes or reducing expenditures elsewhere—but others raised strong opposition to such over-rides. Surprisingly, therefore,

there was less consensus about continuing to keep "systematic" fiscal policies in place than there was about monetary policy.

The Gramm-Rudman-Hollings budget law that was still in force in the United States in the summer of 1990 did not allow for the automatic stabilizers. Increases in the budget deficit whether caused by new programs or by the automatic stabilizers were against this law, and would result in across-the-board cuts in spending. The deficit targets would not change even if an oil-price shock worsened economic conditions. Hence, changes in this law were needed if the automatic stabilizers were to be allowed to help stabilize the economy. The revisions in the budget law worked out in the weeks following the oil-price shock required that the budget targets be adjusted for changes in the economy.

Case study two: the bond market, inflation and German unification

Assessing whether an increase in long-term interest rates is due to an increase in expected inflation or to an increase in the real interest rate is part of the task of operating a systematic monetary policy rule. For example, if the policy is to raise interest rates when inflation picks up, then a rise in long-term interest rates might suggest an incipient rise in inflation and might make policymakers less willing to keep the short-term interest rate steady, even if actual inflation does not change. Even so, that increase in long-term interest rates could be due to other factors, such as a shift in the demand for investment or saving.

Such a situation arose in early 1990 before the oil shock discussed above. After declining in the latter part of 1989, long-term interest rates rose sharply in early 1990. Ten-year Treasury bond yields rose by 75 basis points. Concern about a rise in inflation could have caused this increase, and if so could have called for a postponement of declines in interest rates that the monetary policy rule would have called for. However, considerable evidence suggested that other factors were responsible for the increase in long-term rates.

The United States was not the only country to experience an increase in long-term interest rates. Germany had even larger increases, suggesting the possibility that real factors were behind the increase in interest rates. In an integrated world capital market, an increase in interest rates in Germany could be transmitted to U.S. interest rates.

In fact, there was a major change in Germany at this time that could have had such an impact on German long-term rates—anticipations that East Germany and West Germany would be unified and that the unification would increase the demand for capital in Germany and lead to an increase in the government budget deficit in Germany. Greater investment demand would be expected to raise real interest rates in Germany later in 1990 and in 1991, and with forward-looking expectations raise long-term interest rates

immediately. In fact, the anticipated increase in demand for investment and reduction in national saving occurred in 1990 as the unification took place. In 1989, the West German budget was essentially in balance, with a surplus of .2 percent of GDP. That surplus turned dramatically into a deficit of 3-percent of GDP in 1990. Hence, the timing turned out to be correct and consistent with this explanation.

Monetary policy decisions in early 1990, however, could not wait until 1991 when evidence was available about unification and its impact. In early 1990, the analysis had to rely on forecasts and model simulations to see if the magnitudes were plausible. In other words, would an increase in the demand for capital in Germany of plausible magnitudes cause an increase in interest rates of the magnitudes observed? Was it a quantitatively sufficient explanation? Calculations were made with forward-looking empirical models.

Model simulations suggested that increases in interest rates of about one percentage point were consistent with plausible increases in the demand for capital. Hence, an increase in expected inflation was not needed to explain the increase in long-term interest rates. This gave some guidance that interest-rate policy need not be adjusted.

Concluding remarks

This paper has endeavored to study the role of policy rules in a world where simple, algebraic formulations of such rules cannot and should not be mechanically followed by policymakers. Starting with the assumption that systematic and credible features of rule-like behavior improve policy performance, I considered several ways to incorporate rule-like behavior into actual policymaking. Clarification of terms, distinguishing between the design, the transition, and the operation of policy rules, and actually using specific rules or their general features in policy decisions are some of the ideas considered. Two case studies and a hypothetical policy rule illustrated how the ideas could work in practice.

References

Barro, R.J. and Gordon, D.B., (1983). Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics*, **12**: 101–22.

Blanchard, O. and Fischer, S., (1989). *Lectures in Macroeconomics.* Cambridge, Mass.: MIT Press.

Bryant, R., Hooper, P. and Mann, C., (1993). *Evaluating Policy Regimes: New Research in Empirical Macroeconomics.* Washington, D.C.: Brookings Institution.

Economic Report of the President, (1990). Washington, D.C.: U.S. Government Printing Office.

Kydland, F. and Prescott, E.C., (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, **85**: 473–92.

Laidler, D., (1991). Price Stabilization in the 1990s: A Summary, Kumiharu Shigehara, (ed.) *Price Stabilization in the 1990s*, The MacMillan Press Ltd, 1993.

McCallum, B., (1988). Robustness Properties of a Rule for Monetary Policy. *Carnegie Rochester Conference Series on Public Policy*, **29**: 173–203. Amsterdam: North-Holland.

Taylor, J.B., (1975). Monetary Policy During a Transition to Rational Expectations. *Journal of Political Economy*, **83**: 1009–21.

Taylor, J.B., (1979). Estimation and Control of a Macroeconomic Model with Rational Expectations. *Econometrica*, **47**: 1267–86.

Taylor, J.B., (1992). New Directions in Monetary Policy Research: Comments on the Federal Reserve System's Special Meeting on Operating Procedures, *Federal Reserve System Committee on Financial Analysis*, Federal Reserve Bank of St. Louis, June 18–19, 1992.

Taylor, J.B., (1993). *Macroeconomic Policy in a World Economy: From Econometric Design to Practical Operation*, New York: W.W. Norton.

MONETARY POLICY RULES AND MACROECONOMIC STABILITY: EVIDENCE AND SOME THEORY*

RICHARD CLARIDA
JORDI GALÍ
MARK GERTLER

We estimate a forward-looking monetary policy reaction function for the postwar United States economy, before and after Volcker's appointment as Fed Chairman in 1979. Our results point to substantial differences in the estimated rule across periods. In particular, interest rate policy in the Volcker-Greenspan period appears to have been much more sensitive to changes in expected inflation than in the pre-Volcker period. We then compare some of the implications of the estimated rules for the equilibrium properties of inflation and output, using a simple macroeconomic model, and show that the Volcker-Greenspan rule is stabilizing.

I. INTRODUCTION

From the late 1960s through the early 1980s, the United States economy experienced high and volatile inflation along with several severe recessions. Since the early 1980s, however, inflation has remained steadily low, while output growth has been relatively stable. Many economists cite supply shocks—and oil price shocks, in particular—as the main force underlying the instability of the earlier period. It is unlikely, however, that supply shocks alone could account for the observed differences between the two eras. For example, while jumps in the price of oil might help explain transitory periods of sharp increases in the general price level, it is not clear how they alone could explain persistent high inflation in the absence of an accommodating monetary policy.[1] Furthermore, as De Long [1997] argues, the onset of sustained high inflation occurred prior to the oil crisis episodes.

* We thank seminar participants at Universitat Pompeu Fabra, New York University, the Board of Governors of the Federal Reserve System, University of Maryland, NBER Summer Institute, Université de Toulouse, Bank of Spain, Harvard University, Bank of Mexico, Universitat de Girona, and the Federal Reserve Banks of Atlanta, New York, Richmond, Dallas, and Philadelphia, as well as two anonymous referees and two editors for useful comments. We thank Jean Boivin and Mark Watson for pointing out an error in a previous version of the manuscript. Tommaso Monacelli provided valuable research assistance. The second and third authors are grateful to the C. V. Starr Center (Galí and Gertler) and CREI (Galí) for financial support.

1. In subsection III.D, we provide more detail on the role of oil shocks versus monetary policy and also on how the monetary policy rule influences the effects that oil shocks have on the economy.

© 2000 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.
The Quarterly Journal of Economics, February 2000

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

In this paper we explore the role of monetary policy. We first demonstrate that there is a significant difference in the way monetary policy was conducted pre- and post-1979, the year Paul Volcker was appointed Chairman of the Board of Governors of the Federal Reserve System. We then go on to argue that this difference could be an important source of the shift in macroeconomic behavior. In some ways, our story should not be surprising. Many economists agree that monetary policy in the United States has been relatively well managed from the time Paul Volcker took over the helm, through the current regime of Alan Greenspan. It is also generally agreed that monetary policy was not so well managed in the fifteen or so years prior to Volcker.[2] The contribution of our paper is to add precision to this conventional wisdom.

We identify how monetary policy differed before and after Volcker came to office by estimating policy rules for each era. Specifically, we estimate a general type of rule that treats the Federal Funds rate as the instrument of monetary policy. The rule calls for adjustment of the Funds rate to the gaps between *expected* inflation and output and their respective target levels. It is a version of the kind of policy rule that emerges in both positive and normative analyses of central bank behavior that have appeared in recent literature.[3] A distinctive feature of our specification is that it assumes forward-looking behavior on the part of the central bank.

The key difference in the estimated policy rules across time involves the response to expected inflation. We find (not surprisingly) that the Federal Reserve was highly "accommodative" in the pre-Volcker years: on average, it let real short-term interest rates decline as anticipated inflation rose. While it raised nominal rates, it typically did so by less than the increase in expected inflation. On the other hand, during the Volcker-Greenspan era the Federal Reserve adopted a proactive stance toward controlling inflation: it systematically raised real as well as nominal short-term interest rates in response to higher expected inflation. Our results thus lend quantitative support to the popular view that not until Volcker took office did controlling inflation become the organizing focus of monetary policy.

The second part of the paper presents a theoretical model

2. See, e.g., the recent discussions in Friedman and Kuttner [1996] and Gertler [1996].

3. See Clarida, Galí, and Gertler [1999] for a review of the recent literature on monetary policy.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

designed to flesh out how the observed changes in the policy rule could account for the change in macroeconomic performance. We embed policy rules of the type we estimate within a fairly standard business cycle model and then analyze the dynamics of inflation and output in the resulting equilibrium. We show that the estimated rule for the pre-Volcker period permits greater macroeconomic instability than does the Volcker-Greenspan rule. It does so in two distinct respects.

First, the pre-Volcker rule leaves open the possibility of bursts of inflation and output that result from self-fulfilling changes in expectations. These sunspot fluctuations may arise under this rule because individuals (correctly) anticipate that the Federal Reserve will accommodate a rise in expected inflation by letting short-term real interest rates decline (which in turn stimulates the rise in aggregate demand and inflation).[4] On the other hand, self-fulfilling fluctuations cannot occur under the estimated rule for the Volcker-Greenspan era since, within this regime, the Federal Reserve adjusts interest rates sufficiently to stabilize any changes in expected inflation. Second, the pre-Volcker rule is less effective than the Volcker-Greenspan rule at mitigating the impact of fundamental shocks to the economy. That is, holding constant the volatility of exogenous fundamental shocks, the economy exhibits greater stability under the post-1979 rule than under a rule that closely approximates monetary policy pre-1979.

The plan of the paper is as follows. Section II presents our policy rule specification, and discusses the econometric procedure used to estimate it. Section III reports estimates of this rule for different sample periods, conducts a number of robustness checks, and identifies the main differences in the coefficient estimates across periods. Section IV presents the theoretical model: a (now) conventional New Keynesian framework with money, monopolistic competition, and sticky prices. We then present both a qualitative and quantitative analysis of the model under the pre- and post-1979 policy rules. Section V offers concluding remarks. Here we discuss an important issue that the paper raises but does not resolve: why in the pre-Volcker period the Federal Reserve

4. Chari, Christiano, and Eichenbaum [1997] also suggest that the inflation of the 1970s may have been due mainly to self-fulfilling behavior. Their argument exploits the idea that there may be a multiplicity of equilibria in reputational models of monetary policy. Our analysis is based simply on the implications of the estimated historical policy reaction function.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

appeared to pursue a systematic policy rule that not only accommodated inflation, but did so in a way that was entirely predictable by the private sector (at least with the benefit of hindsight).

II. The Federal Reserve's Policy Reaction Function: A Forward-Looking Model

A. A Simple Forward-Looking Rule

We begin with a baseline specification of the policy reaction function. We take as the instrument of monetary policy the Federal Funds rate. Except possibly for a brief period of reserves targeting at the start of the Volcker era, this seems a reasonable choice (see, e.g., Bernanke and Mihov [1998]). Further, Goodfriend [1991] argues that even under the period of official reserves targeting, the Federal Reserve had in mind an implicit target for the Funds rate.

The baseline policy rule we consider takes a simple form. Let r_t^* denote the target rate for the nominal Federal Funds rate in period t. The target rate each period is a function of the gaps between expected inflation and output and their respective target levels. Specifically, we postulate the linear equation:

$$r_t^* = r^* + \beta\,(E[\pi_{t,k}|\Omega_t] - \pi^*) + \gamma\,E[x_{t,q}|\Omega_t], \tag{1}$$

where $\pi_{t,k}$ denotes the percent change in the price level between periods t and $t + k$ (expressed in annual rates). π^* is the target for inflation. $x_{t,q}$ is a measure of the average output gap between period t and $t + q$, with the output gap being defined as the percent deviation between actual GDP and the corresponding target.[5] E is the expectation operator, and Ω_t is the information set at the time the interest rate is set. r^* is, by construction, the desired nominal rate when both inflation and output are at their target levels.

The policy rule given by (1) has some appeal on both theoretical and empirical grounds. Approximate (and in some cases exact) forms of this rule are optimal for a central bank that

5. The flow nature of GDP forces us to be more precise here: $x_{t,q}$ includes GDP generated between the beginning of period t and the beginning of period $t + q$ (i.e., it includes periods—e.g., quarters—t, $t + 1$, . . . and $t + q - 1$). In our empirical work we account for the fact that period t GDP is not known as of the time the interest rate is set in that period; i.e., $x_{t,1} \notin \Omega_t$. This is not true in our theoretical model of Section IV, where all variables dated in period t are determined simultaneously.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

has a quadratic loss function in deviations of inflation and output from their respective targets, given a generic macroeconomic model with nominal price inertia.[6]

On the empirical side a number of authors have emphasized that policy rules like (1) provide reasonably good descriptions of the way major central banks around the world behave, at least in recent years. It is true that the most notable of these papers, Taylor [1993], proposes a rule where the Funds rate responds to lagged inflation and output rather than their expected future values. However, our forward-looking rule nests the Taylor rule as a special case: if either lagged inflation or a linear combination of lagged inflation and the output gap is a sufficient statistic for forecasting future inflation, then equation (1) collapses to the Taylor rule.[7] On the other hand, our forward-looking specification allows the central bank to consider a broad array of information (beyond lagged inflation and output) to form beliefs about the future condition of the economy, a feature that we find highly realistic.

B. Implied Real Rate Rule

The implications of a policy rule like (1) for the cyclical behavior of the economy will of course depend on the sign and magnitude of the slope coefficients, β and γ. To gain the basic intuition, consider the implied rule for the (ex ante) real rate target, rr_t^*:

$$(2) \qquad rr_t^* = rr^* + (\beta - 1)\,(E[\pi_{t,k}|\Omega_t] - \pi^*) + \gamma E[x_{t,q}|\Omega_t],$$

where $rr_t^* \equiv r_t - E[\pi_{t,k}|\,\Omega_t]$ and where $rr^* \equiv r^* - \pi^*$ is the long-run equilibrium real rate.[8] We assume that the real rate is stationary and is determined by nonmonetary factors in the long run,

6. See, e.g., Svensson [1996], and Clarida, Galí, and Gertler [1999].

7. In this case, however, the estimated coefficients of the Taylor rule may be misleading as indicators of the Fed's intended response to inflation and output changes since, in addition to the size of the policy response, they capture the ability of each variable to forecast the state of the economy.

8. Note that rr_t^* is an "approximate" real rate since the forecast horizon for inflation will generally differ from the maturity of the short-term nominal rate used as a monetary policy instrument. In practice, this is of little relevance, given the high correlation among short-term rates at maturities associated with plausible target horizons. Hence, e.g., the correlation between the federal funds rate and the three-months Treasury Bill rate (our baseline horizon) is 0.982 in our sample.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

consistent with conventional wisdom.[9] Accordingly, rr^* is a constant and is independent of monetary policy.[10]

As equation (2) makes clear, the sign of the response of the *real* rate target to changes in expected inflation and the output gap depends on whether β is greater or less than one and on the sign of γ, respectively. Roughly speaking, to the extent that lower real rates stimulate economic activity and inflation (as implied by standard macroeconomic models and as perceived by policymakers and market participants alike), interest rate rules characterized by $\beta > 1$ will tend to be stabilizing, while those with $\beta \leq 1$ are likely to be destabilizing or, at best, accommodative of shocks to the economy.[11] A similar logic applies to the sign of γ (i.e., stabilizing if $\gamma > 0$; destabilizing if $\gamma \leq 0$). We thus have benchmarks ($\beta = 1$, $\gamma = 0$) to evaluate differences in the estimated policy rules across time.

C. Interest Rate Smoothing and Exogenous Shocks

Absent any further modification, the policy reaction function given by equation (1) is too restrictive to describe actual changes in the Funds rate. There are at least three reasons why. First, the specification assumes an immediate adjustment of the actual Funds rate to its target level, and thus ignores the Federal Reserve's tendency to smooth changes in interest rates.[12] Second, it treats all changes in interest rates over time as reflecting the Federal Reserve's systematic response to economic conditions. Specifically, it does not allow for any randomness in policy actions,

9. In our empirical work, however, we allow for the possibility that there have been shifts over time in the long-run equilibrium real rate, since we assume only that rr^* is constant within the subperiod over which we estimate.

10. In our empirical work we allow for changes in the equilibrium real rate across subperiods.

11. Another way to model accommodation of inflation would be to allow for endogenous adjustment of the target inflation rate, π^*. We view this approach as a sympathetic alternative to the one we take of focusing on whether the slope coefficient β is less than unity. We opt to treat π^* as a constant and instead let β characterize the degree of accommodation for two reasons. First, this approach has the advantage of parsimony—there is no need to model the adjustment of π^*. Second, as a matter of logic, it is a priori reasonable to treat π^* as constant since it is meant to reflect an optimum for inflation that is independent of current economic conditions. In our robustness exercises, however, we do let π^* vary within each subsample. To foreshadow, we find that the magnitude of β seems to do a better job of capturing the degree of accommodation than any time variation in π^*, consistent with the approach we take.

12. See also Rudebusch [1995] for evidence on the serial correlation of interest rate changes. Why this smoothing occurs is beyond the scope of this paper, although a number of explanations are found in the literature, including fear of disruption of financial markets [Goodfriend 1991], or uncertainty about the effects of interest rate changes [Sack 1997].

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

other than that associated with misforecasts of the economy. Third, it assumes that the Federal Reserve has perfect control over interest rates; i.e., it succeeds in keeping them at the desired level (e.g., through necessary open market operations).

We relax the first assumption by extending the model in a straightforward way. In particular, we specify the following relationship for the *actual* Funds rate, r_t:

$$r_t = \rho(L)\, r_{t-1} + (1 - \rho)\, r_t^*, \tag{3}$$

where $\rho(L) = \rho_1 + \rho_2 L + \ldots + \rho_n L^{n-1}$, and where $\rho \equiv \rho(1)$. v_t is a zero mean exogenous interest rate shock, and the Funds rate target r_t^* is given by (1). Equation (3) postulates partial adjustment of the Funds rate to the target r_t^*. Specifically, each period the Federal Reserve adjusts the Funds rate to eliminate a fraction $(1 - \rho)$ of the gap between its current target level and some linear combination of its past values. We interpret ρ as an indicator of the degree of smoothing of interest rate changes.

Combining the partial adjustment equation (3) with the target model (1) yields the policy reaction function,

$$r_t = (1 - \rho)\,[rr^* - (\beta - 1)\pi^* + \beta\pi_{t,k} + \gamma x_{t,q}] + \rho(L)\, r_{t-1} + \epsilon_t, \tag{4}$$

where $\epsilon_t \equiv -(1 - \rho)\,\{\beta(\pi_{t,k} - E[\pi_{t,k}|\Omega_t]) + \gamma(x_{t,q} - E[x_{t,q}|\Omega_t])\}$. Notice that the term in curly brackets is a linear combination of forecast errors and is thus orthogonal to any variable in the information set Ω_t.

Let $\mathbf{z}_t$ denote a vector of instruments known when r_t is set (i.e., $\mathbf{z}_t \in \Omega_t$). Equation (4) then implies the set of orthogonality conditions,

$$E\{[r_t - (1 - \rho)\,(rr^* - (\beta - 1)\pi^* + \beta\pi_{t,k} + \gamma x_{t,q}) + \rho(L)\, r_{t-1}]\, \mathbf{z}_t\} = 0, \tag{5}$$

which provide the basis for the estimation of the parameter vector $(\alpha,\beta,\gamma,\rho)$, using the Generalized Method of Moments [Hansen 1982], with an optimal weighting matrix that accounts for possible serial correlation in $\{\epsilon_t\}$.[13] To the extent that the dimension of vector $\mathbf{z}_t$ exceeds four—the number of parameters being estimated—(5) implies some overidentifying restrictions that we can

13. Note that, by construction, the first component of $\{\epsilon_t\}$ follows an $MA(a)$ process, with $a = \max[k,q] - 1$, and will thus be serially correlated (unless $k = q = 1$).

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

test in order to assess the validity of our specification as well as the set of instruments used.

In the absence of further assumptions our approach identifies only the term $rr^* - (\beta - 1)\pi^*$, but not rr^* or π^* separately. Since target inflation π^* is of some interest in our characterization of monetary policy, we impose an additional restriction that permits us to identify and estimate this parameter. Specifically, we take the observed sample average as a measure of the equilibrium real rate rr^*, an assumption that we view as providing a reasonable first approximation, given our sample size. Imposing this restriction directly in equation (5) allows us to estimate π^* jointly with the parameter vector $(\alpha,\beta,\gamma,\rho)$.

Before proceeding, we briefly address several econometric issues. First, our empirical analysis maintains the assumption that both inflation and the nominal interest rate are stationary. We view this assumption as reasonable for the postwar United States, even though the null of a unit root in either variable is often hard to reject at conventional significance levels, given the persistence of both series and the well-known low power of unit root tests. In addition to its empirical plausibility, stationarity of both inflation and the nominal interest rate is also a property of many of the theoretical models that rationalize the use of the kind of policy rule considered here.[14] In our robustness analysis, however, we do allow for the possibility of drift in the trend rate of inflation by letting the target inflation rate vary across the regimes of different Federal Reserve chairmen.

Second, the sample period must contain sufficient variation in inflation and output and must be sufficiently long in order to identify the slope coefficients in the policy reaction function, as well as the target inflation rate π^*. In particular, estimating the rule over a short sample with little variability in inflation can yield highly misleading results. Suppose, for example, that the Federal Reserve responds aggressively to large deviations of inflation from target but not to small deviations. Then by estimating over a period where inflation does not vary much from its target, one might mistakenly conclude that the Fed is not aggressive in fighting inflation (i.e., one might mistakenly obtain too low an estimate of β). Alternatively, suppose that a central bank confronting high inflation is in the process of raising rates to engineer a disinflation. By estimating the rule only over the

14. See, e.g., Clarida, Gali, and Gertler [1999].

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

current period of high inflation, one might mistakenly conclude that the central bank has a high target inflation rate (i.e., one might mistakenly obtain too high an estimate of π^* and, as a result, too low an estimate of β).[15] The subsamples we consider, however, appear to contain sufficient variation in inflation relative to the sample mean and are sufficiently long to permit us to correctly identify both the slope coefficient on inflation, β, and the target inflation rate π^*.[16]

III. The Federal Reserve's Policy Reaction Function: The Evidence

In this section we report estimates of the policy reaction function defined by equations (1) and (3). We accomplish two main objectives. First, we demonstrate the existence of a systematic relationship between the Funds rate and forecasts of future inflation and output along the lines suggested by our model. Second, we identify differences in the conduct of monetary policy pre- and post-1979. We do so by estimating monetary policy rules for each era and performing tests of structural stability across periods.

The data are quarterly time series spanning the period 1960:1–1996:4. With one exception, we obtain the data from CITIBASE (mnemonics follow in parentheses). We use as the interest rate the average Federal Funds rate (FYFF) in the first-month of each quarter, expressed in annual rates. The baseline inflation measure is the (annualized) rate of change of the GDP deflator (GDPP) between two subsequent quarters. But we also report results using CPI (PUNEW) inflation. The baseline "output gap" measure is the series constructed by the Congressional Budget Office (CBO). We also use two alternative measures (described below) based on the detrended series for GDP (GDPQ) and the unemployment rate (LHUR). The instrument set includes lags of the Funds rate, inflation, and the output gap, as well as the

15. The estimate of π^* will depend heavily on the sample mean of inflation. Thus, estimating over a short period of above target inflation can generate too high an estimate of π^*. The overestimate of π^* will be accompanied by an underestimate of β, since the former will account for the high nominal rates over the short period. By lengthening the sample to include the full disinflation and beyond, one will obtain an estimate of π^* that is closer to the true target, and as a consequence a (higher) estimate of β that is closer to its true value.

16. When we consider subsample stability of our estimates in Section III, we address the short sample problem by restricting some coefficients to be constant across the sample (in the cases where we cannot reject that coefficient is stable).

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE I
AGGREGATE VOLATILITY INDICATORS

	Standard Deviation of:			
	Inflation		Output	
	Level	*hp*	*Gap*	*hp*
Pre-Volcker	2.77	1.48	2.71	1.83
Volcker-Greenspan	2.18	0.96	2.36	1.49
post–82	1.00	0.79	2.06	1.34

same number of lags of commodity price inflation (PSCCOM), M2 growth (FM2), and the "spread" between the long-term bond rate (FYGL) and the three-month Treasury Bill rate (FYGM3).[17]

We divide the sample into two main subperiods. The first (60:1–79:2), encompasses the tenures of William M. Martin, Arthur Burns, and G. William Miller as Federal Reserve chairmen. The second (79:3–96:4) corresponds to the terms of Paul Volcker and Alan Greenspan. As we discussed in the Introduction, these subperiods roughly correspond to the unstable and stable eras of recent macroeconomic history. This characterization, while simplistic, is clearly reflected in the data. Table I reports the standard deviation of inflation (levels and HP-detrended) and output (CBO gap and HP-detrended), for the two subperiods. The reduction in volatility appears substantial for each variable. Not surprisingly, the decline is more dramatic when we begin the second subperiod in 82:4, after Volcker disinflation, as the bottom row of the table indicates.

A. Baseline Estimates

Table II reports GMM estimates of the interest rate rule parameters π^*, β, γ, and ρ for each sample period, using the CBO output gap and GDP deflator inflation (our baseline variables). The target horizon is assumed to be one-quarter for both inflation and the output gap (i.e., $k = q = 1$). Standard errors are reported in parentheses. The right-most column reports the *p*-value associated with a test of the model's overidentifying restrictions (Han-

17. In closely related work Orphanides [1997] estimates a reaction function using more direct measures of the Fed's perception of both the output gap and inflation, based on real time data. His results, by and large, confirm the results we obtain.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE II
BASELINE ESTIMATES

	π^*	β	γ	ρ	p
Pre-Volcker	4.24	0.83	0.27	0.68	0.834
	(1.09)	(0.07)	(0.08)	(0.05)	
Volcker-Greenspan	3.58	2.15	0.93	0.79	0.316
	(0.50)	(0.40)	(0.42)	(0.04)	

Standard errors are reported in parentheses. The set of instruments includes four lags of inflation: output gap, the federal funds rate, the short-long spread, and commodity price inflation.

sen's *J*-test). The estimates are based on a specification of the interest rate rule with two lags of the interest rate ($n = 2$), which seemed to be sufficient to eliminate any serial correlation in the error term. Four lags of the instruments were used.

A number of interesting results stand out. Note first that the model is not rejected at conventional significance levels for any of the specifications or sample periods. The estimates of β and γ, further, generally have the expected sign and are significant in most cases. These estimates also point to substantial differences in the policy reaction function across periods. Most importantly, the estimate of β, the coefficient associated with expected inflation, is significantly below unity for the pre-Volcker period (0.83, with *s.e.* = 0.07), and far greater than one for the Volcker-Greenspan period (2.15, with *s.e.* = 0.40). On the other hand, the estimates of γ—the coefficient measuring the sensitivity to the cyclical variable—are also significant in both periods, but only marginally so for the Volcker-Greenspan era.[18]

The estimates of the inflation target, π^*, seem quite plausible in all cases: roughly four and a quarter percent pre-Volcker and three and half percent post-Volcker. While the point estimates reveal a slight downward trend over time, the difference is not significant. Based on this result and the estimates of α and β across subsamples, it does not seem to be the case that differences in monetary policy pre- and post-1979 simply reflect differences in the target inflation rate. We shortly present some more evidence that bears on this issue that stems from an analysis of (within) subsample stability.

Finally, the estimate of the smoothing parameter ρ is high in all cases, suggesting considerable interest rate inertia: only between 10 and 30 percent of a change in the interest rate target

18. See below for further discussion.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

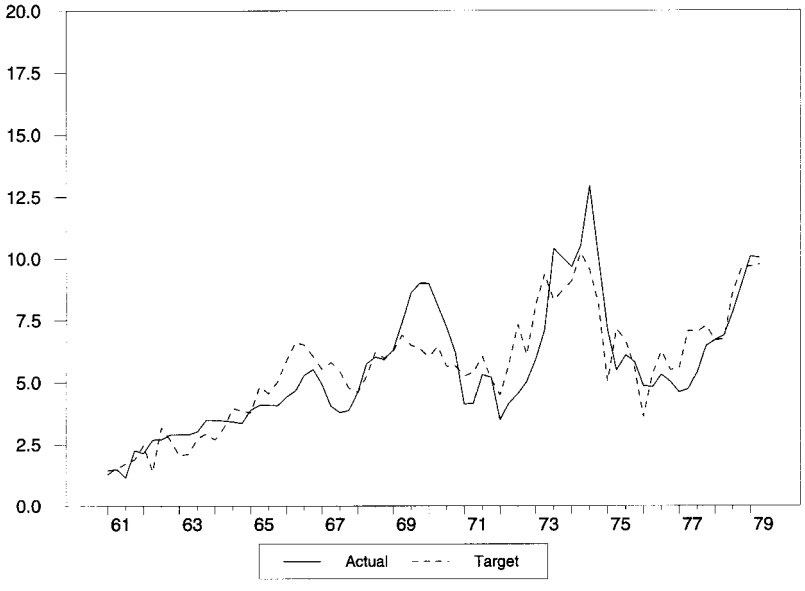


FIGURE I
Actual versus Target Rates:
Pre-Volcker Era

is reflected in the Funds rate within the quarter of the change. Thus, our estimates confirm the conventional wisdom that the Federal Reserve smooths adjustments in the interest rate.

To illustrate how well the model characterizes the behavior of the Funds rate, Figures I and II present the target rate estimates for each subperiod relative to the actual values of the Funds rate, using our baseline estimates. In each subperiod, the target rate captures the broad swings in the actual rate reasonably well. Interestingly, during the 1987–1992 period that Taylor [1993] analyzes, our target rate tracks the actual rate about as well as does the simple Taylor rule.[19]

B. Robustness Analysis

We next explore the robustness of our results along a number of dimensions. We consider (1) alternative measures of inflation

19. We stress that we are comparing the actual rate with the implied target rate, as opposed to the fitted model, which allows for partial adjustment. The fitted model, of course, would track the actual rate even more closely than does the target rate.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

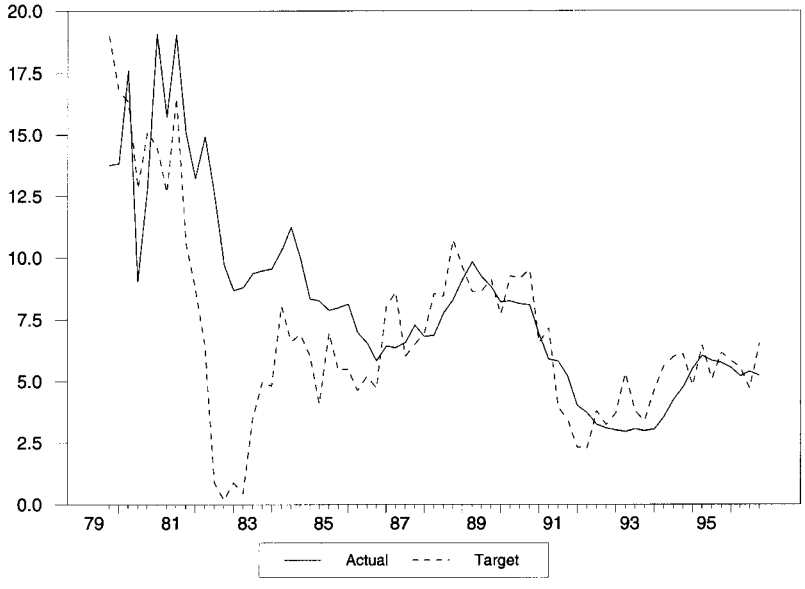


FIGURE II
Actual versus Target Rates: Volcker-Greenspan Era

and the output gap, (2) alternative target horizons for each variable, (3) parameter stability within subsamples, and (4) a backward-looking variation of our policy reaction function. We demonstrate that in each instance, the insights from the baseline case remain intact.

1. Alternative Measures. We first reestimate the reaction function using different measures of the output gap and inflation. We consider two alternative measures of the output gap: (a) the deviation of (log) GDP from a fitted quadratic function of time; and (b) the deviation of the unemployment rate from a similar time trend, with the sign of the resulting series switched.[20] Finally, we consider one alternative measure of inflation: the rate of change of the consumer price index (CPI).

Table III reports the estimates for the two main subperiods. The key results from the baseline case are robust to the use of alternative output gap and inflation measures. In fact, both the signs and magnitudes of the estimated parameters remain largely unchanged. There is, however, one minor difference: the esti-

20. We switch the sign of the series in order to preserve the sign interpretation for parameter γ.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE III
ALTERNATIVE VARIABLES

	π^*	β	γ	ρ	p
Detrended output					
Pre-Volcker	4.17	0.75	0.29	0.67	0.801
	(0.68)	(0.07)	(0.08)	(0.05)	
Volcker-Greenspan	4.52	1.97	0.55	0.76	0.289
	(0.58)	(0.32)	(0.30)	(0.05)	
Unemployment rate					
Pre-Volcker	3.80	0.84	0.60	0.63	0.635
	(0.87)	(0.05)	(0.11)	(0.04)	
Volcker-Greenspan	4.42	2.01	0.56	0.73	0.308
	(0.44)	(0.28)	(0.41)	(0.05)	
CPI					
Pre-Volcker	4.56	0.68	0.28	0.65	0.431
	(0.53)	(0.06)	(0.08)	(0.05)	
Volcker-Greenspan	3.47	2.14	1.49	0.88	0.138
	(0.79)	(0.52)	(0.87)	(0.03)	

Standard errors are reported in parentheses. The set of instruments includes four lags of inflation, output gap, the federal funds rate, the short-long spread, and commodity price inflation.

mated output gap coefficient for the Volcker-Greenspan subperiod, although positive, is now insignificant under the three specifications. At the same time, there remains a striking difference in the estimated slope coefficient for inflation across subperiods: less than one before Volcker, greater than one under Volcker-Greenspan, with point estimates very close to those obtained in the baseline case.

2. Alternative Horizons. In the baseline case we assume that the Federal Reserve looks ahead one quarter for both inflation and the output gap. We now consider allowing for alternative—and, in our opinion, more realistic—target horizons for the same variables. Table IV reports results for ($k = 4$, $q = 1$) as well as ($k = 4$, $q = 2$); i.e., the Fed is assumed to have a target horizon of one year for its inflation target and of one (or two) quarters for the output. One formal rationale is that these horizons are roughly in line with the conventional wisdom regarding the lag with which monetary policy affects either variable (e.g., Bernanke and Mihov [1998]). In addition, these values are roughly consistent with informal discussions of policy tactics by Federal Reserve officials. In either case, the results are qualitatively very similar to those reported in Table II.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE IV
ALTERNATIVE HORIZONS

	π^*	β	γ	ρ	p
$k = 4, q = 1$					
Pre-Volcker	3.58	0.86	0.34	0.73	0.835
	(1.42)	(0.05)	(0.08)	(0.04)	
Volcker-Greenspan	3.25	2.62	0.83	0.78	0.876
	(0.23)	(0.31)	(0.28)	(0.03)	
$k = 4, q = 2$					
Pre-Volcker	3.32	0.88	0.34	0.73	0.833
	(1.80)	(0.06)	(0.09)	(0.04)	
Volcker-Greenspan	3.21	2.73	0.92	0.78	0.886
	(0.21)	(0.34)	(0.31)	(0.03)	

Standard errors are reported in parentheses. The set of instruments includes four lags of inflation, output gap, the federal funds rate, the short-long spread, and commodity price inflation.

3. Subsample Stability. We next explore the stability of parameters within each subsample. Among other things, this exercise permits us to relax the assumption that the inflation target π^* is constant within the estimation period. It is conceivable, for example, that there was an upward shift in the target during the period of rising inflation in the 1970s.

A simple and natural way to proceed is to assume that the policy reaction function is stable during the tenure of the Federal Reserve chairman in charge at the time, but may vary across Chairmen. If we tack the brief period of Miller (78:1–79:3) on to Burns (70:1–78:1), then each subperiod may be divided into two regimes of roughly equal length. For the pre-1979 sample we thus have Martin (60:1–69:4) and Burns and Miller (70:1–79:2); for the post-1979 sample Volcker (79:3–87.2) and Greenspan (87.3–96:4).

As discussed earlier, estimating the policy rule over short samples can generate imprecise estimates, given the limited number of observations. Accordingly, we adopt the following procedure: we first estimate the reaction function for each baseline period (pre-Volcker or Volcker and Greenspan), but allow for a shift across Chairmen in each of the coefficients (by means of appropriate dummies). Second, we reestimate the rule after constraining all the parameters for which the shift was found to be insignificant in the first stage to be constant across Chairmen, while allowing for changes in the remaining parameters. The resulting estimates are reported in Table V. We present estimates

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE V
SUBSAMPLE STABILITY

	π^*	β	γ	ρ	p
Martin					
(1,1)	5.16	0.86	0.14	0.77	0.524
	(1.72)	(0.08)	(0.16)	(0.06)	
(4,1)	7.15	0.92	0.06	0.72	0.719
	(5.55)	(0.08)	(0.07)	(0.05)	
Burns-Miller					
(1,1)	5.16	0.86	0.78	0.69	0.524
	(1.72)	(0.08)	(0.18)	(0.04)	
(4,1)	7.15	0.92	1.24	0.80	0.719
	(5.55)	(0.08)	(0.39)	(0.05)	
Volcker					
(1,1)	3.75	2.02	−0.02	0.63	0.612
	(0.28)	(0.23)	(0.15)	(0.04)	
(4,1)	2.45	2.38	0.68	0.74	0.804
	(0.47)	(0.35)	(0.30)	(0.04)	
Greenspan					
(1,1)	3.75	2.02	0.99	0.63	0.612
	(0.28)	(0.23)	(0.18)	(0.04)	
(4,1)	2.45	2.38	0.68	0.91	0.804
	(0.47)	(0.35)	(0.30)	(0.02)	
Post-82					
(1,1)	3.43	1.58	0.14	0.91	0.416
	(1.24)	(0.72)	(0.42)	(0.03)	
(4,1)	3.16	3.13	0.09	0.82	0.894
	(0.10)	(0.33)	(0.15)	(0.02)	

Standard errors are reported in parentheses. The set of instruments includes two lags of inflation, output gap, the federal funds rate, the short-long spread, and commodity price inflation, as well as the same variables with a multiplicative subperiod dummy.

for two different target horizons for inflation and the output gap: (1,1) and (4,1).[21]

Consider first the pre-Volcker period. Interestingly, no significant difference arises across Chairmen in either the value of the inflation target π^*, or in the inflation coefficient β. The point estimates for the inflation target (in the 5–7 percent range) are somewhat above the baseline estimates for the full pre-Volcker sample, although not significantly since the standard errors are now rather large. The estimated value for the inflation coefficient is below unity and is in line with the baseline estimates. The only

21. Since we also dummy all the instruments, we only use only two instrument lags in our subsample stability analysis, thus keeping the total number of instruments (and the degrees of freedom) comparable to the other specifications.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

notable departure from the baseline case is that the output gap coefficient is insignificant under Martin's tenure, but positive and significant under Burns and Miller.

Our procedure detects few robust differences across the Volcker and Greenspan eras. As in the pre-Volcker era, the estimates π^* and β are stable across chairmen. Estimates for each parameter, further, are close to those obtained in the baseline case. With a one-period horizon for both inflation and the output gap, the coefficient on the latter is close to zero (and insignificant) under Volcker, but is positive and significant under Greenspan. On the other hand, when the target horizon for inflation is four quarters, this difference vanishes. The estimated common value is positive and significant. In addition, the estimates under the (4,1) horizon point to a significant increase in the degree of smoothing of interest rate changes under Greenspan, as reflected in the significantly higher estimate of the parameter ρ.

As a final check on subsample stability, we explore the effects of removing the first three years of the Volcker era from the entire Volcker-Greenspan sample. There are at least two reasons for doing this. First, this period was characterized by a sharp, one-shot disinflation episode, which brought inflation down from roughly 10 percent in 1980 to 4 percent in 1983, a level around which it stabilized. Second, over the period 1979:4–1982:4, the operating procedures of the Federal Reserve involved targeting nonborrowed reserves as opposed to the Federal Funds rate.[22] In order to make sure that neither of those features are overinfluencing our estimates for the Volcker-Greenspan period, we reestimate the reaction function for the period 82:4–96:4, thus excluding the period of nonborrowed reserves targeting. The corresponding results, for two alternative horizons, are shown in the bottom panel of Table V.

Once again, the estimates of β—the coefficient associated with expected inflation—are above unity and not statistically

22. As we noted earlier, Goodfriend [1991] argues forcefully that the Federal Reserve did indeed choose targets for nonborrowed reserves over this period ultimately with an objective for the path of the interest rate in mind. Nevertheless, Bernanke and Mihov [1998] present evidence that over the 1979:10–1982 period nonborrowed reserves was the operating instrument of monetary policy, which accords with conventional wisdom. For the rest of the time they show that it is reasonable to treat the Federal Funds rate as the instrument of monetary policy. In a companion paper [Clarida, Gali, and Gertler 1997] we show that our baseline specification is robust to allowing for the possibility that the Fed may respond to money growth independently of its predictive power for inflation. That is, we reject the hypothesis that the Fed was targeting money growth. Our results are thus consistent with Friedman and Kuttner [1996].

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE VI
BACKWARD-LOOKING ESTIMATES

	π^*	β	γ	ρ	p
Pre-Volcker	5.95	0.86	0.39	0.68	0.590
	(1.92)	(0.07)	(0.08)	(0.05)	
Volcker-Greenspan	4.08	1.72	0.34	0.71	0.307
	(0.56)	(0.28)	(0.19)	(0.05)	
Post-82	2.96	2.55	−0.15	0.89	0.486
	(0.27)	(0.56)	(0.28)	(0.03)	

Standard errors are reported in parentheses. The set of instruments includes four lags of inflation, output gap, the federal funds rate, the short-long spread, and commodity price inflation.

different from the baseline case (although the estimate for the (1,1) horizon is noisy).[23] On the other hand, the estimates of γ—the coefficient measuring the sensitivity to the cyclical variable—all become very small and insignificant when we use post-82 data (in contrast with our results for the full Volcker-Greenspan period). Thus, here we cannot reject the hypothesis that the Fed has effectively pursued a "pure inflation targeting" policy.

In sum, the key insights obtained in the baseline case are robust to allowing for structural changes across Chairmen. In particular, the striking difference in the reaction function across time is the rise in the slope coefficient on inflation from slightly less than unity pre-Volcker to around two in the Volcker-Greenspan era.

4. Backward-looking estimates. We complete our robustness analysis by reporting the estimates for different subperiods of a backward-looking rule of the sort considered by Taylor [1993]. Our version of that rule corresponds to specification (4) with both k and q set to -1. Table VI reports the corresponding results. All the qualitative features of our baseline specification estimates seem to hold here as well, suggesting that they are not inherent to the forward-looking specification of the interest rate rule. In sum, while we view the forward-looking specification as more plausible a priori, our key insights also obtain from the backward-looking specification.

C. Discussion of the Empirical Results

Overall, our estimates point to the existence of important differences across periods in the sensitivity of monetary policy.

23. Notice that the null hypothesis of a unit inflation coefficient cannot be rejected in this case.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

This result, further, appears to be a robust feature of the data. Specifically, during the pre-Volcker period, in response to forecastable inflationary pressures, the Federal Reserve tended to let real interest rates decline or, at best, did not try to raise them. In other words, while the central bank did raise nominal rates, it did not do so sufficiently to raise real rates.[24] This kind of response clearly does not stabilize inflation under any plausible view of the linkages between real rates, aggregate demand, output, and inflation.[25] Thus, according to our results, the persistent and volatile behavior of inflation in the pre-Volcker era may be partly due to the monetary rule in place, independently of the nature of the fundamental shocks that may have impinged on the economy during that period.

By way of contrast, under the Volcker-Greenspan regime, the Federal Reserve has substantially raised target real rates in the wake of an anticipated increase in inflation (on a two-for-one basis, according to a rough average of our point estimates). To the extent that a rise in the real rate slows down the level of economic activity and relieves inflationary pressures, the interest rate policy in the Volcker-Greenspan era provides a natural explanation for the stability of inflation experienced by the U. S. economy in recent years.

D. Oil Shocks versus Monetary Policy

Our analysis has emphasized how differences in monetary policy might account for differences in macroeconomic behavior pre- and post-1979. An alternative (although not incompatible) hypothesis points to the role of the two major oil shocks that occurred in 1973 and in 1979. Indeed, Hamilton [1983] has argued forcefully that oil shocks are a central driving force in business cycles. This raises the question of whether the pattern of oil shocks alone could account for the shift in macroeconomic volatility.

24. Our finding of a less than one-for-one adjustment of the nominal rate to changes in expected inflation is closely related to the finding of a strong negative correlation between the estimated expected inflation rate and the estimated real rate by Mishkin [1981] and others in the context of an exploration of the Fisher hypothesis. Interestingly, the sample period in Mishkin's paper ends in 1979:4, just one quarter after Volcker began his tenure as Fed chairman! Furthermore, in subsequent work Huizinga and Mishkin [1986] show formally that there is a shift in interest rate behavior before and after October 1979.

25. As we noted at the end of Section II, if the Fed responds aggressively only to large deviations of inflation from target but not to small deviations, then it may be possible to obtain an estimate of β that is less than unity, even though the policy regime is geared toward stabilizing inflation. However, for this scenario to be plausible, we should also observe little variation of inflation about target. This was clearly not the case in the Martin-Burns-Miller sample.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

For two distinct reasons, however, we believe that differences in monetary policy remain critical to the story. First, recent work by De Long [1997]; Bernanke, Gertler, and Watson [1997] (BGW); Barsky and Killian [1998] (BK); and others questions the degree to which oil shocks can account for the events of the 1970s. Second, and perhaps more important, even if one accepts the view that the oil shocks were critical to the two major recessions, it would be difficult for these shocks to generate the kind of persistent inflation that arose in this era in the absence of an accommodating monetary policy. We briefly elaborate on each point.

At a minimum, the notion that the oil shocks can largely account for the volatile behavior in output during the 1970s period is open to debate. While a major shock to the real price of oil preceded both the 1974–1975 and 1980–1982 recessions, it is also true that significant tightening of monetary policy preceded each of these downturns (see, e.g., BGW). The bottom line is that it is a nontrivial exercise to sort out the relative importance of the oil shocks versus monetary policy over this period. The original Hamilton [1983] evidence, which is based on a bivariate analysis of oil prices and output, does not confront this identification problem. Multivariate analyses, which among other things allows for an indicator of monetary policy, point to a significantly weaker role for oil shocks in output dynamics (see, e.g., BGW and the references therein). BGW, further, present evidence to suggest that the endogenous tightening of monetary policy in response to the oil price increases accounts for much of the subsequent downturn in output, as opposed to the direct impact of the oil shocks themselves.[26]

While there is room to debate the importance of the oil shocks for real activity over this period, the case that these shocks alone account for the sustained high inflation is completely unpersuasive. Note first that the Hamilton evidence is silent on the link between the oil shocks and inflation. On the other hand, as De Long [1997] emphasizes, timing considerations make the oil

26. Along these lines, both BGW and BK present evidence to suggest that the tightening of monetary policy in response to the nonoil commodity price shocks that preceded the oil shocks had a much greater role in the 1974–1975 recession than did the oil shocks. BK argue further that easy monetary policy during 1972 induced the run-up in commodity prices. Finally, the long lag between the 1979 oil shock and the 1981–1982 recession makes questionable whether the former was central to the latter, especially given the sustained tightening of monetary policy over this period. Further, the Carter credit controls appeared to have been at work during the brief downturn in 1980.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

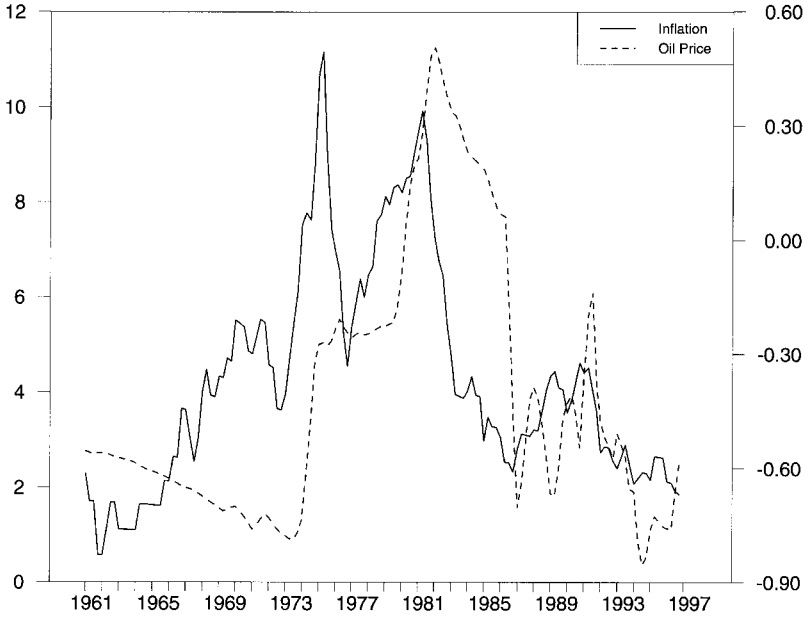


FIGURE III
Real Oil Price and Inflation

shocks suspect as the leading explanation for inflation over this period. Figure III illustrates this point. The figure shows three-quarter centered moving averages of the real oil price against inflation over the period 1960:1–1997:1. As De Long argues, the initial build-up of inflation in the late 1960s and early 1970s occurs prior to the first oil shock. Indeed, until the time of the first oil shock in 1974, the real oil price is steadily declining, while inflation is steadily rising. The real oil price then remains constant until late 1979, the period of the second major oil shock. Note again that there is a steady rise in inflation over a three-year period prior to this shock. In turn, the decline in inflation in the early 1980s appears to lead the decline in the real oil price. Formal evidence in BGW and BK supports this descriptive evidence. The same multivariate analysis that suggests a modest effect of oil prices on output indicates an even weaker effect on inflation.[27]

27. In a multivariate system that includes real GDP, the GDP deflator, nonoil commodity prices, the Federal Funds rate, and Hamilton's preferred index of oil shocks (the minimum of zero and the current oil price minus the largest price over the previous four quarters), the oil shock measure accounts for at best roughly 10

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

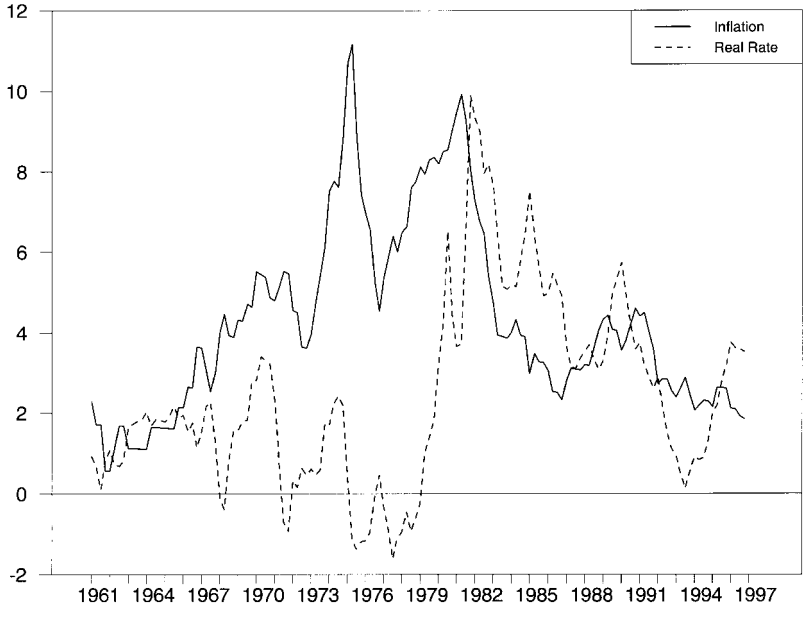


FIGURE IV
The Real Interest Rate and Inflation

Even assuming that the oil shocks played a critical role, the ultimate impact of these disturbances on output and inflation depends very much on the feedback monetary policy rule that is in place. It is hard to imagine, for example, that the 1973 oil shock alone could have generated high inflation up to the time of the second shock in 1979, in the absence of an accommodating monetary policy. We demonstrate this point explicitly in the next section. In particular, we show in the context of a small macro model that a supply shock can indeed induce persistent inflation under the estimated pre-Volcker rule, but not under the estimated Volcker-Greenspan rule. In this respect, the nature of the monetary policy rule must undoubtedly be a critical factor in the 1970s period of stagflation.

Finally, Figure IV plots the behavior of the ex post real interest rate versus inflation (again using three-quarter centered moving averages). Note that pre-1979 the real rate steadily

percent of the variation in the GDP deflator over the period 1960:1 to 1984:4. Extending the sample to 1997:4 points to an even weaker effect of the oil shock.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

declines as inflation rises. Conversely, the real rate rises sharply in late 1979, leading the subsequent decline in inflation. This picture is thus consistent with our story, which emphasizes the switch in monetary policy from accommodating to combating inflation.[28] In the next section we flesh out the implications of the differences in monetary policy pre- and post-Volcker, using a simple macroeconomic framework.

IV. Interest Rate Rules and Economic Fluctuations

In this section we analyze some of the macroeconomic implications of the estimated monetary policy reaction functions. We do so in the context of a monetary business cycle model with sticky prices. We first present the model's equilibrium conditions, and then analyze how the properties depend on the monetary policy rule in place. A comprehensive analysis of the quantitative properties of the model is beyond the scope of the present paper. Instead, we choose to focus our attention on a specific, but (in our opinion) rather important and fascinating issue, namely, the extent to which the change in the systematic component of monetary policy across the pre-Volcker and the Volcker-Greenspan eras may explain some of the differences in the degree of macroeconomic instability, as reflected in the volatility measures reported in Table I.

A. A Baseline Model

Our baseline model is a version of the sticky price models found in King and Wolman [1996], Woodford [1996, 1998], and Yun [1996], among others. After log-linearization around a zero inflation steady state, the model's equilibrium conditions are summarized by the following equations (ignoring uninteresting constants):

$$\pi_t = \delta E\{\pi_{t+1}|\Omega_t\} + \lambda\,(y_t - z_t) \tag{6}$$

$$y_t = E[y_{t+1}|\Omega_t] - (1/\sigma)\,(r_t - E[\pi_{t+1}|\Omega_t]) + g_t \tag{7}$$

$$r_t^* = \beta E[\pi_{t+1}|\Omega_t] + \gamma x_t \tag{8}$$

$$r_t = \rho r_{t-1} + (1 - \rho)\, r_t^*. \tag{9}$$

28. Note that the monetary policy tightening prior to the 1974–1975 recession is quickly reversed, in contrast to the tightening prior to the 1980–1982 downturn. Thus, although policy tightening may have played a role in the 1974–1975 recession, the overall pattern pre-1979 is consistent with accommodation of inflation, as our formal evidence suggests.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

Equation (6) describes the change in the aggregate price level as a function of expected future inflation and the deviation of (log) output y_t from its natural rate z_t, where the latter is defined as the level of output that would obtain under fully flexible prices.[29] It can be derived from the aggregation of optimal price-setting decisions by monopolistically competitive firms, in an environment in which each firm adjusts its price with a constant probability in any given period.[30] Equation (7) combines a standard Euler equation for consumption with a market clearing condition.[31] It is often interpreted as an IS schedule, determining the current output gap as a function of the ex ante real rate and expected future output. In this context g_t can be interpreted as an exogenous demand factor. We assume that both z_t and g_t follow a stationary AR(1) process. Equations (8) and (9) specify the policy rule. They are the theoretical model's counterpart to (1) and (3). For simplicity, we restrict ourselves to the case of a one-period horizon ($k = q = 1$), and assume that all variables dated t or earlier belong to information set Ω_t.

Our objective here is to simulate the model under alternative policy rules. We use standard methods to solve for the equilibrium dynamics (see, e.g., Blanchard and Kahn [1980]). We then choose the nonpolicy parameters as follows. We set the quarterly discount factor equal to 0.99, implying an annual risk-free rate of 4 percent. We set σ, the coefficient of relative risk aversion, to be equal to 1, and λ, the output elasticity of inflation, equal to 0.3,[32] and the autoregressive coefficient of the z_t and g_t processes equal to 0.9. We choose the policy parameters, β, γ, and ρ, to correspond to our estimated policy-reaction functions, depending on the regime we wish to analyze.

We first explore how the estimated rule for the pre-Volcker era opens up the possibility of self-fulfilling fluctuations. We then

29. Galí and Gertler [1999] and Sbordone [1998] present some evidence suggesting that (6) provides a good first approximation to the dynamics of inflation in the United States.

30. Such a price-setting structure was first introduced in Calvo [1983], and has been frequently adopted in macroeconomic applications as a simple, flexible way of introducing price stickiness. A similar forward-looking Phillips curve arises, however, under alternative price-setting assumptions (e.g., quadratic adjustment costs or deterministic time-dependent rules with staggered pricing).

31. It assumes a CRRA period-utility with relative risk aversion parameter σ.

32. There is no widespread consensus on the value of λ. Values found in the literature range from 0.05 [Taylor 1980] to 1.22 [Chari, Kehoe, and McGrattan 1996]. Following Woodford [1996], we choose the intermediate value 0.30, which is consistent with the empirical findings in Roberts [1995].

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

compare how the estimated pre-Volcker and Volcker-Greenspan rules affect the economy's response to fundamental disturbances.

B. Interest Rate Rules and Endogenous Fluctuations

As emphasized by Kerr and King [1996]; Bernanke and Woodford [1997]; and Clarida, Galí, and Gertler [1997], the policy feedback rule itself may be a source of instability if the coefficient on the inflation gap, β, is below unity. Values of β in this range lead to indeterminacy of the equilibrium, and raise the possibility of fluctuations in output and inflation around their steady state values that result from self-fulfilling revisions in expectations.[33] The intuition is straightforward: with β below unity, a rise in anticipated inflation leads to a decline in the real interest rate. The decline in the real rate then stimulates aggregate demand which, in turn, induces a rise in inflation. The initial rise in expected inflation thus becomes self-confirmed.

Since the estimate of β is consistently less than unity for the pre-Volcker rule, self-fulfilling fluctuations are possible in this regime. What do these fluctuations look like? Figure V displays simulated time series of the response of the economy to a sequence of self-fulfilling revisions in inflationary expectations ("sunspot shocks"). We use the baseline estimates of the pre-Volcker regime to fix the parameters of the policy rule (first row of Table II). We draw the "sunspot shocks" to expectations from a standard normal distribution. Our 1997 working paper provides the details.

The figure reports the cyclical behavior of output, inflation, and the nominal interest rate. Note that persistent fluctuations in output and inflation arise, despite the absence of any fundamental shocks.[34] In order to get some intuition for the mechanisms underlying these self-fulfilling fluctuations, Figure VI displays the impulse responses of several variables to a sunspot shock. The

33. In Clarida, Galí, and Gertler [1997] it is shown that the unit threshold value for β obtains exactly only when $\gamma = 0$; i.e., when there is no systematic response to output variations. As we increase γ, the lower bound for β goes down, although the deviation from unity is quantitatively very small (almost negligible), and independent of ρ. In addition, the range of β values for which the equilibrium is unique also has an upper bound. In other words, an "excessive" response to changes in expected inflation may also lead to indeterminacy. This is the case emphasized by Bernanke and Woodford [1997].

34. Generally speaking, the sunspot shocks give rise to positive comovement between output and inflation. To account for the negative comovement that occurred in the 1970s, it appears necessary also to mix in adverse supply shocks. As we make clear in the next subsection, however, supply shocks alone cannot account for the persistent inflation over this era: for them to do so, it is critical also to have an accommodating monetary policy.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

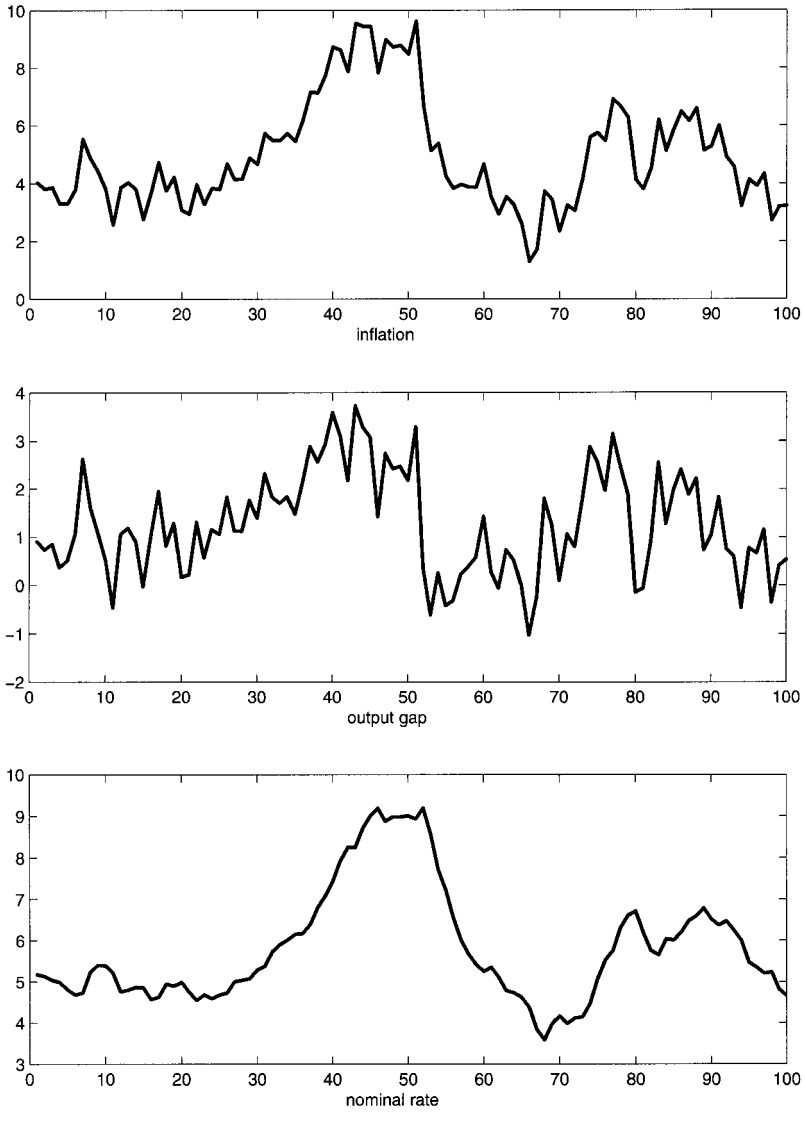


FIGURE V
Simulated Sunspot Fluctuations under Pre-Volcker Rule

sunspot realization generates, on impact, an increase in expected inflation (as well as the anticipation of a slow return to its original level). Given the assumed policy rule, that forecast revision leads to a rise in the nominal rate, but the latter falls short of the

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

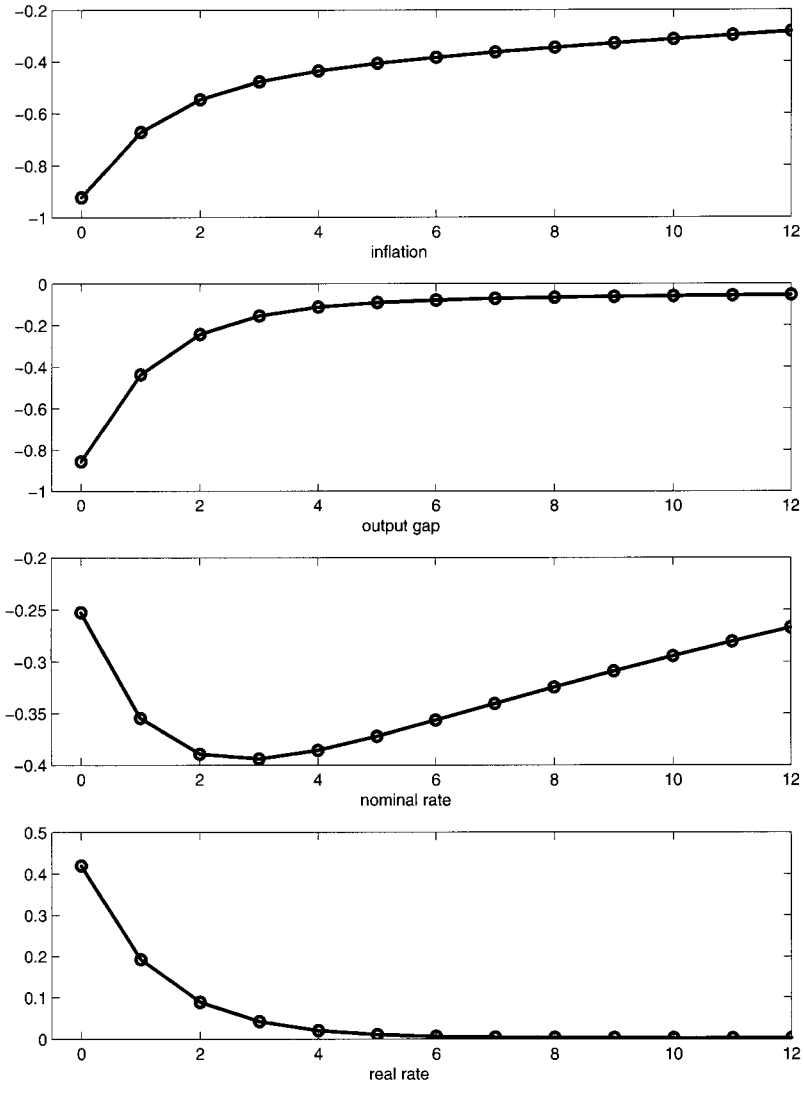


FIGURE VI
Impulse Responses to a Sunspot Shock

increase in expected inflation throughout the entire adjustment process. As a result, the real rate shows a persistent decline, fueling an expansion in output and a rise in inflation, thus validating the initial increase in expected inflation. Over time,

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

output gradually returns to trend, and so do the nominal rate and inflation, as well as the real rate.

By way of contrast, self-fulfilling fluctuations cannot arise under the estimated interest rate rule for the Volcker-Greenspan period. Because β is well above unity in this regime, short-term real rates cannot adjust to accommodate sunspot shifts in inflationary expectations, as they instead do under the estimated pre-Volcker policy. Under this type of regime, accordingly, macroeconomic fluctuations arise only in the presence of shocks to fundamentals. In other words, the monetary policy rule in place is not, *in itself*, a source of macroeconomic instability. On the other hand, the policy rule does affect how the economy responds to fundamental shocks. We explore this issue in the next section.

C. Near-Indeterminacy and Fundamental Shocks

As we have just seen, the point estimates for the inflation and output gap coefficients during the pre-Volcker era fall within the indeterminacy region of the canonical model presented above. This is true for all specifications considered. In some cases, however, the standard errors for the estimate of β are too large to rule out the possibility that the true value is unity or slightly above.[35] In this borderline instance the economy is just outside the region of indeterminacy.

Even though sunspot fluctuations are not feasible in this case, however, the economy is still likely to be relatively unstable under this kind policy regime. We illustrate this issue in the context of our calibrated model. In particular, we compute the implied standard deviation of inflation and output, under alternative values of β, the inflation coefficient and conditional on a given source of fluctuations (supply z_t or demand g_t).[36] Throughout this exercise, we keep both γ and ρ constant and equal to their estimated values for the pre-Volcker period under the baseline specification (i.e., 0.27 and 0.68, respectively).

Table VII reports the results. It shows how the standard deviations of output, inflation, and the output gap vary in

35. In fact, it is possible to find combinations of rule specifications and subperiods within the pre-Volcker era for which the point estimate for β is slightly above one (although not significantly so). We thank an anonymous referee for that observation, which encouraged us to conduct the robustness analysis found above.

36. A general analysis of the effectiveness and desirability of alternative interest rate rules in a model closely related to ours can be found in Rotemberg and Woodford [1997], among others. We differ by focusing on the implications of the feedback parameter on expected inflation.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

TABLE VII
FUNDAMENTAL SHOCKS

	Supply shocks			Demand shocks	
β	$\sigma(\pi)$	$\sigma(x)$	$\sigma(y)$	$\sigma(\pi)$	$\sigma(y)$
2.0	1.00	1.00	1.00	1.00	1.00
1.5	1.48	1.36	1.29	1.61	1.67
1.1	2.57	2.16	2.26	3.04	1.96
1.0	3.20	2.61	2.88	3.88	4.25

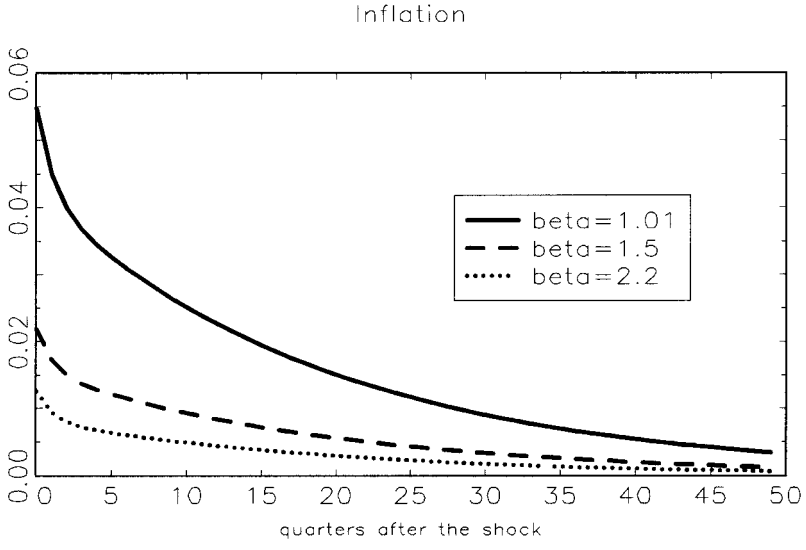
response to shifts in the feedback parameter β. The first three columns display the results conditional on the supply shock z_t being the fundamental driving force. The last three columns repeat the exercise, this time with the demand shock. For ease of exposition, we normalize to unity the standard deviations corresponding to the calibration with $\beta = 2.0$.

The results make it clear that the cyclical response of the economy to fundamental shocks is quite sensitive to β, especially over the estimated range of values across the pre-Volcker and Volcker-Greenspan periods. As β rises from one to two, the volatility of both output and inflation declines by more than half. This occurs for both the supply shock and the demand shock. The reduction in volatility that the Volcker-Greenspan rule achieves is thus quite substantial.

The intuition is straightforward: a value of β equal to or just above unity suggests that the central bank comes close to fully accommodating inflationary pressures, by raising nominal rates to keep real rates roughly constant. The absence of a strong stabilizing adjustment of the real rate suggests, in turn, that fundamental shocks may generate considerable volatility in inflation and the output gap, at least relative to the Volcker-Greenspan type policy, where the central bank adjusts real rates in a strongly countercyclical manner.

We next illustrate explicitly how the ability of supply shocks to generate stagflation (and hence account for the 1970s) depends on the nature of the monetary policy rule. Figure VII reports the response of inflation and output to a negative supply shock (negative shock to z_t) for several values of the feedback coefficient on expected inflation, β, that range from 1.01 to 2.00. The key point is that the supply shock produces a persistent effect on inflation only for values of β near unity. As β rises to 2.00, the

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026



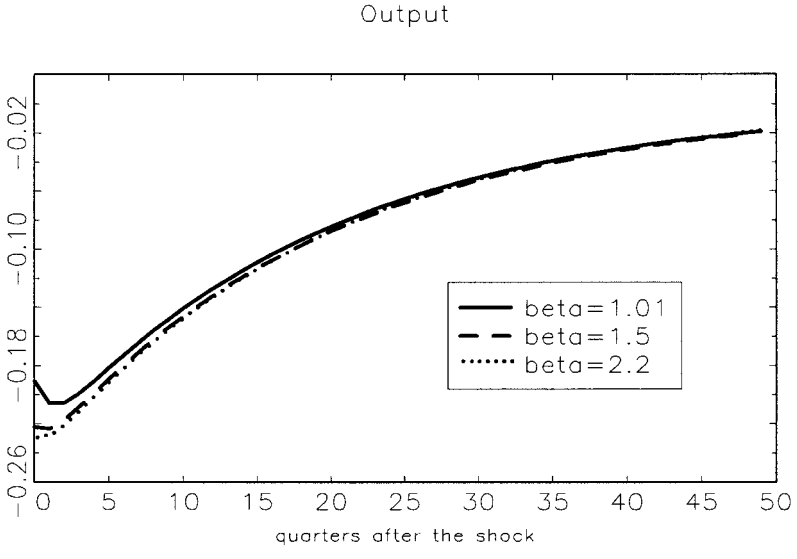


FIGURE VII
Impulse Responses to a Supply Shock

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

supply shock has only a small transitory impact on inflation. In sum, how large and persistent is the response of inflation to the supply shock is quite sensitive to the feedback rule. The accommodating pre-Volcker policy, it appears, can account for a persistent response of inflation. The same does not appear to be true for the Volcker-Greenspan policy, as is consistent with our hypothesis.

Finally, we note that the results we obtain in the context of our simple New Keynesian model are largely robust to using a much broader set of macroeconomic models. Two model features are critical, but they are features that are commonplace in most conventional macroeconomic frameworks. First, there must be an inverse relationship between output and the ex ante real rate (i.e., an IS-type relationship). Second, there must be a positive short-run link between output and inflation (i.e., a Phillips curve). Given these features, a monetary policy rule that accommodates inflationary pressures is, in general, more likely to bring about higher unconditional volatility of inflation and the output gap than otherwise.[37]

V. Summary and Concluding Remarks

In this paper we have provided an empirical characterization of the systematic component of United States monetary policy in the postwar era. In order to do so, we have estimated a simple forward-looking policy reaction function.

Our estimates point to a significant difference in the way monetary policy was conducted pre- and post-late 1979. In the pre-Volcker years the Fed typically raised nominal rates by less than any increase in expected inflation, thus letting real short-term rates decline as anticipated inflation rose. On the other hand, during the Volcker-Greenspan era the Fed raised real as well as nominal short-term interest rates in response to higher expected inflation. Thus, our results lend quantitative support to the view that the anti-inflationary stance of the Fed has been stronger in the past two decades.

Finally, we have argued that the pre-Volcker rule may have contained the seeds of macroeconomic instability that seemed to

37. One difference that arises between the New Keynesian model and a simple backward-looking Keynesian model (see, e.g., Svensson [1996]) is that, in the latter, self-fulfilling fluctuations are not possible since beliefs about the future do not affect behavior. In this instance, a feedback rule with $\beta < 1$ simply leads to explosive behavior, ultimately necessitating a shift in the policy rule.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

characterize the late sixties and seventies. In particular, in the context of a calibrated sticky price model, the pre-Volcker rule leaves open the possibility of bursts of inflation and output that result from self-fulfilling changes in expectations. At best, the pre-Volcker rule does a considerably worse job than the Volcker-Greenspan policy of insulating the economy from fundamental shocks.

One important question our paper raises but does not answer is the following: why is it that during the pre-1979 period the Federal Reserve followed a rule that was clearly inferior? Another way to look at the issue is to ask why it is that the Fed maintained persistently low short-term real rates in the face of high or rising inflation. One possibility, emphasized by De Long [1997], is that the Fed thought the natural rate of unemployment at this time was much lower than it really was (or equivalently, that the output gap was much smaller). There is considerable anecdotal evidence to support this interpretation, although it is not clear why the Fed should have held this view over such a long period of time. Orphanides [1997] emphasizes that preliminary estimates of potential output are often quite different than revised estimates. One possibility, thus, is that during the 1970s the Federal Reserve had consistently overly optimistic preliminary estimates of potential output. Again, one has to explain why the misforecast were so persistently one-sided.

Another somewhat related possibility is that, at that time, neither the Fed nor the economics profession understood the dynamics of inflation very well.[38] Indeed, it was not until the mid-to-late 1970s that intermediate textbooks began emphasizing the absence of a long-run trade-off between inflation and output. The ideas that expectations may matter in generating inflation and that credibility is important in policy-making were simply not well established during that era. What all this suggests is that in understanding historical economic behavior, it is important to take into account the state of policy-maker's knowledge of the economy and how it may have evolved over time. Analyzing policy-making from this perspective, we think, would be a highly useful undertaking.[39]

COLUMBIA UNIVERSITY
NEW YORK UNIVERSITY AND UNIVERSITAT POMPEU FABRA
NEW YORK UNIVERSITY

38. See, e.g., Croushore [1996] for evidence of systematic bias in inflation forecasts during that period.
39. Recent work in this direction can be found in Sargent [1998].

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

REFERENCES

Barsky, Robert, and Lutz Killian, "Money, Stagflation and Oil Prices: A Re-Interpretation," unpublished manuscript, 1998.

Bernanke, Ben, and Ilian Mihov, "Measuring Monetary Policy," *Quarterly Journal of Economics*, CXIII (August 1998), 869–902.

Bernanke, Ben, and Michael Woodford, "Inflation Forecasts and Monetary Policy," *Journal of Money, Credit and Banking*, XXIV (1997), 653–684.

Bernanke, Ben, Mark Gertler, and Mark Watson, "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity* (1997), 91–157.

Blanchard, Olivier, and Charles M. Kahn, "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, XLVIII (1980), 1305–1311.

Calvo, Guillermo, "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, XII (1983), 383–398.

Chari, V. V., Lawrence Christiano, and Martin Eichenbaum, "Expectation Traps and Discretion," Northwestern University, mimeo, 1997.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem," Federal Reserve Bank of Minneapolis, Staff Report, 217, 1996.

Clarida, Richard, Jordi Galí, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," NBER Working Paper No. 6442, 1997.

Clarida, Richard, Jordi Galí, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* (1999), forthcoming.

Croushore, Dean, "Inflation Forecasts: How Good Are They?" *Business Review*, Federal Reserve Bank of Philadelphia (1996), 15–25.

De Long, J. Bradford, "America's Peacetime Inflation: The 1970s," in C. Romer and D. Romer, eds. *Reducing Inflation: Motivation and Strategy* (Chicago: Chicago University Press, 1997).

Friedman, Benjamin M., and Kenneth N. Kuttner, "A Price Target for U. S. Monetary Policy? Lessons from the Experience with Money Growth Targets," *Brookings Papers on Economic Activity* (1996), 126–134.

Galí, Jordi, and Mark Gertler, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* (1999), forthcoming.

Gertler, Mark, "Comments on Friedman and Kuttner," *Brookings Papers on Economic Activity* (1996), 77–125.

Goodfriend, Marvin, "Interest Rates and the Conduct of Monetary Policy," *Carnegie-Rochester Conference Series on Public Policy*, XXXIV (1991), 7–30.

Hamilton, James D., "Oil Price and the Macroeconomy Since World War II," *Journal of Political Economy*, XCI (1983), 228–248.

Hansen, Lars P., "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, L (1982), 1029–1054.

Huizinga, John, and Frederic S. Mishkin, "Monetary Policy Regime Shifts and the Unusual Behavior of Real Interest Rates," *Carnegie-Rochester Conference Series on Public Policy*, XXIV (1986), 231–274.

Kerr, William, and Robert G. King, "Limits on Interest Rate Rules in the IS Model," *Economic Quarterly*, LXXXII (1996), 47–76.

King, Robert G., and Alexander L. Wolman, "Inflation Targeting in a St. Louis Model of the 21st Century," NBER Working Paper No. 5507, 1996.

Mishkin, Frederic S., "The Real Interest Rate: An Empirical Investigation," *Carnegie-Rochester Series on Public Policy*, XV (1981), 151–200.

Orphanides, Athanasios, "Monetary Policy Rules Based on Real-Time Data," Board of Governors of the Federal Reserve System, mimeo, 1997.

Roberts, John M., "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, XXVII (1995), 975–984.

Rotemberg, Julio, and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* (Cambridge, MA: MIT Press, 1997).

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026

Rudebusch, Glenn D., "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure," *Journal of Monetary Economics,* XXXV (1995), 245–274.
Sack, Brian, "Uncertainty and Gradual Monetary Policy," Board of Governors of the Federal Reserve System, mimeo, 1997.
Sargent, Thomas, "The Conquest of American Inflation," unpublished manuscript, 1998.
Sbordone, Argia, "Prices and Unit Labor Costs: A New Test of Sticky Prices," mimeo, 1998.
Svensson, Lars E. O., "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," NBER Working Paper No. 5797, 1996.
Taylor, John B., "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy,* LXXXVIII (1980), 1–23.
____, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Series on Public Policy,* XXXIX (1993), 195–214.
____, "Control of the Public Debt: A Requirement for Price Stability?" NBER Working Paper No. 5684, 1996.
____, "Optimal Monetary Policy Inertia," unpublished manuscript, 1998.
Yun, Tack, "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics,* XXXVII (1996), 345–370.

Downloaded from https://academic.oup.com/qje/article/115/1/147/1842232 by William & Mary user on 10 August 2026



CALL FOR NOMINATIONS

$300,000 Nemmers Prize in Economics

Northwestern University invites nominations for the Erwin Plein Nemmers Prize in Economics, to be awarded during the 2026–27 academic year. The prize pays the recipient $300,000. Recipients of the Nemmers Prize present lectures, participate in department seminars, and engage with Northwestern faculty and students in other scholarly activities.

Details about the prize and the nomination process can be found at **nemmers.northwestern.edu**. Candidacy for the Nemmers Prize is open to those with careers of outstanding achievement in their disciplines as demonstrated by major contributions to new knowledge or the development of significant new modes of analysis.Individuals of all nationalities and institutional affiliations are eligible except current or recent members of the Northwestern University faculty and past recipients of the Nemmers or Nobel Prize. **Nominations will be accepted until January 14, 2026.**

The Nemmers prizes are made possible by a generous gift to Northwestern University by the late Erwin Esser Nemmers and the late Frederic Esser Nemmers.

Northwestern

Nemmers Prizes • Office of the Provost • Northwestern University • Evanston, Illinois 60208
nemmers.northwestern.edu

American Economic Association

The Science of Monetary Policy: A New Keynesian Perspective
Author(s): Richard Clarida, Jordi Galí, Mark Gertler
Source: *Journal of Economic Literature*, Vol. 37, No. 4 (Dec., 1999), pp. 1661-1707
Published by: American Economic Association
Stable URL: http://www.jstor.org/stable/2565488
Accessed: 21/01/2009 20:36

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=aea.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *Journal of Economic Literature.*

http://www.jstor.org

Journal of Economic Literature
Vol. XXXVII (December 1999), pp. 1661–1707

The Science of Monetary Policy: A New Keynesian Perspective

Richard Clarida, Jordi Galí, *and* Mark Gertler[1]

"Having looked at monetary policy from both sides now, I can testify that central banking in practice is as much art as science. Nonetheless, while practicing this dark art, I have always found the science quite useful."[2]
Alan S. Blinder

1. *Introduction*

THERE HAS BEEN a great resurgence of interest in the issue of how to conduct monetary policy. One symptom of this phenomenon is the enormous volume of recent working papers and conferences on the topic. Another is that over the past several years many leading macroeconomists have either proposed specific policy rules or have at least staked out a position on what the general course of monetary policy should be. John Taylor's recommendation of a simple interest rate rule (Taylor 1993a) is a well-known example. So too is the recent widespread endorsement of inflation targeting (e.g., Ben Bernanke and Frederic Mishkin 1997).

Two main factors underlie this rebirth of interest. First, after a long period of near exclusive focus on the role of nonmonetary factors in the business cycle, a stream of empirical work beginning in the late 1980s has made the case that monetary policy significantly influences the short-term course of the real economy.[3] The precise amount remains open to debate. On the other hand, there now seems to be broad agreement that the choice of how to conduct monetary policy has important consequences for aggregate activity. It is no longer an issue to downplay.

Second, there has been considerable improvement in the underlying theoretical frameworks used for policy analysis. To provide theoretical underpinnings, the literature has incorporated the techniques of dynamic general equilibrium theory pioneered in real business cycle

[1] Clarida: Columbia University and NBER; Galí: New York University, Universitat Pompeu Fabra, CEPR, and NBER; Gertler: New York University and NBER. Thanks to Ben Bernanke, Bob King, Ben McCallum, Albert Marcet, Rick Mishkin, Athanasios Orphanides, Glenn Rudebusch, Chris Sims, Lars Svensson, Andres Velasco, and several anonymous referees for helpful comments, and to Tommaso Monacelli for excellent research assistance. Authors Galí and Gertler are grateful to the C.V. Starr Center for Applied Economics, and (Galí) to CREI for financial support. e-mail: mark.gertler@econ.nyu.edu

[2] Blinder 1997, p. 17.

[3] Examples include Romer and Romer (1988), Bernanke and Blinder (1992), Galí (1992), Bernanke and Mihov (1997a), Christiano, Eichenbaum, and Evans (1996, 1998) and Leeper, Sims and Zha (1996). Much of the literature has focused on the effects of monetary policy shocks. Bernanke, Gertler, and Watson (1997) present evidence that suggests that the monetary policy rule may have important effects on real activity.

analysis. A key point of departure from real business cycle theory (as we later make clear) is the explicit incorporation of frictions such as nominal price rigidities that are needed to make the framework suitable for evaluation of monetary policy.

This paper summarizes what we have learned from this recent research on monetary policy. We review the progress that has been made and also identify the central questions that remain. To organize the discussion, we exposit the monetary policy design problem in a simple theoretical model. We start with a stripped-down baseline model in order to characterize a number of broad principles that underlie optimal policy management. We then consider the implications of adding various real world complications. Finally, we assess how the predictions from theory square with policy-making in practice.

Throughout, we concentrate on expositing results that are robust across a wide variety of macroeconomic frameworks. As Ben McCallum (1997b) emphasizes, the key stumbling block for policy formation is limited knowledge of the way the macroeconomy works. Results that are highly model-specific are of limited use. This literature, however, contains a number of useful principles about optimal policy that are reasonably general in applicability. In this respect there is a "science of monetary policy," as Alan Blinder suggests in the quote above. We provide support for this contention in the pages that follow.

At the same time, we should make clear that the approach we take is based on the idea that temporary nominal price rigidities provide the key friction that gives rise to nonneutral effects of monetary policy. The propositions we derive are broadly applicable within this class of models. This approach has widespread support in both theoretical and applied work, as we discuss later.[4] There are, however, important strands of the literature that either reject the idea of nominal price rigidities (e.g., real business cycle theory) or focus on other types of nominal rigidities, such as frictions in money demand.[5] For this reason, we append "New Keynesian Perspective" to the title. In particular, we wish to make clear that we adopt the Keynesian approach of stressing nominal price rigidities, but at the same time base our analysis on frameworks that incorporate the recent methodological advances in macroeconomic modeling (hence the term "New").

Section 2 lays out the formal policy problem. We describe the baseline theoretical model and the objectives of policy. Because we are interested in characterizing policy rules in terms of primitive factors, the model we use evolves from first principles. Though it is quite simple, it nonetheless contains the main ingredients of descriptively richer frameworks that are used for policy analysis. Within the model, as in practice (we argue), the instrument of monetary policy is a short-term interest rate. The policy design problem then is to characterize how the interest rate should adjust to the current state of the economy.

An important complication is that private sector behavior depends on the expected course of monetary policy, as well as on current policy. The credibility of monetary policy thus becomes relevant, as a considerable contemporary literature has emphasized.[6] At issue is

[4] See, for example, the survey by Goodfriend and King (1997).

[5] See, for example, Christiano, Eichenbaum, and Evans (1997). For an analysis of monetary policy rules in these kinds of models—known as "limited participation" frameworks—see Christiano and Gust (1999).

[6] For a recent survey of the credibility literature, see Persson and Tabellini (1997).

whether there may be gains from enhancing credibility either by formal commitment to a policy rule or by introducing some kind of institutional arrangement that achieves roughly the same end. We address the issue by examining optimal policy for both cases: with and without commitment. Along with expositing traditional results, we also exposit some new results regarding the gains from commitment.

Section 3 derives the optimal policy rule in the absence of commitment. If for no other reason, this case is of interest because it captures reality: No major central bank makes any type of binding commitment over the future course of its monetary policy. A number of broad implications emerge from this baseline case. Among these: The optimal policy embeds inflation targeting in the sense that it calls for gradual adjustment to the optimal inflation rate. The implication for the policy rule is that the central bank should adjust the nominal short rate more than one-for-one with expected future inflation. That is, it should adjust the nominal rate sufficiently to alter the real rate (and thus aggregate demand) in the direction that is offsetting to any movement in expected inflation. Finally, how the central bank should adjust the interest rate in response to output disturbances depends critically on the nature of the disturbances: It should offset demand shocks but accommodate supply shocks, as we discuss.

Section 4 turns to the case with commitment. Much of the literature has emphasized that an inefficiently high steady state inflation rate may arise in the absence of commitment, if the central bank's target for real output exceeds the market clearing level.[7] The gain from commitment then is to eliminate this inflationary bias. How realistic it is to presume that a perceptive central bank will try to inadvisedly reap short-term gains from pushing output above its natural level is a matter of recent controversy (e.g., Blinder 1997; McCallum 1997a). We demonstrate, however, that there may be gains from commitment simply if current price setting depends on expectations of the future. In this instance, a credible commitment to fight inflation in the future can improve the current output/inflation trade-off that a central bank faces. Specifically, it can reduce the effective cost in terms of current output loss that is required to lower current inflation. This result, we believe, is new in the literature.

In practice, however, a binding commitment to a rule may not be feasible simply because not enough is known about the structure of the economy or the disturbances that buffet it. Under certain circumstances, however, a policy rule that yields welfare gains relative to the optimum under discretion may be well approximated by an optimal policy under discretion that is obtained by assigning a higher relative cost to inflation than the true social cost. A way to pursue this policy operationally is simply to appoint a central bank chair with a greater distaste for inflation than society as a whole, as Kenneth Rogoff (1985) originally emphasized.

Section 5 considers a number of practical problems that complicate policymaking. These include: imperfect information and lags, model uncertainty and non-smooth preferences over inflation and output. A number of pragmatic issues emerge, such as: whether and how to make use of intermediate targets, the choice of a monetary policy instrument, and why central banks appear to smooth interest rate changes. Among other

[7] The potential inflationary bias under discretion was originally emphasized by Kydland and Prescott (1977) and Barro and Gordon (1983).

things, the analysis makes clear why modern central banks (especially the Federal Reserve Board) have greatly downgraded the role of monetary aggregates in the implementation of policy. The section also shows how the recently advocated "opportunistic" approach to fighting inflation may emerge under a non-smooth policy objective function. The opportunistic approach boils down to trying to keep inflation from rising but allowing it to ratchet down in the event of favorable supply shocks.

As we illustrate throughout, the optimal policy depends on the degree of persistence in both inflation and output. The degree of inflation persistence is critical since this factor governs the output/inflation trade-off that the policy-maker faces. In our baseline model, persistence in inflation and output is due entirely to serially correlated exogenous shocks. In section 6 we consider a hybrid model that allows for endogenous persistence in both inflation and output. The model nests as special cases our forward-looking baseline model and, also, a more traditional backward-looking Keynesian framework, similar to the one used by Lars Svensson (1997a) and others.

Section 7 moves from theory to practice by considering a number of proposed simple rules for monetary policy, including the Taylor rule, and a forward-looking variant considered by Clarida, Galí, and Gertler (1998; forthcoming). Attention has centered around simple rules because of the need for robustness. A policy rule is robust if it produces desirable results in a variety of competing macroeconomic frameworks. This is tantamount to having the rule satisfy the criteria for good policy management that sections 2 through 6 establish. Further, U.S. monetary policy may be judged according to this same metric. In particular, the evidence suggests that U.S. monetary policy in the fifteen years or so prior to Paul Volcker did not always follow the principles we have described. Simply put, interest rate management during this era tended to accommodate inflation. Under Volcker and Greenspan, however, U.S. monetary policy adopted the kind of implicit inflation targeting that we argue is consistent with good policy management.

The section also considers some policy proposals that focus on target variables, including introducing formal inflation or price-level targets and nominal GDP targeting. There is in addition a brief discussion of the issue of whether indeterminacy may cause practical problems for the implementation of simple interest rate rules. Finally, there are concluding remarks in section 8.

2. *A Baseline Framework for Analysis of Monetary Policy*

This section characterizes the formal monetary policy design problem. It first presents a simple baseline macroeconomic framework, and then describes the policy objective function. The issue of credibility is taken up next. In this regard, we describe the distinction between optimal policies with and without credible commitment—what the literature refers to as the cases of "rules versus discretion."

2.1 *A Simple Macroeconomic Framework*

Our baseline framework is a dynamic general equilibrium model with money and temporary nominal price rigidities. In recent years this paradigm has become widely used for theoretical analysis of monetary policy.[8] It has much of the empirical appeal of the traditional

[8] See, e.g, Goodfriend and King (1997), McCallum and Nelson (1997), Walsh (1998), and the references therein.

IS/LM model, yet is grounded in dynamic general equilibrium theory, in keeping with the methodological advances in modern macroeconomics.

Within the model, monetary policy affects the real economy in the short run, much as in the traditional Keynesian IS/LM framework. A key difference, however, is that the aggregate behavioral equations evolve explicitly from optimization by households and firms. One important implication is that current economic behavior depends critically on expectations of the future course of monetary policy, as well as on current policy. In addition, the model accommodates differing views about how the macroeconomy behaves. In the limiting case of perfect price flexibility, for example, the cyclical dynamics resemble those of a real business cycle model, with monetary policy affecting only nominal variables.

Rather than work through the details of the derivation, which are readily available elsewhere, we instead directly introduce the key aggregate relationships.[9] For convenience, we abstract from investment and capital accumulation. This abstraction, however, does not affect any qualitative conclusions, as we discuss. The model is as follows:

Let y_t and z_t be the stochastic components of output and the natural level of output, respectively, both in logs.[10] The latter is the level of output that would arise if wages and prices were perfectly flexible. The difference between actual and potential output is an important variable in the model. It is thus convenient to define the "output gap" x_t:

$$x_t \equiv y_t - z_t$$

In addition, let π_t be the period t inflation rate, defined as the percent change in the price level from $t-1$ to t; and let i_t be the nominal interest rate. Each variable is similarly expressed as a deviation from its long-run level.

It is then possible to represent the baseline model in terms of two equations: an "IS" curve that relates the output gap inversely to the real interest rate; and a Phillips curve that relates inflation positively to the output gap.

$$x_t = -\varphi[i_t - E_t\pi_{t+1}] + E_t x_{t+1} + g_t \quad (2.1)$$

$$\pi_t = \lambda x_t + \beta E_t \pi_{t+1} + u_t \quad (2.2)$$

where g_t and u_t are disturbances terms that obey, respectively:

$$g_t = \mu g_{t-1} + \hat{g}_t \quad (2.3)$$

$$u_t = \rho u_{t-1} + \hat{u}_t \quad (2.4)$$

where $0 \leq \mu, \rho \leq 1$ and where both $\hat{g}_t$ and $\hat{u}_t$ are i.i.d. random variables with zero mean and variances σ_g^2 and σ_u^2, respectively.

Equation (2.1) is obtained by log-linearizing the consumption euler equation that arises from the household's optimal saving decision, after imposing the equilibrium condition that consumption equals output minus government spending.[11] The resulting expression differs from the traditional IS curve mainly because current output depends on expected future output as well as the interest rate. Higher expected future output raises current output: Because individuals prefer to

[9] See, for example, Yun (1996), Kimball (1995), King and Wolman (1995), Woodford (1996), and Bernanke, Gertler, and Gilchrist (1998) for step-by-step derivations.

[10] By stochastic component, we mean the deviation from a deterministic long-run trend.

[11] Using the market clearing condition $Y_t = C_t + E_t$, where E_t is government consumption, we can rewrite the log-linearized consumption Euler equation as:

$$y_t - e_t = -\varphi[i_t - E_t\pi_{t+1}] + E_t\{y_{t+1} - e_{t+1}\}$$

where $e_t \equiv -\log(1 - \frac{E_t}{Y_t})$ is taken to evolve exogenously. Using $x_t \equiv y_t - z_t$, it is then possible to derive the demand for output as

$$x_t = -\varphi[i_t - E_t\pi_{t+1}] + E_t x_{t+1} + g_t$$

where $g_t = E_t\{\Delta z_{t+1} - \Delta e_{t+1}\}$.

smooth consumption, expectation of higher consumption next period (associated with higher expected output) leads them to want to consume more today, which raises current output demand. The negative effect of the real rate on current output, in turn, reflects intertemporal substitution of consumption. In this respect, the interest elasticity in the IS curve, φ, corresponds to the intertemporal elasticity of substitution. The disturbance g_t is a function of expected changes in government purchases relative to expected changes in potential output (see footnote 11). Since g_t shifts the IS curve, it is interpretable as a demand shock. Finally, adding investment and capital to the model changes the details of equation (2.1). But it does not change the fundamental qualitative aspects: output demand still depends inversely on the real rate and positively on expected future output.

It is instructive to iterate equation (2.1) forward to obtain

$$x_t = E_t \sum_{i=0}^{\infty} \{ -\varphi[i_{t+i} - \pi_{t+1+i}] + g_{t+i} \} \qquad (2.5)$$

Equation (2.5) makes transparent the degree to which beliefs about the future affect current aggregate activity within this framework. The output gap depends not only on the current real rate and the demand shock, but also on the expected future paths of these two variables. To the extent monetary policy has leverage over the short-term real rate due to nominal rigidities, equation (2.5) suggests that expected as well as current policy actions affect aggregate demand.

The Phillips curve, (2.2), evolves from staggered nominal price setting, in the spirit of Stanley Fischer (1977) and John Taylor (1980).[12] A key difference is that the individual firm price-setting decision, which provides the basis for the aggregate relation, is derived from an explicit optimization problem. The starting point is an environment with monopolistically competitive firms: When it has the opportunity, each firm chooses its nominal price to maximize profits subject to constraints on the frequency of future price adjustments.

Under the standard scenario, each period the fraction $1/X$ of firms set prices for $X > 1$ periods. In general, however, aggregating the decision rules of firms that are setting prices on a staggered basis is cumbersome. For this reason, underlying the specific derivation of equation (2.2) is an assumption due to Guillermo Calvo (1983) that greatly simplifies the problem: In any given period a firm has a fixed probability θ it must keep its price fixed during that period and, hence a probability $1 - \theta$ that it may adjust.[13] This probability, further, is independent of the time that has elapsed since the last time the firm changed price. Accordingly, the average time over which a price is fixed is $\frac{1}{1-\theta}$. Thus, for example, if $\theta = .75$, prices are fixed on average for a year. The Calvo formulation thus captures the spirit of staggered setting, but facilitates the aggregation by making the timing of a firm's price adjustment independent of its history.

Equation (2.2) is simply a loglinear approximation about the steady state of the aggregation of the individual firm pricing decisions. Since the equation relates the inflation rate to the output gap and expected inflation, it has the flavor of a traditional expectations-augmented Phillips curve (see, e.g., Olivier Blanchard

[12] See Galí and Gertler (1998) and Sbordone (1998) for some empirical support for this kind of Phillips curve relation.

[13] The Calvo formulation has become quite common in the literature. Work by Yun (1996), King and Wolman (1995), Woodford (1996) and others has initiated the revival.

1997). A key difference with the standard Phillips curve is that expected future inflation, $E_t\pi_{t+1}$, enters additively, as opposed to expected current inflation, $E_{t-1}\pi_t$.[14] The implications of this distinction are critical: To see, iterate (2.2) forward to obtain

$$\pi_t = E_t\sum_{i=0}^{\infty}\beta^i[\lambda x_{t+i} + u_{t+i}] \qquad (2.6)$$

In contrast to the traditional Phillips curve, there is no arbitrary inertia or lagged dependence in inflation. Rather, inflation depends entirely on current and expected future economic conditions. Roughly speaking, firms set nominal price based on the expectations of future marginal costs. The variable x_{t+i} captures movements in marginal costs associated with variation in excess demand. The shock u_{t+i}, which we refer to as "cost push," captures anything else that might affect expected marginal costs.[15] We allow for the cost push shock to enable the model to generate variation in inflation that arises independently of movement in excess demand, as appears present in the data (see, e.g., Fuhrer and Moore 1995).

To close the model, we take the nominal interest rate as the instrument of monetary policy, as opposed to a money supply aggregate. As Bernanke and Ilian Mihov (1998) show, this assumption provides a reasonable description of Federal Reserve operating procedures since 1965, except for the brief period of non-borrowed reserves targeting (1980–82) under Paul Volcker.[16] With the nominal rate as the policy instrument, it is not necessary to specify a money market equilibrium condition (i.e., an LM curve).[17] In section 5, we discuss the implications of using instead a narrow monetary aggregate as the policy instrument.

Though simple, the model has the same qualitative core features as more

[14] Another key difference is that the explicit derivation restricts the coefficient λ on the output gap. In particular, λ is decreasing in θ, which measures the degree of price rigidity. Thus, the longer prices are fixed on average, the less sensitive is inflation to movements in the output gap.

[15] The relation for inflation that evolves from the Calvo model takes the form

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \delta\, mc_t$$

where mc_t denotes the deviation of (real) marginal cost from its steady state value. To then relate inflation to the output gap, the literature typically makes assumptions on technology, preferences, and the structure of labor markets to justify a proportionate relation between real marginal cost and the output gap, so that $mc_t = \kappa x_t$ holds, where κ is the output elasticity of real marginal cost. In this instance, one can rewrite the relation for inflation in terms of the output gap, as follows: $\pi_t = \beta E_t\{\pi_{t+1}\} + \lambda x_t$ (see Galí and Gertler (1998) for details). In this context, the disturbance u_t in (2.2) is interpretable as reflecting deviations from the condition $mc_t = \kappa x_t$. (Indeed the evidence in Galí and Gertler 1998 suggests that mc_t does not vary proportionately with x_t). Deviations from this proportionality condition could be caused, for example, by movements in nominal wages that push real wages away from their "equilibrium" values due to frictions in the wage contracting process. On this latter point, see Erceg, Henderson, and Levin (1998). Another interpretation of the u_t shock (suggested by Mike Woodford) is that it could reflect a shock to the gap between the natural and potential levels of output (e.g., a markup shock).

[16] Roughly speaking, Bernanke and Mihov (1998) present formal evidence showing that the Federal Reserve intervenes in the market for non-borrowed bank reserves to support its choice for the level of the Federal Funds rate, the overnight market for bank reserves. (Christiano, Eichenbaum, and Evans 1998, though, take issue with the identifying assumptions in the Bernanke-Mihov test). Informally, Federal Reserve policy actions in recent years routinely take the form of announcing a target for the Federal funds rate (see, e.g, Rudebusch 1995). Policy discussions, further, focus on whether to adjust that target, and by how much. In this context, the view that the Funds rate is the policy instrument is widely held by both practitioners of monetary policy and academic researchers (see, e.g., Goodfriend 1991, Taylor 1993, and Walsh 1998).

[17] With the interest rate as the policy instrument, the central bank adjusts the money supply to hit the interest rate target. In this instance, the condition that money demand equal money supply simply determines the value of the money supply that meets this criteria.

complex, empirically based frameworks that are used for policy analysis.[18] As in these applied frameworks, temporary nominal price rigidities play a critical role. With nominal rigidities present, by varying the nominal rate, monetary policy can effectively change the short-term real rate. Through this classic mechanism it gains leverage over the near term course of the real economy. In contrast to the traditional mechanism, though, beliefs about how the central bank will set the interest rate in the future also matter, since both households and firms are forward looking. In this kind of environment, how monetary policy should respond in the short run to disturbances that buffet the economy is a nontrivial decision. Resolving this issue is the essence of the contemporary debate over monetary policy.

2.2 *The Policy Objective*

The central bank objective function translates the behavior of the target variables into a welfare measure to guide the policy choice. We assume, following much of the literature, that this objective function is over the target variables x_t and π_t, and takes the form:

$$\max -\frac{1}{2} E_t\left\{\sum_{i=0}^{\infty} \beta^i [\alpha x_{t+i}^2 + \pi_{t+i}^2]\right\} \quad (2.7)$$

where the parameter α is the relative weight on output deviations. Since $x_t \equiv y_t - z_t$, the loss function takes potential output z_t as the target. It also implicitly takes zero as the target inflation, but there is no cost in terms of generality since inflation is expressed as a percent deviation from trend.[19]

While there has been considerable progress in motivating behavioral macroeconomic models from first principles, until very recently, the same has not been true about rationalizing the objectives of policy. Over the past several years, there have been a number of attempts to be completely coherent in formulating the policy problem by taking as the welfare criterion the utility of a representative agent within the model.[20]

One limitation of this approach, however, is that the models that are currently available do not seem to capture what many would argue is a major cost of inflation, the uncertainty that its variability generates for lifetime financial planning and for business planning (see, e.g., Brad DeLong 1997).[21] Another issue is that, while the widely used representative agent approach may be a reasonable way to motivate behavioral relationships, it could be highly misleading as a guide to welfare analysis. If some groups suffer more in recessions than others (e.g. steel workers versus professors) and there are incomplete insurance and credit markets, then the utility of a hypothetical representative agent might not provide an accurate barometer of cyclical fluctuations in welfare.

With certain exceptions, much of the

[18] Some prominent examples include the recently renovated large scale model used by the Federal Reserve Board, the FRB-US model (see Brayton, Levin, Tyron, and Williams 1997), and the medium scale models of Taylor (1979, 1993b) and Fuhrer and Moore (1995a,b).

[19] Put differently, under the optimal policy, the target inflation rate pins down the trend inflation rate. The loss function thus penalizes deviations from this trend.

[20] Some examples of this approach include Aiyagari and Braun (1997), King and Wolman (1995), Ireland (1996a), Carlstrom and Fuerst (1995), and Rotemberg and Woodford (1997).

[21] Underlying this kind of cost is the observation that contracts are typically written in nominal terms and, for reasons that are difficult to explain, not perfectly indexed to the price level. On this issue, see the discussion in Shiller (1997) and the associated comment by Hall (1997).

literature takes a pragmatic approach to this issue by simply assuming that the objective of monetary policy is to minimize the squared deviations of output and inflation from their respective target levels. However, Julio Rotemberg and Michael Woodford (1999) and Woodford (1998) provide a formal justification for this approach. These authors show that an objective function looking something like equation (2.7) may be obtained as a quadratic approximation of the utility-based welfare function. In this instance, the relative weight, α, is a function of the primitive parameters of the model.

In what follows, we simply adopt the quadratic objective given by (2.7), appealing loosely to the justification offered in Rotemberg and Woodford (1999). Judging by the number of papers written by Federal Reserve economists that follow this lead, this formulation does not seem out of sync with the way monetary policy operates in practice (at least implicitly).[22] The target level of output is typically taken to be the natural level of output, based on the idea that this is the level of output that would obtain absent any wage and price frictions. Yet, if distortions exist in the economy (e.g., imperfect competition or taxes), a case can be made that the welfare maximizing level of output may exceed its natural level. This issue becomes important in the context of policy credibility, but we defer it for now.

What should be the target rate of inflation is perhaps an even more ephemeral question, as is the issue of what should be the relative weight assigned to output and inflation losses. In the U.S., policy-makers argue that "price stability" should be the ultimate goal. But they define price stability as the inflation rate at which inflation is no longer a public concern. In practice, it is argued that an inflation rate between one and three percent seems to meet this definition (e.g., Bernanke and Mishkin 1997). A further justification for this criteria is that the official price indices may be overstating the true inflation rate by a percent or two, as argued recently by the Boskin Commission. In this regard, interestingly, the Bundesbank has had for a long time an official inflation target of two percent.[23] They similarly argue that this positive rate of inflation is consistent with price stability, and cite measurement error as one of the reasons (Clarida and Gertler 1997).

It is clear that the experience of the 1970s awakened policy-makers to the costs of high inflation (DeLong 1997). Otherwise, there is no directly observable indicator of the relative weights assigned to output and inflation objectives. Nor, argues Blinder (1997), is there any obvious consensus among policy-makers about what these weights really are in practice. It is true that there has been a growing consensus that the primary aim of monetary policy should be to control inflation (see, e.g., Bernanke and Mishkin 1997). But this discussion in many respects is about what kind of policy rule may be best, as opposed to what the underlying welfare function looks like.

For our purposes, however, it is reasonable to take the inflation target and preference parameters as given and simply explore the implications for optimal policy rules.

[22] See, for example, Williams (1997) and references therein.

[23] Two percent is also the upper bound of the inflation target range established by the European Central Bank. On the other hand, Feldstein (1997) argues that the tax distortions that arise because corporate and personal income taxes are not indexed to inflation justify moving from three percent to zero inflation.

2.3 *The Policy Problem and Discretion versus Rules*

The policy problem is to choose a time path for the instrument i_t to engineer time paths of the target variables x_t and π_t that maximize the objective function (2.7), subject to the constraints on behavior implied by (2.1) and (2.2). This formulation is in many ways in the tradition of the classic Jan Tinbergen (1952)/Henri Theil (1961) (TT) targets and instruments problem. As with TT, the combination of quadratic loss and linear constraints yields a certainty equivalent decision rule for the path of the instrument. The optimal feedback rule, in general, relates the instrument to the state of the economy.

There is, however, an important difference from the classic problem: The target variables depend not only on the current policy but also on expectations about future policy: The output gap depends on the future path of the interest rate (equation 2.5); and, in turn, inflation depends on the current and expected future behavior of the output gap (equation 2.6). As Finn Kydland and Edward Prescott (1977) originally emphasized, in this kind of environment, credibility of future policy intentions becomes a critical issue. For example, a central bank that can credibly signal its intent to maintain inflation low in the future may be able to reduce current inflation with less cost in terms of output reduction than might otherwise be required.[24] In section 4, we illustrate this point explicitly.

From the standpoint of policy design, the issue is to identify whether some type of credibility-enhancing commitment may be desirable. Answering this question boils down to comparing optimal policy under discretion versus rules (using the terminology of the literature). In our context, a central bank operating under discretion chooses the current interest rate by reoptimizing every period. Any promises made in the past do not constrain current policy. Under a rule, it chooses a plan for the path of the interest rates that it sticks to forever. The plan may call for adjusting the interest rate in response to the state of the economy, but both the nature and size of the response are etched in stone.

Two points need to be emphasized. First, the key distinction between discretion and rules is whether current commitments constrain the future course of policy in any credible way. In each instance, the optimal outcome is a feedback policy that relates the policy instrument to the current state of the economy in a very specific way. The two approaches differ, however, in their implications for the link between policy intentions and private sector beliefs. Under discretion, a perceptive private sector forms its expectations taking into account how the central bank adjusts policy, given that the central bank is free to reoptimize every period. The rational expectations equilibrium thus has the property that the central bank has no incentive to change its plans in an unexpected way, even though it has the discretion to do so. (For this reason, the policy that emerges in equilibrium under discretion is termed "time consistent.") In contrast, under a rule, it is simply

[24] In this regard, we stress further that, in contrast to conventional wisdom, the issue of credibility in monetary policy is not tied to central bank objectives over output. In the classic, Barro/Gordon (1983) formulation (and countless papers thereafter), the central bank's desire to push output above potential output gives rise to the credibility problem. However, as we make clear in section 4, gains from commitment potentially emerge whenever private sector behavior depends on beliefs about the future, even if central bank objectives over output are perfectly aligned.

the binding commitment that makes the policy believable in equilibrium.

Second, (it should almost go without saying that) the models we use are nowhere near the point where it is possible to obtain a tightly specified policy rule that could be recommended for practical use with great confidence. Nonetheless, it is useful to work through the cases of discretion and rules in order to develop a set of normative guidelines for policy behavior. As Taylor (1993a) argues, common sense application of these guidelines may improve the performance of monetary policy. We expand on this point later. In addition, understanding the qualitative differences between outcomes under discretion versus rules can provide lessons for the institutional design of monetary policy. For example, as we discuss, Rogoff's (1985) insightful analysis of the benefits of a conservative central bank chair is a product of this type of analysis. Finally, simply understanding the qualitative aspects of optimal policy management under discretion can provide useful normative insights, as we show shortly.

We proceed in the next section to derive the optimal policy under discretion. In a subsequent section we then evaluate the implications of commitment.

3. *Optimal Monetary Policy without Commitment*

We begin with the case without commitment ("discretion") for two reasons. First, at a basic level this scenario accords best with reality. In practice, no major central bank makes any kind of binding commitment over the course of its future monetary policy. In this respect, it seems paramount to understand the nature of optimal policy in this environment. Second, as we have just discussed, to fully comprehend the possible gains from commitment to a policy rule and other institutional devices that might enhance credibility, it is necessary to understand what the benchmark case of discretion yields.

Under discretion, each period the central bank chooses the triplet $\{x_t, \pi_t, i_t\}$, consisting of the two target variables and the policy instrument, to maximize the objective (2.7) subject to the aggregate supply curve (2.2) and the IS curve, (2.1). It is convenient to divide the problem into two stages: First, the central bank chooses x_t and π_t to maximize the objective (2.7), given the inflation equation (2.2).[25] Then, conditional on the optimal values of x_t and π_t, it determines the value of i_t implied by the IS curve (2.1) (i.e., the interest rate that will support x_t and π_t).

Since it cannot credibly manipulate beliefs in the absence of commitment, the central bank takes private sector expectations as given in solving the optimization problem.[26] (Then, conditional on the central bank's optimal rule, the private sector forms beliefs rationally.) Because there are no endogenous state variables, the first stage of the policy problem reduces to the following sequence of static optimization

[25] Since all the qualitative results we derive stem mainly from the first stage problem, what is critical is the nature of the short run Phillips curve. For our baseline analysis, we use the Phillips curve implied the New Keynesian model. In section 6 we consider a very general Phillips curve that is a hybrid of different approaches and show that the qualitative results remain intact. It is in this sense that our analysis is quite robust.

[26] We are ignoring the possibility of reputational equilibria that could support a more efficient outcome. That is, in the language of game theory, we restrict attention to Markov perfect equilibria. One issue that arises with reputational equilibria is that there are multiplicity of possible equilibria. Rogoff (1987) argues that the fragility of the resulting equilibria is an unsatisfactory feature of this approach. See also, Ireland (1996b). On the other hand, Chari, Christiano, and Eichenbaum (1998) argue that this indeterminacy could provide a source of business fluctuations.

problems:[27] Each period, choose x_t and π_t to maximize

$$-\frac{1}{2}[\alpha x_t^2 + \pi_t^2] + F_t \tag{3.1}$$

subject to

$$\pi_t = \lambda x_t + f_t \tag{3.2}$$

taking as given F_t and f_t, where

$$F_t \equiv -\frac{1}{2} E_t\left\{\sum_{i=1}^{\infty} \beta^i[\alpha x_{t+i}^2 + \pi_{t+i}^2]\right\} \text{ and}$$

$f_t \equiv \beta E_t \pi_{t+1} + u_t$. Equations (3.1) and (3.2) simply reformulate (2.7) and (2.2) in a way that makes transparent that, under discretion, (a) future inflation and output are not affected by today's actions, and (b) the central bank cannot directly manipulate expectations.

The solution to the first stage problem yields the following optimality condition:

$$x_t = -\frac{\lambda}{\alpha}\pi_t \tag{3.3}$$

This condition implies simply that the central bank pursue a "lean against the wind" policy: Whenever inflation is above target, contract demand below capacity (by raising the interest rate); and vice-versa when it is below target. How aggressively the central bank should reduce x_t depends positively on the gain in reduced inflation per unit of output loss, λ, and inversely on the relative weight placed on output losses, α.

To obtain reduced form expressions for x_t and π_t, combine the optimality condition (fonc) with the aggregate supply curve (AS), and then impose that private sector expectations are rational:

$$x_t = -\lambda q\, u_t \tag{3.4}$$

$$\pi_t = \alpha q\, u_t \tag{3.5}$$

where

$$q = \frac{1}{\lambda^2 + \alpha(1 - \beta\rho)}$$

The optimal feedback policy for the interest rate is then found by simply inserting the desired value of x_t in the IS curve (2.1):

$$i_t = \gamma_\pi E_t \pi_{t+1} + \frac{1}{\varphi} g_t \tag{3.6}$$

where

$$\gamma_\pi = 1 + \frac{(1-\rho)\lambda}{\rho\varphi\alpha} > 1$$

$$E_t\pi_{t+1} = \rho\, \pi_t = \rho\alpha q\, u_t$$

This completes the formal description of the optimal policy.

From this relatively parsimonious set of expressions there emerge a number of key results that are reasonably robust findings of the literature:

Result 1: *To the extent cost push inflation is present, there exists a short run trade-off between inflation and output variability.*

This result was originally emphasized by Taylor (1979) and is an important guiding principle in many applied studies of monetary policy that have followed.[28] A useful way to illustrate the trade-off implied by the model is to construct the corresponding efficient policy frontier. The device is a locus of points that characterize how the unconditional standard deviations of output and inflation under the optimal policy, σ_x and σ_π, vary with central bank preferences, as defined by α. Figure 1 portrays the efficient policy frontier for our

[27] In section 6, we solve for the optimum under discretion for the case where an endogenous state variable is present. Within the Markov perfect equilibrium, the central bank takes private sector beliefs as a given function of the endogenous state.

[28] For some recent examples, see Williams (1997), Fuhrer (1997a) and Orphanides, Small, Wilcox and Wieland (1997). An exception, however, is Jovanovic and Ueda (1997) who demonstrate that in an environment of incomplete contracting, increased dispersion of prices may reduce output. Stabilizing prices in this environment then raises output.

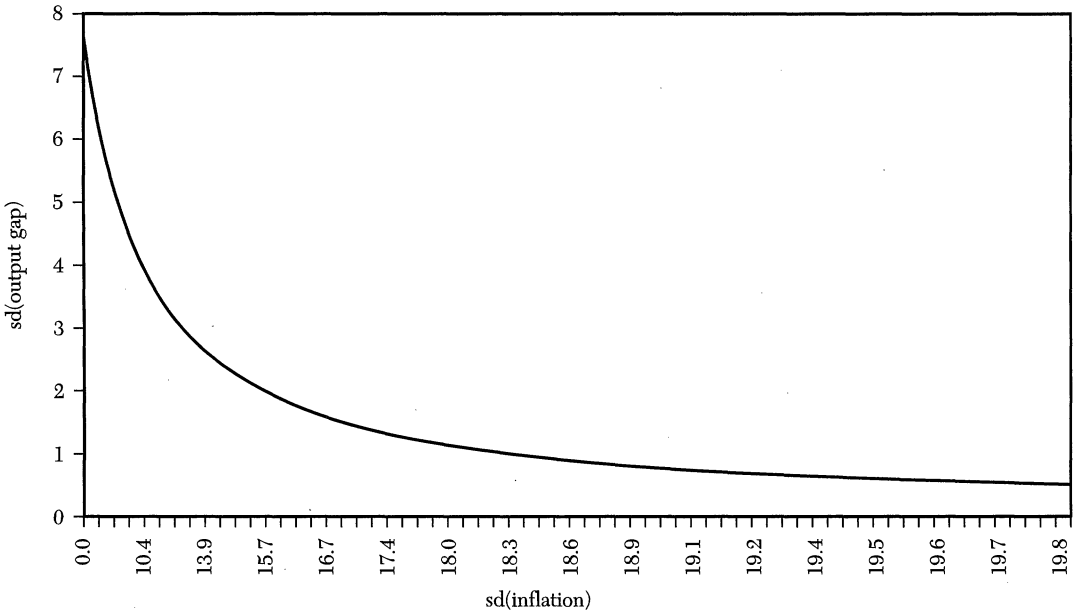


Figure 1. Efficient Policy Frontier for the Baseline Model

baseline model.[29] In (σ_x,σ_π) space the locus is downward sloping and convex to the origin. Points to the right of the frontier are inefficient. Points to the left are infeasible. Along the frontier there is a trade-off: As α rises (indicating relatively greater preference for output stability), the optimal policy engineers a lower standard deviation of output, but at the expense of higher inflation volatility. The limiting cases are instructive:

$$\text{As } \alpha \to 0: \quad \sigma_x = \frac{\sigma_u}{\lambda}; \quad \sigma_\pi = 0 \tag{3.7}$$

$$\text{As } \alpha \to \infty: \quad \sigma_x = 0; \quad \sigma_\pi = \frac{\sigma_u}{1-\beta\rho} \tag{3.8}$$

where σ_u is the standard deviation of the cost push innovation.

It is important to emphasize that the trade-off emerges only if cost push inflation is present. In the absence of cost inflation (i.e., with $\sigma_u = 0$), there is no trade-off. In this instance, inflation depends only on current and future demand. By adjusting interest rates to set $x_t = 0, \forall\, t$, the central bank is able to hit its inflation and output targets simultaneously, all the time. If cost push factors drive inflation, however, it is only possible to reduce inflation in the near term by contracting demand. This consideration leads to the next result:

Result 2: *The optimal policy incorporates inflation targeting in the sense that it requires to aim for convergence of inflation to its target over time. Extreme inflation targeting, however, i.e., adjusting policy to immediately reach an inflation target, is optimal under only one of two circumstances: (1) cost push inflation is absent; or (2) there is no concern for output deviations (i.e.,* $\alpha = 0$).

In the general case, with $\alpha > 0$ and $\sigma_u > 0$, there is gradual convergence of inflation back to target. From equations (3.5) and (2.4), under the optimal policy

$$\lim_{i \to \infty} E_t\{\pi_{t+i}\} = \lim_{i \to \infty} \alpha q \rho^i u_t = 0$$

[29] Equations (3.4) and (3.5) define the frontier for the baseline model.

In this formal sense, the optimal policy embeds inflation targeting.[30] With exogenous cost push inflation, policy affects the gap between inflation and its target along the convergent path, but not the rate of convergence. In contrast, in the presence of endogenous inflation persistence, policy will generally affect the rate of convergence as well, as we discuss later.

The conditions for extreme inflation targeting can be seen immediately from inspection of equations (3.7) and (3.8). When $\sigma_u = 0$ (no cost push inflation), adjusting policy to immediately hit the inflation target is optimal, regardless of preferences. Since there is no trade-off in this case, it is never costly to try to minimize inflation variability. Inflation being the only concern of policy provides the other rationale for extreme inflation targeting. As equation (3.7) indicates, it is optimal to minimize inflation variance if $\alpha = 0$, even with cost push inflation present.

Result 2 illustrates why some conflicting views about the optimal transition path to the inflation target have emerged in the literature. Marvin Goodfriend and Robert King (1997), for example, argue in favor of extreme inflation targeting. Svensson (1997a,b) and Laurence Ball (1997) suggest that, in general, gradual convergence of inflation is optimal. The difference stems from the treatment of cost push inflation: It is absent in the Goodfriend-King paradigm, but very much a factor in the Svensson and Ball frameworks.

Results 1 and 2 pertain to the behavior of the target variables. We now state several results regarding the behavior of the policy instrument, i_t.

Result 3: *Under the optimal policy, in response to a rise in expected inflation, nominal rates should rise sufficiently to increase real rates. Put differently, in the optimal rule for the nominal rate, the coefficient on expected inflation should exceed unity.*

Result 3 is transparent from equation (3.6). It simply reflects the implicit targeting feature of optimal policy described in Result 2. Whenever inflation is above target, the optimal policy requires raising real rates to contract demand. Though this principle may seem obvious, it provides a very simple criteria for evaluating monetary policy. For example, Clarida, Galí, and Gertler (forthcoming) find that U.S. monetary policy in the pre-Volcker era of 1960–79 violated this strategy. Federal Reserve policy tended to accommodate rather than fight increases in expected inflation. Nominal rates adjusted, but not sufficiently to raise real rates. The persistent high inflation during this era may have been the end product of the failure to raise real rates under these circumstances. Since 1979, however, the Federal Reserve appears to have adopted the kind of implicit inflation targeting strategy that equation (3.6) suggests. Over this period, the Fed has systematically raised real rates in response to anticipated increases in inflationary expectations. We return to this issue later.

Result 4: *The optimal policy calls for adjusting the interest rate to perfectly offset demand shocks, g_t, but perfectly accommodate shocks to potential output, z_t, by keeping the nominal rate constant.*

That policy should offset demand shocks is transparent from the policy rule (3.6). Here the simple idea is that countering demand shocks pushes both output and inflation in the right direction. Demand shocks do not force a

[30] Note here that our definition is somewhat different from Svensson (1997a), who defines inflation targeting in terms of the weights on the objective function, i.e., he defines the case with $\alpha = 0$ as corresponding to strict inflation targeting and $\alpha > 0$ as corresponding to flexible inflation targeting.

short run trade-off between output and inflation.

Shocks to potential output also do not force a short run trade-off. But they require a quite different policy response. Thus, e.g., a permanent rise in productivity raises potential output, but it also raises output demand in a perfectly offsetting manner, due to the impact on permanent income.[31] As a consequence, the output gap does not change. In turn, there is no change in inflation. Thus, there is no reason to raise interest rates, despite the rise in output.[32] Indeed, this kind of scenario seems to describe well the current behavior of monetary policy. Output growth was substantially above trend in recent times, but with no apparent accompanying inflation.[33] Based on the view that the rise in output may mainly reflect productivity movements, the Federal Reserve has resisted large interest rate increases.

The central message of Result 4 is that an important task of monetary policy is to distinguish the sources of business cycle shocks. In the simple environment here with perfect observability, this task is easy. Later we explore some implications of relaxing this assumption.

4. *Credibility and the Gains from Commitment*

Since the pioneering work of Kydland and Prescott (1977), Robert Barro and David Gordon (1983), and Rogoff (1985), a voluminous literature has developed on the issue of credibility of monetary policy.[34] From the standpoint of obtaining practical insights for policy, we find it useful to divide the papers into two strands. The first follows directly from the seminal papers and has received by far the most attention in academic circles. It emphasizes the problem of persistent inflationary bias under discretion.[35] The ultimate source of this inflationary bias is a central bank that desires to push output above its natural level. The second is emphasized more in applied discussions of policy. It focuses on the idea that disinflating an economy may be more painful than necessary, if monetary policy is perceived as not devoted to fighting inflation. Here the source of the problem is simply that wage and price setting today may depend upon beliefs about where prices are headed in the future, which in turn depends on the course of monetary policy.

These two issues are similar in a sense: They both suggest that a central bank that can establish credibility one way or another may be able to reduce inflation at lower cost. But the source of the problem in each case is different in subtle but important ways. As a consequence the potential empirical relevance may differ, as we discuss below.

We first use our model to exposit the famous inflationary bias result. We then illustrate formally how credibility can reduce the cost of maintaining low inflation, and also discuss mechanisms in

[31] In this experiment we are holding constant the IS shock g_t. Since $g_t = [(e_t - z_t) - E_t(e_{t+1} - z_{t+1})]$, (see footnote 9), this boils down to assuming either that the shock to z_t is permanent (so that $E_t z_{t+1} - z_t = 0$) or that e_t adjusts in a way to offset movements in g_t.

[32] That monetary policy should accommodate movements in potential GDP is a theme of the recent literature (e.g., Aiyagari and Braun 1997; Carlstrom and Fuerst 1995; Ireland 1996a; and Rotemberg and Woodford 1997). This view was also stressed in much earlier literature. See Friedman and Kuttner (1996) for a review.

[33] See Lown and Rich (1997) for a discussion of the recent "inflation puzzle."

[34] For recent surveys of the literature, see Fischer (1995), McCallum (1997) and Persson and Tabellini (1997).

[35] While the inflationary bias result is best known example, there may also be other costs of discretion. Svennson (1997c), for example, argues also that discretion may lead to too much inflation variability and too little output variability.

the literature that have been suggested to inject this credibility. An important result we wish to stress—and one that we don't think is widely understood in the literature—is that gains from credibility emerge even when the central bank is not trying to push output above its natural level.[36] That is, as long as price setting depends on expectations of the future, as in our baseline model, there may be gains from establishing some form of credibility to curtail inflation. Further, under certain plausible restrictions on the form of the feedback rule, the optimal policy under commitment differs from that under discretion in a very simple and intuitive way. In this case, the solution with commitment resembles that obtained under discretion using a higher effective cost applied to inflation than the social welfare function suggests.[37] In this respect, we think, the credibility literature may have some broad practical insights to offer.

4.1 *The Classic Inflationary Bias Problem*

As in Kydland and Prescott (1979), Barro and Gordon (1983), and many other papers, we consider the possibility that the target for the output gap may be $k > 0$, as opposed to 0. The policy objective function is then given by

$$\max -\frac{1}{2} E_t\left\{\sum_{i=o}^{\infty} \beta^i[\alpha(x_{t+i} - k)^2 + \pi_{t+i}^2]\right\} \quad (4.1)$$

The rationale for having the socially optimal level of output exceed its natural level may be the presence of distortions such as imperfect competition or taxes. For convenience, we also assume that price setters do not discount the future, which permits us to fix the parameter β in the Phillips curve at unity.[38]

In this case, the optimality condition that links the target variables is given by:

$$x_t^k = -\frac{\lambda}{\alpha}\pi_t^k + k \quad (4.2)$$

The superscript k indicates the variable is the solution under discretion for the case $k > 0$. Plugging this condition into the IS and Phillips curves, (2.1) and (2.2), yields:

$$x_t^k = x_t \quad (4.3)$$

$$\pi_t^k = \pi_t + \frac{\alpha}{\lambda} k \quad (4.4)$$

where x_t and π_t are the equilibrium values of the target variables for the baseline case with $k = 0$ (see equations 3.4 and 3.5).

Note that output is no different from the baseline case, but that inflation is systematically higher, by the factor $\frac{\alpha}{\lambda} k$. Thus, we have the familiar result in the literature:

Result 5. *If the central bank desires to push output above potential (i.e., $k > 0$), then under discretion a suboptimal equilibrium may emerge with inflation persistently above target, and no gain in output.*

The model we use to illustrate this

[36] A number of papers have shown that a disinflation will be less painful if the private sector perceives that the central bank will carry it out. But they do not show formally that, under discretion, the central bank will be less inclined to do so (see., e.g. Ball 1995, and Bonfim and Rudebusch 1997).

[37] With inflationary bias present, it is also possible to improve welfare by assigning a higher cost to inflation, as Rogoff (1985) originally emphasized. But it is not always possible to obtain the optimum under commitment. The point we emphasize is that with inflationary bias absent, it is possible to replicate the solution under commitment (for a restricted family of policy rules) using the algorithm to solve for the optimum under discretion with an appropriately chosen relative cost of inflation. We elaborate on these issues later in the text.

[38] Otherwise, the discounting of the future by price-setters introduces a long-run trade-off between inflation and output. Under reasonable parameter values this tradeoff is small and its presence merely serves to complicate the algebra. See Goodfriend and King (1997) for a discussion.

result differs from the simple expectational Phillips curve framework in which it has been typically studied. But the intuition remains the same. In this instance, the central bank has the incentive to announce that it will be tough in the future to lower current inflation (since in this case, current inflation depends on expected future inflation), but then expand current demand to push output above potential. The presence of k in the optimality condition (4.2) reflects this temptation. A rational private sector, however, recognizes the central bank's incentive. In mechanical terms, it makes use of equation (4.2) to forecast inflation, since this condition reflects the central bank's true intentions. Put simply, equilibrium inflation rises to the point where the central bank no longer is tempted to expand output. Because there is no long-run trade-off between inflation and output (i.e., x_t converges to zero in the long run, regardless of the level of inflation), long-run equilibrium inflation is forced systematically above target.

The analysis has both important positive and normative implications. On the positive side, the theory provides an explanation for why inflation may remain persistently high, as was the case from the late 1960s through the early 1980s. Indeed, its ability to provide a qualitative account of this inflationary era is a major reason for its popularity.

The widely stressed normative implication of this analysis is that there may be gains from making binding commitments over the course of monetary policy or, alternatively, making institutional adjustments that accomplish the same purpose. A clear example from the analysis is that welfare would improve if the central bank could simply commit to acting as if k were zero. There would be no change in the path of output, but inflation would decline.

Imposing binding commitments in a model, however, is much easier than doing so in reality. The issue then becomes whether there may be some simple institutional mechanisms that can approximate the effect of the idealized policy commitment. Perhaps the most useful answer to the question comes from Rogoff (1985), who proposed simply the appointment of a "conservative" central banker, taken in this context to mean someone with a greater distaste for inflation (a lower α), than society as a whole:

Result 6: *Appointing a central bank chair who assigns a higher relative cost to inflation than society as a whole reduces the inefficient inflationary bias that is obtained under discretion when* $k > 0$.

One can see plainly from equation (4.4) that letting someone with preferences given by $\alpha^R < \alpha$ run the central bank will reduce the inflationary bias.[39] The Rogoff solution, however, is not a panacea. We know from the earlier analysis that emphasizing greater reduction in inflation variance may come at the cost of increased output variance. Appointing an extremist to the job (someone with α at or near zero) could wind up reducing overall welfare.

How important the inflationary bias problem emphasized in this literature is in practice, however, is a matter of controversy. Benjamin Friedman and Kenneth Kuttner (1996) point out that inflation in the major OECD countries now appears well under control, despite the absence of any obvious institutional changes that this literature argues are needed to enhance credibility. If this theory is robust, they argue, it should account not only for the high inflation of the 1960s and 1970s, but also for the

[39] See Svensson (1997) and Walsh (1998) for a description of how incentive contracts for central bankers may reduce the inflation bias; also, Faust and Svensson (1998) for a recent discussion of reputational mechanisms.

transition to an era of low inflation during the 1980s and 1990s. A possible counterargument is that in fact a number of countries, including the U.S., effectively adopted the Rogoff solution by appointing central bank chairs with clear distaste for inflation.

Another strand of criticism focuses on the plausibility of the underlying story that leads to the inflationary bias. A number of prominent authors have argued that, in practice, it is unlikely that $k > 0$ will tempt a central bank to cheat. Any rational central bank, they maintain, will recognize the long-term costs of misleading the public to pursue short-term gains from pushing output above its natural level. Simply this recognition, they argue, is sufficient to constrain its behavior (e.g. McCallum 1997a; Blinder 1997). Indeed, Blinder argues, based on his own experience on the Federal Reserve Board, that there was no constituency in favor of pursuing output gains above the natural rate. In formal terms, he maintains that those who run U.S. monetary policy act as if they were instructed to set $k = 0$, which eliminates the inflationary bias.

What is perhaps less understood, however, is that there are gains from enhancing credibility even when $k = 0$. To the extent that price setting today depends on beliefs about future economic conditions, a monetary authority that is able to signal a clear commitment to controlling inflation may face an improved short-run output/inflation trade-off. Below we illustrate this point. The reason why this is not emphasized in much of the existing literature on this topic is that this work either tends to focus on steady states (as opposed to short-run dynamics), or it employs very simple models of price dynamics, where current prices do not depend on beliefs about the future. In our baseline model, however, short-run price dynamics depend on expectations of the future, as equation (2.2) makes clear.[40]

4.2 *Improving the Short-Run Output/Inflation Trade-off: Gains from Commitment with $k = 0$.*

We now illustrate that there may be gains from commitment to a policy rule, even with $k = 0$. The first stage problem in this case is to choose a state contingent sequence for x_{t+i} and π_{t+i} to maximize the objective (2.7) assuming that the inflation equation (2.2) holds in every period $t + i, i \geq 0$. Specifically, the central bank no longer takes private sector expectations as given, recognizing instead that its policy choice effectively determines such expectations.

To illustrate the gains from commitment in a simple way, we first restrict the form of the policy rule to the general form that arises in equilibrium under discretion, and solve for the optimum within this class of rules. We then show that, with commitment, another rule within this class dominates the optimum under discretion. Hence this approach provides a simple way to illustrate the gains from commitment. Another positive byproduct is that the restricted optimal rule we derive is simple to interpret and implement, yet still yields gains relative to the case of discretion. Because the policy is not a global optimum, however, we conclude the section by solving for the unrestricted optimal rule.

4.2.1 *Monetary Policy under Commitment: The Optimum within a Simple Family of Policy Rules (that includes the optimal rule under discretion)*

In the equilibrium without commitment, it is optimal for the central bank

[40] This section is based on Galí and Gertler (1999).

to adjust x_t solely in response to the exogenous cost push shock, u_t. We accordingly consider a rule for the target variable xt that is contingent on the fundamental shock u_t, in the following way:

$$x_t^c = -\omega u_t \tag{4.5}$$

for all t, where $\omega > 0$ is the coefficient of the feedback rule, and where x_t^c denotes the value of x_t conditional on commitment to the policy.[41] Note that the rule includes the optimum under discretion as a special case (i.e., the case with $\omega = \lambda q$ shown in 3.4).

Combining equation (4.5) with the Phillips curve (2.2), in turn, implies that inflation under the rule, π_t^c, is also a linear function of the cost push shock:

$$\pi_t^c = \lambda x_t^c + \beta E_t \pi_{t+1}^c + u_t \tag{4.6}$$

$$= E_t \sum_{i=0}^{\infty} \beta^i [\lambda x_{t+i}^c + u_{t+i}] \tag{4.7}$$

$$= E_t \sum_{i=0}^{\infty} \beta^i [-\lambda\omega u_{t+i} + u_{t+i}] \tag{4.8}$$

$$= \frac{1-\lambda\omega}{1-\beta\rho} u_t \tag{4.9}$$

The problem for the central bank is to choose the optimal value of the feedback parameter ω. Relative to the case of discretion, the ability to commit to a feedback policy provides the central bank with an improved short-run output/inflation trade-off. To this end, note that it is possible to express equation (4.9) as

$$\pi_t^c = \frac{\lambda}{1-\beta\rho} x_t^c + \frac{1}{1-\beta\rho} u_t \tag{4.10}$$

In this case, a one percent contraction in x_t^c reduces π_t^c by the factor $\frac{\lambda}{1-\beta\rho}$. Under discretion, reducing x_t by one percent only produces a fall in π_t of $\lambda < \frac{\lambda}{1-\beta\rho}$. The extra kick in the case with commitment is due to the impact of the policy rule on expectations of the future course of the output gap. In particular, the choice of ω affects not only x_t but also beliefs about the course of x_{t+i}^c, $i = 1,2, \ldots$, since $E_t x_{t+i}^c = -\omega u_t$. A central bank that commits to a tough policy rule (high ω), for example, is able to credibly signal that it will sustain over time an aggressive response to a persistent supply shock. Since inflation depends on the future course of excess demand, commitment to the tough policy rule leads to a magnified drop in inflation per unit of output loss, relative to the case of discretion.

To find the optimal value of ω, note first that since x_{t+i}^c and π_{t+i}^c are each a constant multiple of the cost push shock u_{t+i}, it is possible to express the objective function as a multiple of period t loss:

$$\max -\frac{1}{2} E_t \left\{ \sum_{i=0}^{\infty} \beta^i [\alpha (x_{t+i}^c)^2 + (\pi_{t+i}^c)^2] \right\}$$

$$\longleftrightarrow \max -\frac{1}{2} [\alpha (x_t^c)^2 + (\pi_t^c)^2] L_t \tag{4.11}$$

with $L_t \equiv E_t \left\{ \sum_{i=0}^{\infty} \beta^i (u_{t+i}/u_t)^2 \right\} > 0$. The problem then is to choose ω to maximize (4.11), subject to (4.10). In this instance, the optimality condition is given by:

$$x_t^c = -\frac{\lambda}{\alpha^c} \pi_t^c \tag{4.12}$$

where

$$\alpha^c \equiv \alpha(1-\beta\rho) < \alpha \tag{4.13}$$

Since $\alpha^c < \alpha$, relative to the case of discretion, commitment to the rule implies that it is optimal for the central bank to engineer a greater contraction in output in response to inflationary pressures. Intuitively, the more aggressive response

[41] The policy rule only depends on u_t because the central bank can adjust i_t to offset any impact of movements in g_t on aggregate demand. See equation (4.16).

to inflation is the product of the improved output/inflation trade-off that commitment affords. Specifically, the output cost of lowering inflation declines from α to α^c per unit, since reducing inflation a given amount requires, ceteris paribus, only a fraction $(1 - \beta\rho)$ of the output loss required under discretion. The decline in the effective cost of reducing inflation, in turn, induces the more aggressive policy response to inflation, as comparing equation (4.12) with equation (3.3) makes clear.

The equilibrium solutions for x_t^c and π_t^c are easily obtained by combining equations (4.12) and (4.10):

$$x_t^c = -\lambda q^c u_t \qquad (4.14)$$

$$\pi_t^c = -\alpha^c q^c u_t \qquad (4.15)$$

with

$$q^c = \frac{1}{\lambda^2 + \alpha^c(1 - \beta\rho)}$$

It is interesting to observe that the solution under commitment in this case perfectly resembles the solution obtained under discretion that arises when α is replaced with $\alpha^c < \alpha$ in the objective function. It follows that, conditional on the value of the cost push shock, u_t, inflation is closer to target and output is further, relative to the outcome under discretion.[42]

It is straightforward to verify that commitment to the policy rule raises welfare.[43] The tension produced by such gains from commitment, we think, is compelling from an empirical standpoint. Because inflation depends on expected future output gaps, the central bank would like to convince the private sector that it will be tough in the future, but at the same time, not to have to contract demand much today. As the future comes to pass, the central bank has the incentive to renege on its planned toughness and, instead, promise again to undertake contractionary policy down the road. To see this, suppose that there is a positive cost push shock. If the central bank is free to deviate from the rule, it will always choose the optimal policy under discretion, which calls for a smaller contraction of output, relative to the case of commitment (again, compare 4.1 and 3.3). A rational private sector will recognize that incentive and, unless the central bank is able to commit credibly, will not expect large contractions in demand in the future either. As a result, the cost push shock generates higher inflation in the absence of commitment. We stress again that, in contrast to the traditional analysis, this gain from commitment is not tied to the desire of the central bank to push output above potential, but to the forward-looking nature of inflation (and, thus, the importance of expectations about future policy) in our baseline model.

From a policy standpoint, Rogoff's rationale for a conservative central banker carries over perfectly to this case. Indeed with omniscience (i.e. exact knowledge of α^c and the true model), an appropriately chosen central banker could replicate the outcome under commitment.

We summarize the findings in Result 7:

Result 7: *If price-setting depends on expectations of future economic conditions, then a central bank that can credibly commit to a rule faces an improved short-run trade-off between inflation*

[42] Importantly, with endogenous inflation persistence, commitment produces a faster transition of inflation to target, as we discuss later.

[43] To verify that commitment raises welfare, simply substitute the implied values for x_t and π_t under the optimal rule for each case into the policy objective function. However, it should be obvious that commitment raises welfare, since the optimal rule under discretion falls within the class of rules that we permitted the central bank to choose in the case with commitment: Yet we found that with commitment it is optimal to choose a different parameterization of the rule than arises in the optimum under discretion.

and output. This gain from commitment arises even if the central bank does not prefer to have output above potential (i.e., even when $k = 0$). The solution under commitment in this case perfectly resembles the solution that would obtain for a central bank with discretion that assigned to inflation a higher cost than the true social cost.

One additional interesting feature of this case with commitment involves the behavior of interest rates. This can be seen formally by simply replacing α with α^c in the interest rate rule under discretion (given by equation 3.6) to obtain

$$i_t = \gamma_\pi^c E_t \pi_{t+1} + \frac{1}{\varphi} g_t \qquad (4.16)$$

with

$$\gamma_\pi^c \equiv 1 + \frac{(1-\rho)\lambda}{\rho\varphi\alpha^c} > 1 + \frac{(1-\rho)\lambda}{\rho\varphi\alpha} \equiv \gamma_\pi$$

In particular, relative to the case of discretion, the central bank increases the nominal interest rate by a larger amount in response to a rise in expected inflation.

4.2.2 *Monetary Policy under Commitment: The Unconstrained Optimum*

We now provide a brief description of the general solution for the optimal policy under commitment.[44] Because the derivation is more cumbersome than for the restricted case just described, we defer most of the details to an appendix. As with the simple fundamental based policy, however, the general solution exploits the ability that commitment affords to manipulate private sector expectations of the future.

The first stage problem remains to choose a state-contingent sequence for x_{t+i} and π_{t+i} to maximize the objective (2.7) given that the aggregate supply curve (2.2) holds in every period $t+i, i \geq 0$. We no longer restrict the choice of x_t to depend on the contemporaneous value of the shock (i.e., u_t), but allow instead for rules that are a function of the entire history of shocks. To find the globally optimal solution to the linear quadratic policy problem under commitment, we follow David Currie and Paul Levine (1993) and Woodford (1998), and form the Lagrangian:[45]

$$\max -\frac{1}{2} E_t \Big\{ \sum_{i=0}^{\infty} \beta^i [\alpha x_{t+i}^2 + \pi_{t+i}^2 + \phi_{t+i}(\pi_{t+i} - \lambda x_{t+i} - \beta\pi_{t+i+1} - u_{t+i})] \Big\} \qquad (4.17)$$

where $\frac{1}{2}\phi_{t+i}$ is the (state-contingent) multiplier associated with the constraint at $t+i$. It is straightforward to show that the first order conditions yield the following optimality conditions

$$x_{t+i} - x_{t+i-1} = -\frac{\lambda}{\alpha}\pi_{t+i}, \quad \text{for } i = 1,2,3,... \qquad (4.18)$$

and

$$x_t = -\frac{\lambda}{\alpha}\pi_t \qquad (4.19)$$

Recall that under discretion the optimal policy has the central bank adjust the level of the output gap in response to inflation. The optimal policy under commitment requires instead adjusting the *change* in the output gap in response to inflation. In other words, commitment changes the level rule for x_t under discretion into a difference rule for x_t, as a comparison of equations

[44] We thank Chris Sims and Albert Marcet for calling to our attention that the globally optimal rule under commitment would likely not fall within the restricted family of rules considered in the previous sub-section.

[45] See also King and Wolman, who analyze the optimal monetary policy under commitment in a version of Taylor's (1980) staggered contracts model.

(3.3) and (4.8) indicates.[46] The one caveat is that in the initial period the policy is implemented (i.e., period t) the central bank should simply adjust the level of the output gap x_t is response to π_t, *as if* it were following the optimal policy under discretion, but for that period only.

Because x_{t+i} depends in general on x_{t+i-1}, the (unconstrained) optimal policy under commitment is in general not simply a function of the contemporaneous state variable u_{t+i}. As Woodford (1998) emphasizes in a related context, the lagged dependence in the policy rule arises as a product of the central bank's ability under commitment to directly manipulate private sector expectations.[47] To see this for our framework, keep in mind that π_t depends not only on current x_t but also on the expected future path of x_{t+i}. Then suppose, for example, that there is a cost push shock that raises inflation above target at time t. The optimal response under discretion, as we have seen, is to reduce x_t, but then let x_{t+i} revert back to trend over time as π_{t+i} falls back to target. The optimal policy under commitment, however, is to continue to reduce x_{t+i} as long as π_{t+i} remains above target. The (credible) threat to continue to contract x_t in the future, in turn, has the immediate effect of dampening current inflation (given the dependency of π_t on future values of x_t). Relative to the case of discretion, accordingly, the cost push shock has a smaller impact in current inflation.[48]

As with the constrained policy, the globally optimal policy under commitment exploits the ability of the central bank to influence π_t with expected future values of x_{t+i} as well as current x_t. It is also easy to see that, as was the case with the more restrictive rule, the policy is not time consistent. Clearly, if it could reoptimize at $t+i$, the central bank would choose the same policy it implemented at t, the one which mimics the rule under discretion for the first period only.

A disadvantage of the unconstrained optimal policy under commitment is that it appears more complex to implement than the constrained one (described by equation 4.12). As we have seen, the constrained rule resembles in every dimension the optimal policy under discretion, but with relatively more weight placed on fighting inflation. Accordingly, as we discussed, it is possible to approximate this policy under discretion with an appropriately chosen central banker. The same is not true, however, for the unconstrained optimal policy. A conservative central banker operating with discretion has no obvious incentive to stick to the difference rule for the output gap implied by equation (4.18).

A further complication, discussed at length in Woodford (1998), is that the interest rate rule that implements the optimal policy might have undesirable

[46] Woodford (1998) makes the connection between the lagged dependence in the optimal rule under commitment and the lagged dependence that appears to arise in interest rate behavior under practice (see section 5.2). Roughly speaking, since the interest rate affects the output gap, lagged dependence in the latter translates into lagged dependence in the former.

[47] Woodford (1998) considers a closely related environment. The difference is that in his framework the policy-maker confronts a trade-off between inflation and the output gap ultimately because his objective function includes a target for the nominal interest rate (along with targets for the output gap and inflation), whereas in our framework the trade-off arises due to the cost push shock.

[48] On the surface it appears that the difference rule for x_t might be unstable. However, π_t adjusts to ensure that this is not the case. In particular, the optimal response to a positive cost push shock is to contract x_t sufficiently to push π_t below target. x_t then adjusts back up to target over time. The appendix provides the details.

side effects. To see this, combine (4.18) and (2.1) to obtain the implied optimal interest rate rule

$$i_t = \left(1 - \frac{\lambda}{\alpha\varphi}\right) E_t\pi_{t+1} + \frac{1}{\varphi} g_t$$

Notice that the coefficient associated with expected inflation is less than one. Under this rule, accordingly, a rise in anticipated inflation leads to a decline in the real interest rate. As we discuss in section 7, if inflationary pressures vary inversely with the real rate, a rule of this type may permit self-fulfilling fluctuations in output and inflation that are clearly suboptimal.[49]

Overall, we have:

Result 8: *The globally optimal policy rule under commitment has the central bank partially adjust demand in response to inflationary pressures. The idea is to exploit the dependence of current inflation on expected future demand. In addition, while appointing a conservative central banker may raise welfare under discretion (see Result 7), it does not appear that it is possible to attain the globally optimal rule with this strategy. Finally, there may be some practical complications in implementing the globally optimal interest rate rule that involve potential indeterminacy, as discussed in Woodford (1998).*

We conclude that, though substantial progress has been made, our understanding of the full practical implications of commitment for policy-making is still at a relatively primitive stage, with plenty of territory that is worth exploring.

[49] Indeterminacy does not arise in the case of discretion or in the case of the constrained optimum under commitment, since in each instance the implied interest rate rule has an inflation coefficient greater than one. To the extent that such coefficient is not too large, implementation of such a rule will result in a unique equilibrium (see the discussion in section 7 and also in Clarida, Galí, and Gertler (1998).

5. *Practical Complications*

In this section we consider a number of important practical issues that complicate the implementation of monetary policy. While they may not be as exotic as the question of credibility, they are no less important for the day-to-day formulation of policy.

5.1 *Imperfect Information*

Thus far we have assumed that the central bank is able to control perfectly the paths of the key target variables. In practice, of course, this is not the case. One important reason is imperfect observability. At the time it sets interest rates, a central bank may not have all the relevant information available about the state of the economy. Certain data take time to collect and process. Sampling is imperfect. Even if it has access to data in real time, some key variables such as the natural level of output are not directly observable and are likely measured with great error (see, e.g., the discussion in Arturo Estrella and Mishkin 1999 and Orphanides 1998).

Beyond limiting the efficacy of policy, imperfect information has several specific implications. First, it is no longer possible to specify rules simply in terms of target variables. With perfect information, a policy may be expressed equivalently in terms of targets or instruments since a one-to-one relationship generally exists between these variables. With imperfect information, rules for targets can be expressed only in terms of the respective forecasts, as opposed to the ex-post values. An alternative is to use an intermediate target that is directly observable, such as a broad monetary aggregate.

Second, imperfect information makes the policy instrument choice non-trivial. With perfect information, for example, it does not matter whether the central

bank uses the short-term interest rate or a monetary aggregate as the policy instrument, so long as the money demand function yields a monotonic relation between the two variables.[50] With imperfect information, the ex post volatility of a variety of key variables hinges on the instrument choice, as originally argued by William Poole (1970). We illustrate each of these issues below.[51]

5.1.1 *Forecasts as Targets and Intermediate Targets*

We now return to the baseline model with no commitment, and modify it as follows. Suppose that the central bank cannot observe the contemporaneous values of output, inflation, or any of the random shocks. Then let Ω_t be the central bank's information set at the time it fixes the interest rate that prevails at time t.[52] The optimality condition for policy now is expressed in terms of the expected as opposed to realized target variables.

$$E[x_t \mid \Omega_t] = -\frac{\lambda}{\alpha} E[\pi_t \mid \Omega_t] \qquad (5.1)$$

Equation (5.1) is the certainty equivalent version of the condition for the case of perfect information, given by equation (3.3). Certainty equivalence applies here because of the linear quadratic setup (that gives linear decision rules under perfect information) and because the errors in forecasting the target variables are additive.

For ease of exposition, assume that there is no serial correlation in the cost push shock; that is, $\rho = 0$, so that $u_t = \hat{u}_t$. The implied equilibrium values of the target variables under imperfect information, x_t^I and π_t^I, are given by

$$x_t^I = x_t + \left(\frac{\lambda}{\lambda^2 + \alpha}\hat{u}_t + \hat{g}_t\right) = \hat{g}_t \qquad (5.2)$$

$$\pi_t^I = \left(1 + \frac{\lambda^2}{\alpha}\right)\pi_t + \lambda\hat{g}_t = \hat{u}_t + \lambda\hat{g}_t \qquad (5.3)$$

where x_t and π_t are the optimal values of the target variables that emerge in case of perfect information (when u_t is serially uncorrelated),[53] and where $\hat{u}_t$ and $\hat{g}_t$ are the unexpected movements in the cost push and demand shocks, respectively. Imperfect information clearly implies greater volatility of inflation, since the central bank cannot immediately act to offset the impact of the shocks. The net effect on the volatility of the output gap is unclear: the inability to offset the demand shock clearly raises output volatility. On the other hand, the central bank cannot offset the inflationary impact of the cost push shock, which works to reduce the volatility of the output. There is, however, an unambiguous reduction in welfare.[54]

One additional result is worth noting. Since demand shocks now affect the behavior of output, a positive short-run co-movement between inflation and output can emerge if $\hat{g}_t$ has a variance sufficiently large relative to that of $\hat{u}_t$.

It is straightforward to generalize the analysis to a setting where the imperfect observability stems from lags in the

[50] To clarify, a money aggregate can serve as an instrument only if it is directly controllable. A candidate aggregate then would be bank reserves. A broad aggregate such as M3 would not qualify.

[51] For a broad survey of the literature on monetary policy targets and instruments, see Friedman (1991).

[52] Thus, Ω_t is similarly the private sector's information set. Specifically, we let firms observe the current values of their marginal costs, but neither firms nor households can observe contemporaneous aggregate variables. In this instance, the IS and Phillips curve equations are respectively given by

$$x_t = -\varphi[(i_t \mid \Omega_t) - E_{t-1}\pi_{t+1}] + E_{t-1}x_{t+1} + g_t$$

$$\pi_t = \lambda x_t + \beta E_{t-1}\pi_{t+1} + u_t$$

[53] When u_t is serially uncorrelated, $x_t = \frac{-\lambda}{\lambda^2 + \alpha} u_t$ and $\pi_t = \frac{\alpha}{\lambda^2 + \alpha} u_t$.

[54] To prove that imperfect information leads to a reduction in welfare, evaluate the welfare function with x_t^I and π_t^I versus x_t and π_t.

transmission of monetary policy. This case is of interest since much of the available evidence suggests a lag of six to nine months in the effect of a shift in interest rates on output.[55] The lag in the effect on inflation is around a year and a half. Suppose, for example, that it takes j periods for a shift in the current interest rate to affect output and another k periods for an impact on inflation. In the left side of equation (5.1) would appear the j period ahead forecast of the output gap, and on the right would be the (suitably discounted) $j+k$ period ahead forecast of inflation.

Svensson (1997a,b) has emphasized the practical importance of this result for the mechanics of inflation targeting (specifically, the kind of inflation targeting that the theory implies (see Result 2 in section 3). A standard criticism of employing an inflation target is that information about the impact of current monetary policy on inflation is only available with a long lag. This information lag, it is argued, makes it impossible to monitor policy performance. It is possible to circumvent this problem, according to Svensson, by focusing instead on the inflation forecast. The forecast is immediately available. It thus provides a quick way to judge the course of policy. A caveat to this argument is that to generate the correct inflation forecast, the central bank must have a good structural model of the economy.[56] VAR-based forecasts are reasonable only if the economy has attained a stationary equilibrium.

A traditional alternative to using the target variable forecasts is to focus on the behavior of a variable that is correlated with the underlying targets but is instead observable and controllable. Broad monetary aggregates are the best known examples of intermediate targets. If demand for a particular aggregate is stable, then this aggregate is likely to have a stable covariance with nominal GDP. In practice, however, experience with monetary targeting has not been successful. The U.S. and the U.K., for example, attempted to regulate the growth of money aggregates in the early 1980s and then quickly abandoned the policy after the aggregates went haywire.[57] Financial innovation in each instance was the underlying culprit. Even in Germany, long considered a bastion of money targeting, there have been problems. Unstable movements in money demand have forced a retreat from strict money growth targeting. A number of recent papers go further by arguing that in practice Bundesbank policy looks more like inflation targeting (as defined in Result 2) than money targeting (Clarida and Gertler 1997; Bernanke and Mihov 1997b).

For similar reasons, policies that target other kinds of simple indicators, such as commodity prices or long term interest rates, have not been widely employed. As Woodford (1994a) has emphasized, the correlation properties of these simple indicators with output and

[55] Galí (1992), Christiano, Eichenbaum, and Evans (1996), and Bernanke and Mihov (1997a) document the slow response of GDP to a policy shock, and the even slower response of prices. Bernanke and Gertler (1995) show that, while the overall response of output is sluggish, certain components of spending do respond quickly, such as housing and consumer durables. Inventories adjust to reconcile the gap between spending and output.

[56] Bernanke and Woodford (1997) emphasize the need to make structural forecasts. They also raise some other related criticisms of using forecast-based targets, including the possibility of indeterminacy under this kind of policy rule. We discuss this issue in section 7.

[57] See Friedman and Kuttner (1996) for a detailed accounting of the failure of monetary targeting to take hold in the U.S. See also Estrella and Mishkin (1996). On the other hand, Feldstein and Stock (1997) argue that, with periodic adjustment, a broad monetary aggregate can still be a useful intermediate target.

inflation is likely to vary with changes in the policy rule. In the end, there is no simple substitute for employing a structural model.

To summarize, we have

Result 9: *With imperfect information, stemming either from data problems or lags in the effect of policy, the optimal policy rules are the certainty equivalent versions of the perfect information case. Policy rules must be expressed in terms of the forecasts of target variables as opposed to the ex post behavior. Using observable intermediate targets, such as broad money aggregates is a possibility, but experience suggests that these indirect indicators are generally too unstable to be used in practice.*

5.1.2 *The Instrument Choice Problem: The Interest Rate versus a Narrow Monetary Aggregate*

We now turn to the issue of instrument choice. In practice, the interest rate that major central banks adjust is an overnight rate on interbank lending of funds to meet reserve requirements.[58] They control this rate by manipulating the supply of bank reserves, i.e., the quantity of high-powered money available for meeting bank reserve requirements. The issue that arises is whether, from an operational standpoint, policy should prescribe paths (or rules) for bank reserves or for interest rates. Suppose that the demand for bank reserves m_t is given by[59]

$$m_t - p_t = \kappa\, y_t - \eta\, i_t + v_t \qquad (5.4)$$

where p_t is the price level and v_t is a random disturbance to money demand. If v_t is perfectly observable then it does not matter whether i_t or m_t is employed as the policy instrument. Given the time path of i_t implied by the optimal policy, it is possible to back out a time path for m_t that supports this policy from equation (5.4).

Matters change if v_t is not observable. With the interest rate as the instrument, the central bank lets the money stock adjust to the money demand shock. There is no impact of money demand shocks on output or inflation because the central bank perfectly accommodates them. With money targeting, the reverse is true: the interest rate and (possibly) output adjust to clear the money market. Assume for simplicity that demand and cost push shocks are absent (i.e., $g_t = 0$, $u_t = 0$), so that the only shock is the innovation to money demand. Then the interest rate implied by a money supply instrument i_t^m, is given by

$$i_t^m = i_t + \frac{1}{\eta + \varphi(\kappa + \lambda)}\,\hat{v}_t \qquad (5.5)$$

where i_t is the rate that would arise under interest rate targeting and $\hat{v}_t$ is the unexpected movement in money demand.

The key point is that money demand shocks can induce volatile behavior of interest rates. This is particularly true if money demand is relatively interest inelastic in the short run, as is the case for bank reserves. This short run volatility in interest rates will then feed into output volatility, via the aggregate demand relation, equation (2.1). It is for this reason that in practice central banks use interbank lending rates as the policy instrument, an insight due originally to Poole (1970).[60] Recent empirical

[58] See Bernanke and Mihov (1997a) for a discussion of Federal Reserve operating procedures and how they have changed over time.

[59] In the optimizing IS/LM framework of section 2, it is possible to motivate this specification of the money demand function from first principles, assuming that utility is separable in consumption and real money balances and that consumption is the only type of good (see, e.g., Woodford 1996).

[60] Poole also argued that if unobservable demand shocks were large relative to money demand shocks, then it may be preferable to use a money

work by Bernanke and Mihov (1997a) confirms that except for the brief period of non-borrowed reserve targeting under Volcker (1979:10–1982:10), the Federal Reserve Board has indeed treated the Funds rate as the policy instrument. In summary, we have

Result 10: *Large unobservable shocks to money demand produce high volatility of interest rates when a monetary aggregate is used as the policy instrument. It is largely for this reason that an interest rate instrument may be preferable.*

The analysis thus makes clear why the new Federal Reserve Board model does not even bother to include a money aggregate of any form (see Flint Brayton et al. 1997). Narrow aggregates are not good policy instruments due to the implied interest rate volatility. Broad aggregates are not good intermediate targets because of their unstable relation with aggregate activity.

5.2 *Policy Conservatism: Model Uncertainty vs. Exploitation of Forward-Looking Behavior*

In practice, central banks adjust interest rates more cautiously than standard models predict. Put differently, optimal policies derived in a certainty equivalent environment generally predict a much more variable path of interest rates than is observed in practice. An interesting illustration of this point is Rotemberg and Woodford (1997) who estimate a model very similar to our baseline model, and then compute an optimal interest rate policy. The historical interest rate displays much less volatility than the optimal interest rate. This finding is not uncommon. The FRB-US model also generates high interest rate volatility under an optimal rule. Because this degree of volatility seems greater than monetary policy makers seem willing to tolerate in practice, optimal rules are also computed with constraints on the volatility of interest rate changes (see, e.g., John Williams 1997).[61]

The tendency of the Federal Reserve to adjust rates cautiously is generally referred to as "interest rate smoothing." To be precise, as a number of authors have shown, a monetary policy rule of the following form captures the last twenty or so years of data fairly well:

$$i_t = (1-\rho)[\alpha + \beta\,\pi_t + \gamma\,x_t] + \rho\,i_{t-1} + \varepsilon_t \quad (5.6)$$

where α is a constant interpretable as the steady state nominal interest rate[62] and where $\rho \in [0,1]$ is a parameter that reflects the degree of lagged dependence in the interest rate.[63] Interest rate smoothing is present in distinct respects. First, the estimated slope coefficients on inflation and the output gap, β and γ, are typically smaller than what the optimal rule would suggest. Second, there is typically partial adjustment to movements in π_t and x_t, reflected by the presence of the lagged interest in the fitted rule. That is, i_t is a weighted average of some desired value that depends on the state economy (given by the term $[\alpha + \beta\pi_t + \gamma x_t]$) and the lagged interest rate, where the relative weights depend on the smoothing parameter ρ. Estimates of ρ for quarterly data are typically on the order of 0.8 or 0.9, which suggests very slow adjustment

supply instrument. With a money supply instrument, interest rates will naturally move in an offsetting direction in response to unobserved demand shocks (see Result 4). In practice, the high variability of money demand shocks seems to dominate the instrument choice, however.

[61] An alternative is to penalize large changes in the nominal interest rate by including the squared deviations of the change in the interest rate (i.e, $(i_t - i_{t-1})^2$) in the function, as in Rudebusch and Svensson (1998).

[62] Recall that π_t represents deviations of inflation from its average (target) level.

[63] See Rudebusch (1995), for example, for a discussion of the persistence in short term interest rates.

in practice. The existing theory, by and large, does not readily account for why the central bank should adjust rates in such a sluggish fashion.

Indeed, understanding why central banks choose a smooth path of interest rates than theory would predict is an important unresolved issue. One implication is that the standard certainty equivalence models may not adequately capture the constraints policy-makers face in practice. A natural possibility is that policy-makers know far less about the way the world works than is presumed in simple policy experiments.

In general, model uncertainty is a formidable problem. Ideally, one would like to take into account that the central bank is continually learning about the economy as it adjusts its policy. Performing this exercise in a clean way is beyond the frontier of current knowledge. Though, advances in computational methodology have allowed some progress to be made with relatively simple frameworks.[64]

It is possible to illustrate how model uncertainty could in principle introduce at least some degree of policy caution. Suppose the values of several parameters in the model are random. The central bank knows the distribution of these parameters but not the realization. When it adjusts policy, accordingly, it cannot be sure of the impact on the economy. As originally demonstrated by William Brainard (1969), this kind of uncertainty can introduce caution in policy responses. In contrast to the case of certainty-equivalence, policy actions now affect the conditional variance of inflation and output, as well as the conditional mean.

To be concrete, suppose that the two parameters of the model, the interest elasticity in the IS equation and the slope coefficient on the output gap are random variables, now given by $\tilde{\varphi}_t = \varphi + \varepsilon_t$ and by $\tilde{\lambda}_t = \lambda + \eta_t$.[65] Assume further that ε_t and η_t are i.i.d random variables with zero means. The optimality condition for policy then becomes:

$$E\{x_t \mid \Omega_t\} = \frac{\lambda}{\alpha + \lambda^2\sigma_\eta^2} E\{\pi_t \mid \Omega_t\} + (\alpha + \lambda^2)\frac{\sigma_\varepsilon^2}{\varphi} r_t \quad (5.7)$$

where $r_t \equiv i_t - E\{\pi_{t+1} \mid \Omega_t\}$ is the ex ante real interest rate. This condition leads to the following result:

Result 11: *Parameter uncertainty may reduce the response of the policy instrument to disturbances in the economy. It can thus motivate a smoother path of the interest rate than the certainty equivalent policy implies.*

Comparing equations (5.1) and (5.7) reveals how parameter uncertainty reduces policy activism. Under certainty equivalence, a rise in inflation above target requires the central bank to raise interest rates to contract demand.[66]

With an uncertain slope coefficient on the output gap in the AS curve, however, contraction of output below potential raises the variability of inflation. This induces the central bank to moderate the contraction in demand, as reflected by the presence of the term $\lambda^2\sigma_\eta^2$ in the coefficient on $E\{\pi_t \mid \Omega_t\}$. Similarly,

[64] Wieland (1997) analyzes policy in a framework where the central bank has to learn the value of the natural rate of unemployment (which, in our analysis, corresponds to having to learn about potential GDP.)

[65] We are assuming that the policy-maker knows the first two moments of the random parameters. It may be more plausible to argue that the policy-maker in fact has little idea what the true distribution looks like. See Onatski and Stock (1999) who analyze the policy problem in this kind of environment using robust control methods.

[66] It should also be clear from equation (5.7) that with parameter uncertainty the interest rate no longer adjusts to perfectly offset demand shocks. Suppose, for example, that there is a positive demand shock. The interest rate goes up, but the parameter uncertainty moderates the extent of the rise, relative to the certainty equivalence case.

uncertainty about the impact of an increase in the interest rate on the output gap moderates the extent of adjustment in i_t. The second term on the right side of equation (5.7) captures this latter dampening effect.

This simple form of model uncertainty thus may help explain the relatively low variability of interest rates in the data. One feature of interest rate smoothing it does not appear to capture, however, is the strong lagged dependence in the interest rate. Put differently, the kind of parameter uncertainty we have discussed may explain why the slope coefficients on inflation and the output gap, α and β, are small relative to the case of certainty equivalence. But it does not explain the partial adjustment, given by the dependence of i_t on i_{t-1}.[67]

Rotemberg and Woodford (1997) offer a novel explanation for the lagged dependence that is based on the leverage that this kind of adjustment rule may provide the central bank over the long term interest rate. The idea is that lagged dependence in i_t permits the central bank to manipulate long term rates, and hence aggregate demand, with more modest movements in the short term rate than would be otherwise be required. This kind of rule is thus desirable to the extent the central bank may care about avoiding excessive volatility in the short term interest rate in pursuing its stabilization goals.

To illustrate, consider the special case of equation (5.6) with $\rho = 1$. In this instance, the difference in the interest rate $(i_t - i_{t-1})$, as opposed to the level, is a linear function of π_t and x_t. Under the difference rule, the expected future short rate at $t+i$, $E_t\{i_{t+i}\}$, is given by

$$E_t\{i_{t+k}\} = E_t\left\{\sum_{j=1}^{k}(i_{t+j} - i_{t+j-1})\right\} + i_t \qquad (5.8)$$
$$= E_t\left\{\sum_{j=1}^{k}[\alpha + \beta\pi_{t+j} + \gamma x_{t+j}]\right\} + i_t$$

Assume that the long-term rate depends on the sum of expected short rates over the same horizon, in keeping with the expectations hypothesis of the term structure. Then, in comparison with the level rule, the difference rule increases the responsiveness of the long term rate under the feedback policy. Suppose for example that, in reaction to a rise in inflation above target at time t, the central bank raises i_t above its steady state value. Under the difference rule the increase in the interest rate has a persistent effect on the path of the expected short rate, since $E_t\{i_{t+i}\}$ depends additively on i_t. Further, if changes in inflation and output are persistent, then the path of expected short rates will actually be rising, as equation (5.8) makes clear.[68] The difference rule thus enhances the countercyclical movement of the long rate relative to the movement of short rate. Given that aggregate demand depends on the long rate, this kind of rule thus enables the central bank to

[67] Sack (1997a,b) argues, nonetheless, that parameter uncertainty can explain this phenomenon if the uncertainty of the impact of the interest rate on the economy is based on the change in the interest rate $(i_t - i_{t-1})$ as opposed to the deviation from trend i_t. In the former instance, changes in i_t raise the conditional variability of output, which induces the central bank to keep i_t close to i_{t-1}. On the other hand, it is not well understood how the link between model uncertainty and policy conservatism is affected when there is active learning about the economy. Some results suggest that learning should induce active adjustments of the policy instrument to facilitate estimating the true model. See the discussion in Wieland (1997), for example. Also, it is possible to construct examples where parameter uncertainty leads to increased activism. See, for example, Thomas Sargent (1998).

[68] On the surface it appears that the interest rate might explode under the difference rule, since it will continue to increase so long as inflation is above target. However, the rise in the interest rate will dampen demand and inflation. In the context of our model, it does so sufficiently to preclude explosive behavior.

stabilize the economy with relatively modest movements in the short rate.[69]

Overall, Rotemberg and Woodford provide a plausible explanation for why central banks may want to introduce lagged dependence in the interest rate. Whether this story can also account for the empirically observed modest response of the short rate to inflation and the output gap (i.e., the low values of β and γ, the slope coefficients on π_t and x_t) remains to be seen.

Another explanation for policy conservatism and the associated interest rate smoothing includes fear of disrupting financial markets (see, e.g., Goodfriend 1991). Sharp unanticipated increases in interest rates can generate capital losses, particularly for commercial banks and other financial institutions that may be exposed to interest rate risk. This consideration might explain why the Federal Reserve chose to raise rates only very gradually during 1994, the tail end of a period of considerable financial distress (see, e.g., the discussion in John Campbell 1995). Disagreement among policy-makers is another explanation for slow adjustment of rates. Neither of these alternative stories have been well developed, however. In general, understanding why interest rate smoothing occurs in practice is an important unresolved issue.

5.3 *Non-Smooth Preferences and Opportunism*

Another aspect of policy that has received considerable attention involves the process of disinflation. In the baseline model, if inflation is above target, it is always optimal to tighten monetary policy to gradually bring inflation back to the optimum (see Result 2 in Section 3). During his tenure at the Federal Reserve Board, however, Blinder proposed the following alternative: If inflation is above but near the optimum, policy should not contract demand. Rather, it should take an "opportunistic" approach. Roughly speaking, being opportunistic boils down to waiting until achieving the inflation target could be done at the least cost in terms of incremental output reduction. Blinder's original concept was vague as to the details. Recent work by researchers at the Federal Reserve Board has filled in a number of the missing pieces.

Athanasios Orphanides and David Wilcox (1996) show that it is possible to rationalize something like opportunistic policy by making a small adjustment of the policy objective function. In particular, suppose that policy-makers care quite a lot about small departures of output from target, at least relative to small departures of inflation. An example of an objective function that capture this phenomenon is given by

$$\max -\frac{1}{2}E_t\left\{\sum_{i=0}^{\infty}\beta^i(\alpha\,|\,x_{t+i}\,|+\pi_{t+i}^2)\right\} \qquad (5.9)$$

With this objective function, the optimality condition for policy becomes:

$$\begin{aligned} x_t &= 0, \text{ if } |\,\pi_t\,| < \frac{\alpha}{\lambda} \\ |\,\pi_t\,| &= \frac{\alpha}{\lambda}, \text{ otherwise} \end{aligned} \qquad (5.10)$$

Thus, if inflation is within $\frac{\alpha}{\lambda}$ units of the target, the optimal policy is to simply stabilize output. Otherwise, policy should keep inflation at most $\frac{\alpha}{\lambda}$ units from target and then wait for favorable supply shocks that move it closer to target (e.g., favorable movements in the cost push shock u_t). In this respect the

[69] The idea that the central bank should pursue a partial adjustment rule to exploit the dependence of demand on future policy is reminiscent of the globally optimal policy under commitment (see section 4.2.2). Indeed, Woodford (1998) makes this connection formally.

policy is opportunistic. A better term for it, however, might be "inflation zone targeting" (Bernanke and Mishkin 1997). What the policy really amounts to is keeping inflation with a certain range, as opposed to trying to move it to an exact target.

Variations on this theme allow for preferences that generate an inflation zone target, but then has policy trade off between inflation and output goals when inflation is outside the target zone. Orphanides, David Small, Volcker Wieland and Wilcox (1997) (OSWW) provide an example of this more general setup.

It is important to emphasize, though, that opportunistic policy behavior that is distinct from the gradualism of the baseline model only arises if cost push factors are present in inflation. This is true because only with cost push inflation present does a trade-off between output and inflation emerge (see Result 1). Indeed, OSWW show that opportunistic policy rules are equivalent to conventional gradualist rules in the presence of demand shocks, but differ when there are supply shocks.[70]

In summary, we have

Result 12: *If there is more cost associated with small departures of output from target than with small departures of inflation, then an opportunistic approach to disinflation may be optimal. This policy, further, is equivalent to targeting inflation around a zone as opposed to a particular value.*

6. *Implications of Endogenous Inflation and Output Persistence*

Within our baseline model, the dynamics of output and inflation are due entirely to exogenous force processes. We now consider an alternative framework that allows for endogenous persistence in output and inflation. Our purpose is to show that the results derived in the baseline framework extend to this more general setting. In this regard, we show that our results are not specific to the particular benchmark model we employed, but instead hold across a reasonably broad class of models that are used for applied macroeconomic analysis. The major difference is that with endogenous persistence in inflation, the equilibrium feedback monetary policy now influences the speed of convergence of inflation to its target.

Consider the following generalizations of the IS and aggregate supply curves:

$$x_t = -\varphi[i_t - E_t\pi_{t+1}] + \theta\, x_{t-1} + (1-\theta)E_t x_{t+1} + g_t \quad (6.1)$$

$$\pi_t = \lambda\, x_t + \phi\, \pi_{t-1} + (1-\phi)\beta E_t\pi_{t+1} + u_t \quad (6.2)$$

Equation (6.1) incorporates the lagged output gap in the IS curve. Equation (6.2) adds lagged inflation to the aggregate supply curve. The parameters θ and ϕ index the influence of lagged versus expected future variables. As a result the model nests some important special cases. With $\theta = 0$ and $\phi = 0$, we recover the baseline model. Conversely, with $\theta = 1$ and $\phi = 1$, the model becomes (approximately) the backward-looking framework that Svensson (1997a,b) and Ball (1997) have used to analyze monetary policy. For simplicity we assume that the disturbances g_t and u_t are serially uncorrelated (i.e., we set μ and ρ in equation (2.3) and (2.4) equal to zero). This simple formulation does not allow for delays in the effect of policy, but we show later that it is easy to amend the analysis to incorporate delayed policy effects.

As we noted earlier, virtually all the

[70] For an alternative description of the opportunistic approach, see Bomfin and Rudebusch (1997). These authors emphasize the ratcheting down of inflation and, in particular, explore the role of imperfect credibility.

major applied macroeconomic models allow for some form of lagged dependence in output and inflation. The primary justification is empirical.[71] By appealing to some form of adjustment costs, it may be feasible to explicitly motivate the appearance of x_{t-1} within the IS curve. Motivating the appearance of lagged inflation in the aggregate supply curve, however, is a more formidable challenge.[72] Some frameworks do so by effectively appealing to costs of changing the rate of inflation.[73] This assumption, though, is clearly unattractive. In the spirit of robustness, however, it is important to understand the implications of lagged dependence. This is particularly true given the empirical appeal of this formulation.

We begin with the case of discretion, and then later describe briefly how the results are affected when the central bank can make credible promises.[74] An analytical solution is not available, except in the polar cases of $\phi = 0$ and $\phi = 1$. It is, however, possible to provide an intuitive description of the optimum. Let a_π be a parameter that measures the serial dependence of inflation in the reduced form. Then the optimality condition that governs policy is given by:

$$x_t = -\frac{\lambda}{\alpha}\left[\pi_t + \sum_{k=1}^{\infty} \beta^k E_t \pi_{t+k}\right] \quad (6.3)$$

$$= -\frac{\lambda}{\alpha(1-\beta a_\pi)}\pi_t \quad (6.4)$$

with

$$\pi_t = a_\pi \pi_{t-1} + a_u u_t \quad (6.5)$$

and

$$0 \leq a_\pi < 1$$

With inertia present, adjustments in current monetary policy affect future time path of inflation. As consequence, policy now responds not only to current inflation but also to forecasts of inflation into the indefinite future. How much depends positively on a_π, which measures the degree of inflationary persistence.

The coefficients a_π and a_u are functions of the underlying parameters $(\alpha,\lambda,\beta,\phi)$.[75] The former, a_π, is key, since it measures the speed of convergence to inflation under the optimal policy. It is possible to show that this parameter lies between zero and unity, implying convergence. The magnitude of a_π depends positively on the degree of inflation inertia ϕ. In the baseline case of no inflation inertia, $\phi = 0$, implying $a_\pi = 0$. a_π

[71] For an empirical justification for including lagged dependent variables, see Fuhrer (1996).

[72] It is possible to motivate a dependency of current inflation on lagged inflation by appealing to adaptive expectations (e.g., suppose $E_{t-1}\pi_t = \kappa\pi_{t-1}$). Indeed, this is the traditional approach (see the discussion in Blanchard 1997). The issue then becomes motivating the assumption of adaptive expectations.

[73] See, for example, Fuhrer and Moore (1995a,b) and Brayton, Levin, Tyron, and Williams (1997). Galí and Gertler (forthcoming) criticize the existing empirical literature on inflation dynamics, and provide new evidence which suggests that (2.2) is a good first approximation to the data.

[74] As in section 3, we restrict attention to Markov perfect equilibria. In this case, however, we must take into account that inflation is an endogenous state variable. In any stationary equilibrium, therefore, expected inflation will depend on lagged inflation. What the policy maker takes as given, accordingly, is not the level of expected inflation, but rather how private sector expectations of inflation tomorrow respond to movements in inflation today. Simply put, to solve for the equilibrium under discretion, we assume that private sector forecast of π_{t+1} takes the form $v_\pi\pi_t + v_u u_t$, where v_π and v_u are arbitrary constants that the policy-maker takes as given. In the rational expectations equilibrium v_π and v_u equal the true fundamental parameters in the reduced form inflation, a_π and a_u.

[75] To obtain solutions for a_π and a_u, substitute the optimality condition $x_t = -\frac{\lambda}{\alpha(1-\beta a_\pi)}\pi_t$ and the conjectured solution for π_t, (6.5), into the aggregate supply curve. Then use the methods of undetermined coefficients to solve for a_π and a_u. The equation for a_π is a cubic. The solution is the unique value between zero and unity, which corresponds to the unique stable root.

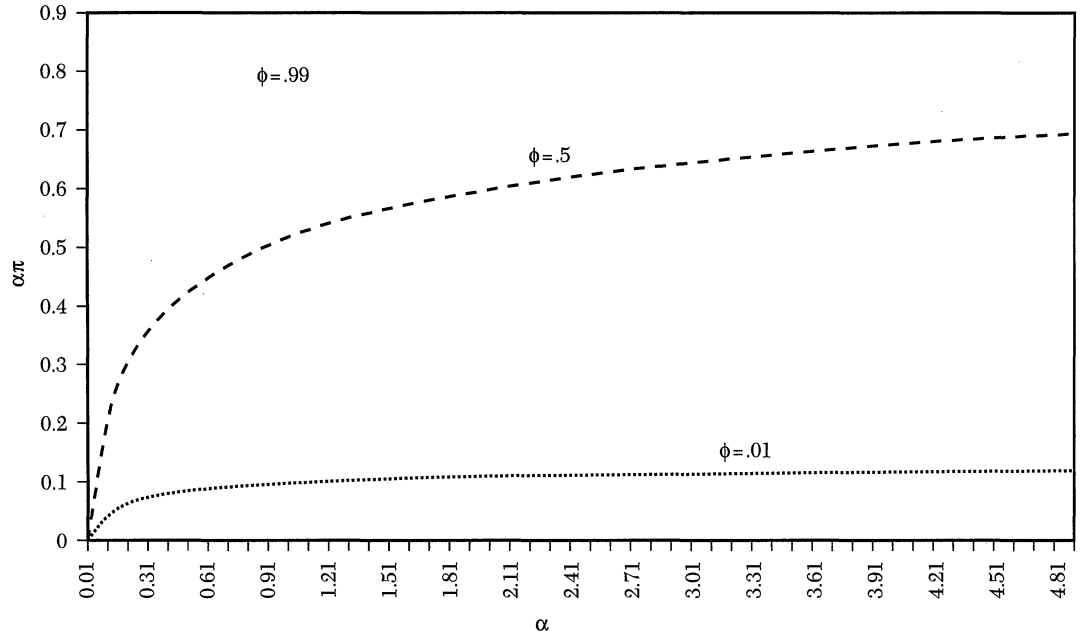


Figure 2. α_π for different values of ϕ

also depends negatively on the relative cost of inflation, measured by $1/\alpha$. As in the baseline case, if the distaste for inflation is high (α is low), the optimal policy aggressively contracts demand whenever inflation is above target: With endogenous persistence, this contraction not only reduces inflation but also increases the speed of convergence to target. Figure 2 illustrates the relation between a_π and α for three different values of ϕ: $\phi = 0.01$ (low inertia), $\phi = 0.5$ (medium) and $\phi = 0.99$ (high).

Combining (6.3) with (6.1) yields the implied optimal interest rate rule:

$$i_t = \gamma_\pi E_t\pi_{t+1} + \gamma_x x_{t-1} + \frac{1}{\varphi} g_t \qquad (6.6)$$

with

$$\gamma_\pi = 1 + \frac{\lambda(1 - a_\pi)}{\varphi\alpha a_\pi(1 - \beta a_\pi)}$$

$$\gamma_x = \frac{\theta}{\varphi}$$

$$E_t\pi_{t+1} = a_\pi\pi_t$$

Most of the qualitative results obtained in the baseline case extend to this more general setting. As in the baseline case, the policy-maker faces a short-run trade-off between output and inflation (Result 1). The effect of inflation inertia is to make this trade-off less favorable. Equation (6.3) shows that relative to the baseline case of $\phi = 0$, the optimal policy requires a more aggressive response to any burst of inflation. The problem is that any inflation not eliminated today persists into the future, potentially requiring more output contraction. Figure 3 illustrates how the trade-off becomes less favorable in this case by plotting the efficient policy frontier for the three benchmark values of ϕ. In addition, since $0 \leq a_\pi < 1$, the optimal policy calls for gradual adjustment of inflation to target (Result 2). With $\phi > 0$, further, extreme inflation targeting is only optimal if $\alpha = 0$, as equation (6.3) and Figure 2 suggest.

From the interest rate rule given by

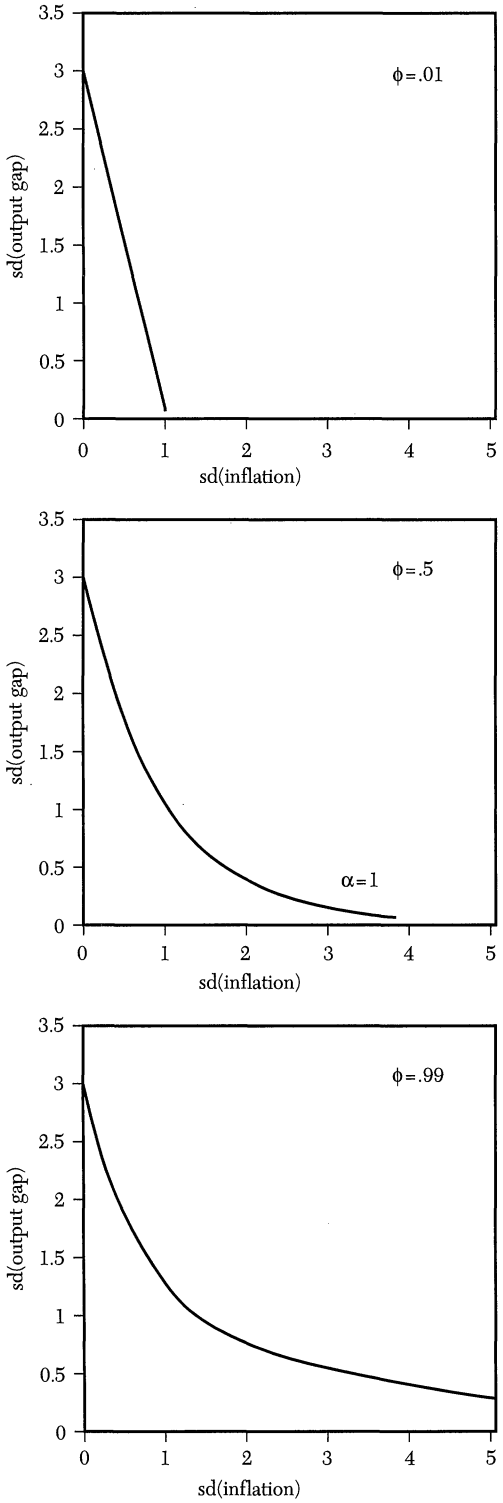


Figure 3. Efficient Policy Frontier for Different Values of ϕ

equation (6.6) it is apparent that the coefficient on expected inflation exceeds unity, implying that the ex ante real rate must rise in response to higher expected inflation (Result 3). Finally, the interest rate should also adjust to perfectly offset demand shocks, but should not respond to movements in potential output (Result 4.) One interesting difference in this case is that the interest rate responds to the lagged output gap, since this variable now enters the IS curve. Thus, the optimal interest rate rule now resembles the simple gap rules that have been discussed in the literature. We return to this point later. In summary, we have

Result 13: *Results 1 through 4 that describe optimal monetary policy under discretion within the baseline model also apply in the case with endogenous output and inflation persistence.*

In addition to allowing for lagged dependence in output and inflation, there is also strong empirical justification for incorporating delays in the effect of policy. It is straightforward to extend the analysis to include this real world feature. Suppose, following Svensson (1997a,b) and Ball (1997), that there is a one-period delay in the effect of the real interest rate on the output gap and, in turn, a one-period delay in the effect of the output gap on inflation. Then the optimality condition becomes[76]

$$E_t\{x_{t+1}\} = -\frac{\lambda}{\alpha(1-\beta a_\pi^l)} E_t\{\pi_{t+2}\} \quad (6.7)$$

where the parameter a_π^l measures the serial dependence in inflation for this case. It has qualitatively similar properties to a_π in equation (6.5), with $0 \le a_\pi^l < 1$. The left side of (6.7) reflects the one-period delay in the impact of policy on output,

[76] In this case, the IS curve is given by $x_t = -\varphi[i_{t-1} - E_{t-1}\pi_t] + \theta x_{t-1} + (1-\theta)E_{t-1}x_{t+1} + g_t$ and the aggregate supply curve is given by $\pi_t = \lambda x_{t-1} + \phi\pi_{t-1} + (1-\phi)\beta E_t\pi_{t+1} + u_t$.

and the right side reflects the two-period delay on inflation.

Due to the delayed impact of policy, the central bank takes both the output gap at t, x_t, and the forecast of inflation at $t+1$, $E_t\{\pi_{t+1}\}$, as predetermined from the vantage of time t. The rest of the solution may thus be expressed in terms of these predetermined variables:

$$E_t\{\pi_{t+2}\} = a_\pi^l E_t\{\pi_{t+1}\} \quad (6.8)$$

$$i_t = \gamma_\pi^l E_t\pi_{t+1} + \gamma_x x_t \quad (6.9)$$

with

$$\gamma_\pi^l = 1 + \frac{(\gamma_\pi - 1)\beta}{a_\pi^l} > 1$$

The solution closely resembles the case without delay. Any differences just reflect the lagged influence of policy in this environment. The nominal rate still adjusts more than one-for-one with expected inflation. Due to the lag structure, though, it adjusts to the current output gap, as opposed to one from the previous period.

We conclude this section with brief discussion of the gains from commitment. It is possible to show that, as in the baseline model, the policy rule under commitment resembles the rule under discretion that would obtain if the policy-maker assigned a higher relative cost to inflation (lower value of α) than the true social cost. Because inflation inertia is endogenous in this case, the optimal policy with commitment implies a faster transition of inflation to the optimum relative to what occurs under discretion. This can be seen by noting that the parameter which governs the speed of convergence of inflation, a_π, is decreasing in the relative cost of inflation $1/\alpha$ (see Figure 4).[77] Simply put, disinflations will be swifter than otherwise if credible commitment is possible either directly or indirectly by installing a conservative central bank chair.

[77] Note that the speed of convergence of inflation is decreasing in a_π.

7. *Simple Rules for Monetary Policy*

We next discuss some normative and positive aspects of simple feedback rules for the interest rate that have been discussed in the literature. We then discuss how these instrument-based rules are related to simple rules for targets that have been recently proposed, including inflation targeting and nominal GDP targeting. Finally, we conclude with a brief discussion of the issue of possible indeterminacy of interest rate rules.

7.1 *Simple Interest Rate Rules*

Taylor (1993a) ignited the discussion of simple interest rate rules.[78] He proposed a feedback policy of the following form:

$$i_t^* = \alpha + \gamma_\pi (\pi_t - \bar{\pi}) + \gamma_x x_t \quad (7.1)$$

with

$$\alpha = \bar{r} + \bar{\pi}$$

$$\gamma_\pi > 1, \gamma_x > 0$$

where i_t^* is the target interest rate the feedback rule defines, $\bar{\pi}$ is the target inflation rate, and $\bar{r}$ is the long-run equilibrium real interest rate.[79] Also, we now express all variables in levels, as opposed to deviations from trend.

A number of other researchers have considered rules like (7.1) (see, e.g., Henderson and Mckibbon 1993). Taylor's contribution is to spell out the normative and positive implications. On

[78] McCallum (1988) proposed a simple rule for the monetary base. The rule is less popular in policy circles due to the implied interest rate volatility (see Result 9). McCallum (1997) argues, however, that the concern about interest rate volatility is not well understood, a point with which we agree.

[79] The inflation rate Taylor uses is actually the rate over the previous year (as opposed to the previous quarter).

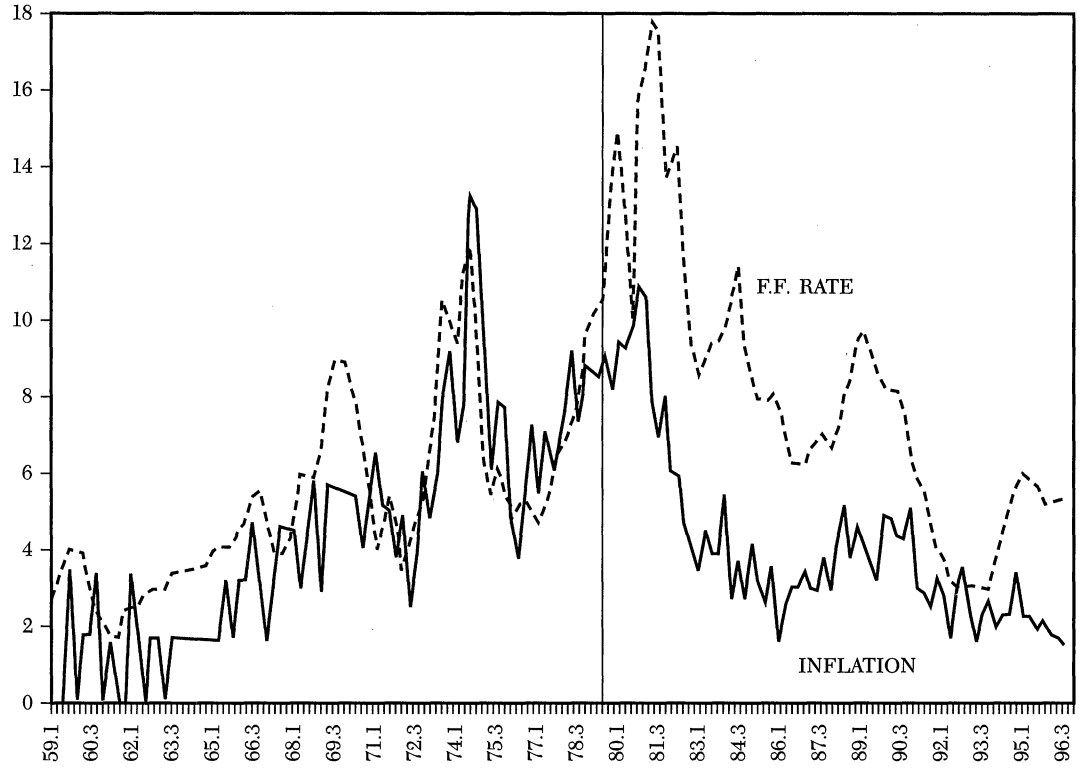


Figure 4. The Federal Funds Rate and the Inflation Rate

the normative side, the rule is consistent with the main principles for optimal policy that we have described. It calls for gradual adjustment of inflation to its target (see Result 2). Specifically, it has the nominal rate adjust more than one-for-one with the inflation rate. To the extent lagged inflation is a good predictor of future inflation, the rule thus has real rates adjusting to engineer inflation back to target (see Result 3). Finally, note that the interest rate responds to the output gap as opposed to the level of output. Thus, in at least an approximate sense, the rule calls for a countercyclical response to demand shocks and accommodation of shocks to potential GDP that do not affect the output gap (see Result 4).

On the positive side, Taylor showed that with certain parameter values, the rule provides a reasonably good description of policy over the period 1987–92. These are: $\gamma_\pi = 1.5$, $\gamma_x = 0.5$, $\bar{\pi} = 2$, and $\bar{r} = 2$. Taylor used informal judgement to pick them. An interesting question is whether a formal methodology would yield something different.

In this spirit, Clarida, Galí, and Gertler (forthcoming) estimate a simple rule for U.S. monetary policy, and consider how this rule has evolved over time. The specific formulation is a "forward looking" version of the simple Taylor rule:

$$i_t^* = \alpha + \gamma_\pi (E_t \pi_{t+1} - \bar{\pi}) + \gamma_x x_t \quad (7.2)$$

Under this rule, policy responds to expected inflation as opposed to lagged inflation. In this respect, the formulation is consistent with the optimal rules derived for both the baseline and hybrid models (see equations 3.6 and 6.6).

TABLE 1
ESTIMATES OF POLICY REACTION FUNCTION

	γ_π	γ_x	ρ
Pre-Volcker	0.83	0.27	0.68
	(0.07)	(0.08)	(0.05)
Volcker–Greenspan	2.15	0.93	0.79
	(0.40)	(0.42)	(0.04)

Another virtue is that this formulation nests the simple Taylor rule as a special case. If either inflation or a linear combination of lagged inflation and the output gap is a sufficient statistic for future inflation, then the specification collapses to the Taylor rule.

Because of the Federal Reserve's tendency to smooth interest rate adjustments (see the discussion in section 5), a static relation like equation (7.2) cannot capture the serial correlation present in the data. We thus allow for the possibility of partial adjustment to the target rate, according to:

$$i_t = \rho\, i_{t-1} + (1-\rho) i_t^* \qquad (7.3)$$

where ρ is a parameter that measures the degree of interest rate smoothing.

We estimate different rules for the pre-Volcker (1960:1–79:2) and Volcker–Greenspan (1979:3–96:4). We do so because it is widely believed that U.S. monetary policy took an important turn for the better with the appointment of Paul Volcker as Fed Chairman (see Friedman and Kuttner 1996 and Gertler 1996). Among other things, this period marks the beginning of an apparently successful and long-lasting disinflation.

We find that the simple rule given by equation (7.2) does a good job of characterizing policy in the Volcker–Greenspan era. Further, it adheres to the guidelines for good policy that we have established. The estimated pre-Volcker rule violates these guidelines. Specifically, the parameter estimates along with standard errors are given by Table 1.[80]

The key lesson involves the parameter γ_π, the coefficient on the inflation gap. The estimate for the pre-Volcker rule is significantly less than unity. This suggests that monetary policy over this period was accommodating increases in expected inflation, in clear violation of the guidelines suggested by Results 2 and 3. For the post-1979 rule the estimate is significantly above unity. It thus incorporates the implicit inflation targeting feature that we have argued is a critical feature of good monetary policy management. It is also true that in the Volcker–Greenspan era the Federal Reserve was only responding to the output gap to the extent it had predictive power for inflation:[81] The estimated coefficient on the output gap, γ_x, is not significantly different from zero. Pre 1979:4 it is positive and significant. This outcome is consistent with the conventional view that pre-1979, the Federal

[80] The estimates of the parameters in equation (7.2) are obtained by using an instrumental variables procedure based on Generalized Methods of Moments (GMM). See Clarida, Galí, and Gertler (forthcoming) for details. The specific numbers reported here are based on a version of this policy reaction function that has the Funds rate respond to expected inflation a year ahead and the current output gap (reported in Table 2 of that paper). The results, however, are robust to reasonable variations in the horizons for the gap variables.

[81] In particular, the output gap enters the instrument set for expected inflation. Thus, the coefficent γ_x reflects the influence of the output gap on the interest rate that is independent of its predictive power for inflation.



Figure 5. Target Based On Estimated Post-October '79 Rule vs. Actual Funds Rate

Reserve was relatively more focused on output stabilization and less focused on inflation.

The finding that the Fed responded differently to inflation in the two eras is apparent from inspection of the data. Figure 4 plots the Federal Funds rate and the rate of CPI inflation from 1965 to the present. The graph shows a clear break in the Funds rate process around 1979.[82] During most of the 1970s, the ex post real rate was zero or negative. After 1979 it becomes positive. While many factors influence the real rate, the tight monetary policy engineered by Paul Volcker surely provides the most logical explanation for this initial run-up.

Figure 5 illustrates the policy change by plotting the estimated target value of the interest rate under the Volcker–Greenspan rule over the entire sample period. The target rule does a good job of capturing the broad movements in the Funds rate for the second half of the sample, for which it was estimated. For the pre-Volcker period, matters are different. The target (generated by the estimated Volcker–Greenspan rule) is systematically well above the historical series. In this concrete respect, policy was far less aggressive in fighting inflation in the earlier period.[83]

Figure 6 compares the ability of the forward and backward looking (Taylor) target rules to explain the post 1979 data. Though we find that the data rejects the backward looking rule in favor of the forward looking one,[84] the two do a roughly similar job of accounting for the behavior of the Funds rate. This occurs probably because, with U.S. data,

[82] Huizinga and Mishkin (1986) present formal evidence of a structural break at this time.

[83] Some but not nearly all the difference between rates pre-1979 and the target values under a post-1979 rule could be accounted for by a secular change in the real rate.

[84] See Clarida, Galí, and Gertler (1998).

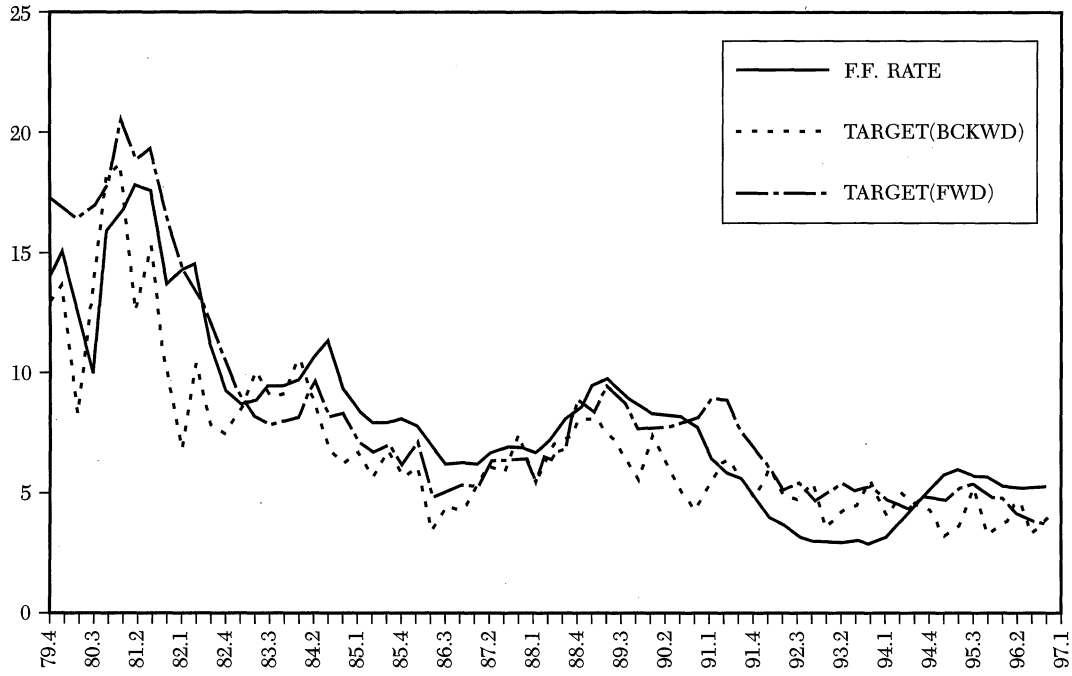


Figure 6. Targets from Forward vs. Backward Looking Rules

not much besides lagged inflation is useful for predicting future inflation.

Finally, it is interesting to observe that the other major central banks, the Bundesbank and the Bank of Japan, have behaved very similarly in the post-1979 era. In Clarida, Galí, and Gertler (1998), we estimate our specification for these central banks. The estimated parameters in each case are quite close to those obtained for the Federal Reserve during the Volcker–Greenspan period. Thus, good policy management appears to have been a global phenomenon. Perhaps this is not surprising since the successful disinflation has also been a world-wide event.

7.2 *Simple Target Rules*

There have also been proposed simple rules for targets, as opposed to instruments. Of these proposed policies, inflation targeting has received by far the most attention (see Bernanke and Mishkin 1997 for a recent survey). Indeed a number of central banks, most notably the Bank of England, have recently adopted formal inflation targets (see, e.g., Andrew Haldane 1996).

In one sense, inflation targeting involves nothing more than pursuing the kind of gradualist policy that our optimal policy calculation implies (see Result 2). Indeed, all the leading real-world proposals call for gradual convergence of inflation to target. None recommend trying to hit the inflation target continuously, which is consistent with our analysis. In this respect, the rule we estimate for the period is perfectly consistent with inflation targeting.

The rationale for inflation targeting, we think, is twofold. The first is simply to guarantee that monetary policy avoids the mistakes of the pre-Volcker era by identifying a clear nominal anchor for policy. (After all, Alan Greenspan will not be around forever). The

inflation target is in effect the nominal anchor. Since the anchor is directly in terms of inflation, it avoids the potentially instability problems associated with alternatives such as money growth that are only indirectly linked to inflation. For example, if there are large shocks to money demand, then a money growth target may fail precisely to pin down the equilibrium inflation rate.

The second rationale has to do with credibility and commitment. We have seen that it is in general optimal for policy-makers to place a higher weight on the costs of inflation than the true social loss function suggests (see Results 6 and 7). The focus on inflation targets may be viewed as a way to instill a higher effective weight on inflation in the policy choice.

Price level targeting is another type of simple rule that has been discussed in the literature. This policy, which may be thought of as a more extreme version of inflation targeting, has not received much support among policy-makers and applied economists. There are several problems: First, if the price level overshoots its target, the central bank may have to contract economic activity in order to return the price level to its goal. That is, inflation above the amount implied by the price level target must be followed by inflation below this desired amount in order to return the target. Under inflation targeting, bygones are bygones: overshooting of inflation in one year does not require forcing inflation below target in the following year. Second, the source of positive drift in the price level maybe measurement error (see the discussion in section 2.) It would be unfortunate to have measurement error induce tightening of monetary policy. Third, as McCallum (1997b) shows, the net reduction in price uncertainty under a price level target rule, may be small relative that obtained under an inflation targeting policy. For all these reasons, it is perhaps not surprising that no major central bank has adopted a price level target.

Another candidate variable for targeting is nominal GDP. This approach has also received less attention in the recent literature, however. One problem is that if there are shifts in the trend growth of real GDP, the rule does not provide a precise nominal anchor. Another problem, emphasized by Ball (1997), is that the policy may be overly restrictive. In the hybrid model of section 5, for example, the optimal policy in general has the interest rate adjust to some linear combination of expected inflation, the output gap and demand disturbances. The weights depend upon the underlying structural parameters of the model. Under nominal GDP targeting, the central bank adjusts the interest rate to the sum of inflation and real GDP growth. It thus arbitrarily applies an equal weight to each component of nominal GDP. High nominal GDP growth, further, could occur when the economy is recovering from a recession and is still well below full capacity. A rule that calls for raising interest rates in response to above-target nominal GDP growth in these circumstances could stifle the recovery.[85]

7.3 *Indeterminacy under Interest Rate Rules*

One criticism of simple interest rate rules is that, under certain circumstances, they may induce instability. That is, in many models there may not be a determinate equilibrium under particular parametrizations of the policy

[85] See Ball (1997) and Svensson (1997b) for explicit examples of how nominal GDP targeting could produce adverse outcomes. McCallum (1997c), however, argues that these results are sensitive to the use of a backward-looking Phillips curve. For the case in favor of nominal GDP targeting, see Hall and Mankiw (1994).

rule. In a classic paper, Thomas Sargent and Neil Wallace (1975), illustrated how nominal indeterminacy may arise if prices are perfectly flexible. Under an interest rate rule the equilibrium pins down the level of real money balances. However, there are an infinite number of combinations of the nominal money stock and the price level that satisfy this equilibrium condition.[86] In this respect, the interest rate rule produces nominal indeterminacy.[87]

When there is sluggish price adjustment, the problem of nominal indeterminacy vanishes. Last period's price level effectively serves a nominal anchor. Simple interest rate rules thus do not produce price level indeterminacy in the frameworks we have analyzed. More generally, since there is little reason to believe that prices are perfectly flexible, the issue of nominal indeterminacy does not seem important in practice. On the other hand, there is potentially a problem of real indeterminacy in the case of price stickiness, as William Kerr and Robert King (1996), Bernanke and Woodford (1997) and Clarida, Galí and Gertler (forthcoming) have recently emphasized.[88] Two types of indeterminacy are possible. First, if in response to a rise in expected inflation, the nominal rate does not increase sufficiently to raise the real rate, then self-fulfilling bursts of inflation and output are possible. A rise in expected inflation, leads to a fall in real rates that, in turn, fuels the boom. Indeed, the monetary policy rule that Clarida, Galí, and Gertler (forthcoming) estimate for the pre-Volcker period permits exactly this kind of sunspot behavior. The lesson here is simply that a good monetary policy rule should not accommodate rise in expected inflation. It should instead pursue the implicit kind of inflation targeting that we have been emphasizing. This boils down to raising nominal rates sufficiently to increase real rates whenever expected inflation goes up.

As Bernanke and Woodford (1998) emphasize, indeterminacy is also possible if the rule calls for an overly aggressive response of interest rates to movements in expected inflation. In this instance, there is a "policy overkill" effect that emerges that may result in an oscillating equilibrium. Clarida, Galí and Gertler (forthcoming) show, however, the magnitude of the policy response required to generate indeterminacy of this type greatly exceeds the estimates obtained in practice. This potential indeterminacy however does suggest another reason why a gradual approach to meeting an inflation target may be desirable.

8. *Concluding Remarks*

We conclude by describing several areas where future research would be quite useful:

(1) It is always the case that more knowledge of the way the macroeconomy works can improve the performance of monetary policy. Particularly critical, however, is a better understanding

[86] McCallum (1997), however, argues that the price level is in fact determined in this kind of environment.

[87] A recent literature shows that the government's intertemporal budget constraint may restore uniqueness under an interest rule, even in an environment with flexible prices. What is critical is whether the interest on the debt is financed by taxes or money creation. See, for example, Woodford (1994), Sims (1994), and Leeper (1991).

[88] These papers focus on local indeterminacy. See Jess Benhabib, Stephanie Schmidt-Grohe, and Martin Uribe (1998) for a discussion of global indeterminacy. To avoid global indeterminacy, the central bank may have to commit to deviate from a simple interest rate rule if the economy were to get sufficiently off track. This threat to deviate can be stabilizing, much the way off the equilibrium path threats induce uniqueness in game theory. Because the threat is sufficient to preclude indeterminate behavior, further, it may never have to be implemented in practice.

of the determinants of inflation. As we have emphasized, the output/inflation trade-off is highly sensitive to both the degree and nature of the persistence in inflation. As a consequence, so too is the speed at which monetary policy should try to reach the optimal inflation rate. Rationalizing the observed persistence in inflation is thus a high priority. Work by Galí and Gertler (forthcoming) and Argia Sbordone (1998) suggests that the short-run aggregate supply curve employed in our baseline model may provide a reasonable approximation of reality, so long as real marginal cost (specifically real unit labor costs) is used as the relevant real sector forcing variable instead of the output gap, as the theory suggests. Galí and Gertler (forthcoming) argue further that persistence in inflation may be related to sluggish adjustment of unit labor costs vis-a-vis movements in output. Sorting out this issue will have important repercussions for monetary policy.

(2) Our analysis of monetary policy, as in much of the literature, was restricted to closed economy models. Extensions to open economy frameworks are likely to provide new insights on the desirability of alternative monetary policy rules, and raise a number of issues of great interest, including: the choice of exchange rate regime, the potential benefits from monetary policy coordination, the optimal response to shocks originating abroad, and consumer price index versus domestic inflation targeting. Recent work by Ball (1998), Svensson (1998), and Monacelli (1999) along these lines will undoubtedly lay the ground for further research on this front.

(3) Throughout the analysis, we assumed that the lower bound of zero on the nominal interest rate was not a constraint on the performance of monetary policy. In Japan, for example, the short-term nominal rate has fallen to the point where this constraint clearly is a consideration for policy management. Similarly, in the U.S. and Europe, the inflation rates have fallen to the point where the zero bound limit could conceivably affect the ability to ease rates in the event of a downturn. Understanding how monetary policy should proceed in this kind of environment is an important task. When the nominal rate is at zero, the only way a central bank can reduce the real interest rate is to generate a rise in expected inflation (see the discussion in Alexander Wolman 1998, and the references therein). How the central bank should go about this and whether cooperation from fiscal policy is necessary are important open questions. As Wolman (1998) suggests, the conclusions are quite sensitive to the nature of the inflationary process.

(4) A more specific issue, but nonetheless an important one, is to understand why central banks smooth interest rate adjustments. As we discussed in section 5, optimal policies implied by most existing macroeconomic frameworks generate paths for the interest rate that are much more volatile than what is observed in reality. The possibility thus arises that existing models may fail to adequately characterize the constraints that policy-makers face in practice. We suggested in section 5 that some form of model uncertainty might be able to account for this phenomenon. Another alternative is that central banks may be exploiting the dependency of demand on expected future interest rates, as argued by Rotemberg and Woodford (1999). Whether these explanations or any others, such as fear of disruption of financial markets, can account for interest rate smoothing needs to be determined.

(5) A somewhat related issue involves how a central bank should deal with

financial stability. The policy rules discussed in the literature do include contingencies for financial crises. A frequently cited reason for why monetary policy should not adhere tightly to a simple rule is the need for flexibility in the event of a financial collapse. In the wake of the October 1987 stock market crash, for example, most economists supported the decision of the Federal Reserve Board to reduce interest rates. This support was based largely on instinct, however, since there is virtually no formal theoretical work that rationalizes this kind of intervention. More generally, concern about financial stability appears to be an important constraint on policy-making. As we suggested in section 5, it is one possible reason why central banks smooth interest rate changes. Understanding the nature of this concern is clearly a fertile area for research.

(6) Finally, with few exceptions, virtually all the literature ignores the issue of transition to a new policy regime.[89] In particular, the rational expectations assumption is typically employed. Policy simulations thus implicitly presume that the private sector catches on immediately to any regime change. In reality, however, there may be a period of transition where the private sector learns about the regime change. This kind of scenario may be highly relevant to a central bank that has accommodated inflation for a sustained period of time but is intent on embarking on a disinflation. Modeling private sector learning is a challenging but nonetheless important task. Sargent (1999) provides a promising start in this direction. More work along these lines would be highly desirable.

[89] An exception is Brayton, Levin, Tyron, and Williams (1997) who present simulations of policy regime changes under different assumptions about the behavior of private sector expectations.

Appendix: The General Solution under Commitment

At time t, the central bank commits to a state contingent sequence for x_{t+i} and π_{t+i} to maximize

$$\max -\frac{1}{2} E_t\left\{\sum_{i=o}^{\infty} \beta^i[\alpha\, x_{t+i}^2 + \pi_{t+i}^2]\right\}$$

subject to the short-run aggregate supply curve

$$\pi_{t+i} = \lambda\, x_{t+i} + \beta\, E_t\{\pi_{t+1+i}\} + u_{t+i}$$

with

$$u_{t+i} = \rho u_{t+i-1} + \varepsilon_{t+i}$$

Following Currie and Levine (1993) and Woodford (1998), form the Lagrangian:

$$\max -\frac{1}{2} E_t\left\{\sum_{i=o}^{\infty} \beta^i\{[\alpha x_{t+i}^2 + \pi_{t+i}^2] + \phi_{t+i}[\pi_{t+i} - \lambda x_{t+i} - \beta\pi_{t+1+i} - u_{t+i}]\}\right\}$$

where $\frac{1}{2}\phi_{t+i}$ is the multiplier associated with the constraint at $t+i$.

The first order necessary conditions yield:

$$\alpha\, x_{t+i} - \frac{\lambda}{2}\phi_{t+i} = 0, \quad \forall i \geq 0$$

$$\pi_{t+i} + \frac{1}{2}\phi_{t+i} - \frac{1}{2}\phi_{t+i-1} = 0, \quad \forall i \geq 1$$

$$\pi_t + \frac{1}{2}\phi_t = 0$$

Combining the first order necessary conditions to eliminate ϕ_{t+i} then yields the optimality conditions

$$x_{t+i} - x_{t+i-1} = -\frac{\lambda}{\alpha}\pi_{t+i}, \quad \forall i \geq 1$$

$$x_t = -\frac{\lambda}{\alpha}\pi_t$$

Substituting the optimality conditions in the aggregate supply curve to eliminate π_{t+i} then yields a stochastic difference equation for x_t:

$$x_t = a\, x_{t-1} + a\beta\, E_t\{x_{t+1}\} - \frac{\lambda a}{\alpha} u_t$$

where $a \equiv \dfrac{\alpha}{\alpha(1+\beta)+\lambda^2}$. The stationary solution to this difference equation is given by:

$$x_t = \delta\, x_{t-1} - \frac{\lambda\delta}{\alpha(1-\delta\beta\rho)} u_t \qquad (8.1)$$

where $\delta \equiv \dfrac{1-\sqrt{1-4\beta a^2}}{2a\beta} \in (0,1)$, implying the process for x_t is stable. Substituting the solution for x_t

in the aggregate supply curve then yields a solution for π_t.

$$\pi_t = \delta\,\pi_{t-1} + \frac{\delta}{(1-\delta\beta\rho)}(u_t - u_{t-1})$$

Since $\pi_t = p_t - p_{t-1}$, the solution implies a stationary process for the price level:

$$p_t = \delta\,p_{t-1} + \frac{\delta}{(1-\delta\beta\rho)}u_t$$

The stationary behavior of the price level results from the fact that the optimality condition effectively has the central bank adjust demand in response to movements in the price level relative to trend. Given $\pi_t = p_t - p_{t-1}$, the optimality condition may be expressed as

$$x_{t+i} = -\frac{\lambda}{\alpha}p_{t+i} \qquad \forall i \geq 1$$

Thus, for example, the central bank contracts demand when the price level rises above trend: hence, the trend-reverting behavior of the price level.

References

Aiyagari, Rao S. and R. Anton Braun. 1998. "Some Models to Guide the Fed," *Carnegie-Rochester Conf. Ser. Public Policy*, 48, pp. 1–42.

Ball, Laurence. 1995. "Disinflation with Imperfect Credibility," *J. Monet. Econ.*, 35:1, pp. 5–23.

———. 1997. "Efficient Rules for Monetary Policy," 1997. NBER Working Paper 5952.

———. 1998. "Policy Rules for Open Economies," NBER Working Paper 6760.

Barro, Robert J. and David B. Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Model," *J. Polit. Econ.*, 91:4, pp. 589–610.

Benhabib, Jess; Stephanie Schmitt-Grohe, and Martin Uribe. 1998. "The Perils of Taylor Rules," mimeo, New York U.

Bernanke, Ben S. and Alan Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission," *Amer. Econ. Rev.*, 82:4, pp. 901–21.

Bernanke, Ben. S. and Mark Gertler. 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," *J. Econ. Perspect.*, 9:2, pp. 27–48.

Bernanke, Ben. S.; Mark Gertler, and Simon Gilchrist. 1998. "The Financial Accelerator in a Quantitative Business Cycle Framework," NBER Working Paper 6455. Forthcoming in *The Handbook of Macroeconomics*. John Taylor and Michael Woodford, eds.

Bernanke, Ben S.; Mark Gertler, and Mark Watson. 1997. "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Pap. Econ. Act.*, 0:1, pp. 91–142.

Bernanke, Ben S. and Ilian Mihov. 1997. "What Does the Bundesbank Target?" *Europ. Econ. Rev.*, 41:6, pp. 1025–53.

———. 1998. "Measuring Monetary Policy," *Quart. J. Econ.*, 113:3, pp. 869–902.

Bernanke, Ben S. and Frederic Mishkin. 1997. "Inflation Targeting: A New Framework for Monetary Policy?" *J. Econ. Perspect.*, 11:2, pp. 97–116.

Bernanke, Ben S. and Michael Woodford. 1997. "Inflation Forecasts and Monetary Policy," *J. Money, Credit, Banking*, 29:4, pp. 653–84.

Blanchard, Olivier J. 1997. *Macroeconomics*. Upper Saddle River, NJ: Prentice-Hall.

Blinder, Alan S. 1997. "What Central Bankers Can Learn from Academics—and Vice-Versa," *J. Econ. Perspect.*, 11:2, pp. 3–19.

Bomfin, Antulio N. and Glenn D. Rudebusch. 1997. "Opportunistic and Deliberate Disinflation Under Imperfect Credibility," mimeo, Fed. Res. Bank San Francisco.

Brainard, William C. 1967. "Uncertainty and the Effectiveness of Policy," *Amer. Econ. Rev.*, 57, pp. 411–25.

Brayton, Flint; Andrew Levin, Ralph Tryon, and John C. Williams. 1997. "The Evolution of Macro Models at the Federal Reserve Board," *Finance Econ. Discuss. Paper Series*, 1997–29, Fed. Res. Board.

Calvo, Guillermo. 1983. "Staggered Prices in a Utility Maximizing Framework," *J. Monet. Econ.*, 12:3, pp. 383–98.

Campbell, John Y. 1995. "Some Lessons from the Yield Curve," *J. Econ. Perspect.*, 9:3, pp. 129–52.

Carlstrom, Charles T. and Timothy S. Fuerst. 1995. "Interest Rate Rules vs. Money Supply Rules: A Welfare Comparison in a Cash-in-Advance Model," *J. Monet. Econ.*, 36:2, pp. 247–68.

Chari, V.V.; Lawrence J. Christiano, and Martin Eichenbaum, 1998. "Expectation Traps and Discretion," *J. Econ. Theory*, 81:2, pp. 462–92.

Christiano, Lawrence J. and Christopher J. Gust. 1999. "Taylor Rules in a Limited Participation Model," NBER Working Paper 7017.

Christiano, Lawrence J.; Martin Eichenbaum, and Charles Evans. 1996. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds," *Rev. Econ. Statist.*, 78:1, pp. 16–34.

———. 1997. "Sticky Price and Limited Participation Models of Money: A Comparison," *Europ. Econ. Rev.*, 41:6, pp. 1201–49.

———. 1998. "Monetary Policy Shocks: What Have We Learned and To What End?" NBER Working Paper 6400.

Clarida, Richard and Mark Gertler. 1997. "How the Bundesbank Conducts Monetary Policy," in *Reducing Inflation: Motivation and Strategy*. Christina Romer and David Romer, eds. Chicago: NBER, pp. 363–412.

Clarida, Richard; Jordi Galí, and Mark Gertler. Forthcoming. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quart. J. Econ.*

———. 1998. "Monetary Policy Rules in Practice: Some International Evidence," *Europ. Econ. Rev.*, 42:6, pp. 1033–67.

Currie, David and Paul Levine. 1993. *Rules, Repu-*

tation and Macroeconomic Policy Coordination. Cambridge: Cambridge U. Press.

DeLong, J. Bradford. 1997. "America's Peacetime Inflation: The 1970s," in *Reducing Inflation: Motivation and Strategy*. C. Romer and D. Romer, eds. Chicago: NBER, pp. 247–80.

Erceg, Christopher J.; Dale W. Henderson, and Andrew T. Levin. 1998. "Output-Gap and Price Volatilities: Reaffirming Tradeoffs in an Optimizing Model," mimeo, Fed. Res. Board.

Estrella, Arturo and Frederic S. Mishkin. 1997. "Is There a Role for Monetary Aggregates in the Conduct of Monetary Policy?" *J. Monet. Econ.*, 40:2, pp. 279–304.

———. 1999. "Re-Thinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty," forthcoming in *Monetary Policy Rules*. John B. Taylor, ed.

Faust, John and Lars E. O. Svensson. 1998. "Credibility and Transparency: Monetary Policy with Unobservable Goals," mimeo, Fed. Res. Board.

Feldstein, Martin S. 1997. "The Costs and Benefits of Going from Low Inflation to Price Stability," in *Reducing Inflation: Motivation and Strategy*. C. Romer and D. Romer, eds. Chicago: NBER, pp. 123–66.

Feldstein, Martin and James H. Stock. 1996. "Measuring Money Growth when Financial Markets Are Changing," *J. Monet. Econ.*, 37:1, pp. 3–27.

Fischer, Stanley. 1995. "Modern Approaches to Central Banking. NBER Working Paper 5064.

Friedman, Benjamin M. 1990. "Targets and Instruments of Monetary Policy," in *Handbook of Monetary Economics*. Friedman and Frank Hahn, eds. Amsterdam: North-Holland.

Friedman, Benjamin M. and Kenneth N. Kuttner. 1996. "A Price Target for U.S. Monetary Policy? Lessons from the Experience with Money Growth Targets," *Brookings Pap. Econ. Act.*, 1, pp. 77–146.

Fuhrer, Jeffrey C. 1996. "Towards a Compact, Empirically Verified Rational Expectations Model for Monetary Policy Analysis," mimeo, Federal Reserve Bank of Boston.

——→.. 1997a. "Inflation/Output Variance Tradeoffs and Optimal Monetary Policy," *J. Money, Credit, Banking*, 29:2, pp. 214–23.

——→.. 1997b. "The (Un)Importance of Forward Looking Behavior in Price Setting," *J. Money, Credit, Banking*, 29, Aug., pp. 338–50.

Fuhrer, Jeffrey C., and George R. Moore. 1995a. "Inflation Persistence," *Quart. J. Econ.*, 440, Feb., pp. 127–59.

———. 1995b. " Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output," *Amer. Econ. Rev.*, 85:1, pp. 219–39.

Galí, Jordi, 1992. "How Well Does the IS/LM Model Fit Post-War U.S. Data?" *Quart. J. Econ.*, 92, pp. 709–38.

Galí, Jordi and Mark Gertler. Forthcoming. "Inflation Dynamics: A Structural Econometric Analysis," *J. Monet. Econ.*

———. 1999. "Rules vs. Discretion Revisited: The Gains from Commitment in the Absence of an Inflationary Bias," work in progress.

Gertler, Mark. 1996. "Comments on Friedman and Kuttner," *Brookings Pap. Econ. Act.*, 1, pp. 77–125.

Goodfriend, Marvin. 1991. "Interest Rates and the Conduct of Monetary Policy," *Carnegie-Rochester Conf. Ser. Public Policy*, 34, pp. 7–37.

Goodfriend, Marvin and Robert G. King. 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy," in *NBER Macroeconomics Annual*. Ben Bernanke and Julio Rotemberg, eds.

→ Gordon, Robert J. 1997. "The Time-Varying NAIRU and Its Implications for Economic Policy," *J. Econ. Perspect.*, *11*, Winter, pp. 11–32.

Haldane, Andrew. 1995. *Inflation Targeting*. Bank of England.

Hall, Robert E. 1997. "Comments on Shiller," *Brookings Pap. Econ. Act.*, pp. 219–23.

Hall, Robert E. and N. Gregory Mankiw. 1994. "Nominal Income Targeting," in *Monetary Policy*. N. G. Mankiw, ed. Chicago: U. Chicago Press.

Henderson, Dale W. and Warren J. Mckibbon. 1993. "A Comparison of Some Basic Monetary Policy Regimes for Open Economies," *Carnegie-Rochester Conf. Ser. Public Policy*, 39, pp. 221–317.

Huizinga, John and Frederic S. Mishkin. 1986. "Monetary Policy Regime Shifts and the Unusual Behavior of Real Interest Rates," *Carnegie-Rochester Conf. Ser. Public Policy*, 24, pp. 231–74.

→ Ireland, Peter N. 1996a. "The Role of Countercyclical Monetary Policy," *J. Polit. Econ.*, 104:4, pp. 704–23.

———. 1996b. "Expectations, Credibility, and Time-Consistent Monetary Policy," mimeo, Rutgers U.

Jovanovic, Boyan and Masako Udea. 1997. "Contracts and Money," *J. Polit. Econ.*, 105, Aug., pp. 700–708.

Kerr, William and Robert G. King. 1996. "Limits on Interest Rate Rules in the IS Model," *Econ. Quart.*, LXXXII, pp. 47–76.

King, Robert G. and Alexander L. Wolman. 1996. "Inflation Targeting in a St. Louis Model of the 21st Century," NBER Working Paper 5507.

———. 1998. "What Should the Monetary Authority Do When Prices Are Sticky?" Forthcoming in *Monetary Policy Rules*. John B. Taylor, ed.

→ Kydland, Finn E. and Edward C. Prescott. 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *J. Polit. Econ.*, 85, pp. 473–91.

→ Leeper, Eric M. 1991. "Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies," *J. Monet. Econ.*, 27. pp. 129–47.

Leeper, Eric M.; Christopher Sims, and Tao Zha. 1996. "What Does Monetary Policy Do?" *Brookings Pap. Econ. Act.*, 2, pp. 1–63.

Lown, Cara S. and Robert Rich. 1997. "Is There

an Inflation Puzzle?" *Fed. Res. Bank New York Econ. Policy Rev.*, 3:4, pp. 51–69.

McCallum, Bennett T. 1988. "Robustness Properties of a Rule for Monetary Policy," *Carnegie-Rochester Conf. Ser. Public Policy*, 29, pp. 173–204.

———. 1997a. "Crucial Issues Concerning Central Bank Independence," *J. Monet. Econ.*, 39, June, pp. 99–112.

———. 1997b. "Issues in the Design of Monetary Policy Rules," NBER Working Paper 6016. Forthcoming in *The Handbook of Macroeconomics*.

———. 1997c. "The Alleged Instability of Nominal Income Targeting," NBER Working Paper 6291.

McCallum, Bennett T. and Edward Nelson. 1997. "An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis," NBER Working Paper 5875.

Monacelli, Tommaso. 1999. "Into the Mussa Puzzle: Monetary Policy Regimes and the Real Exchange Rate in a Small Open Economy," mimeo, Boston College.

Onatski Alexei, and James H. Stock. 1999. "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy," mimeo, Harvard U.

Orphanides, Athanasios. 1998. "Monetary Policy Evaluation with Noisy Information," mimeo, Fed. Res. Board.

Orphanides, Athanasios and David Wilcox. 1996. "The Opportunistic Approach to Disinflation," *Finance Econ. Discuss. Paper Series*, 96–24, Fed. Res. Board.

Orphanides, Athanasios; David H. Small, Volker Weiland, and David W. Wilcox. 1997. "A Quantitative Exploration of the Opportunistic Approach to Disinflation," *Finance Econ. Discuss. Paper Series*, 36, Fed. Res. Board, Washington D.C.

Persson, Torsten and Guido Tabellini. 1997. "Political Economics and Macroeconomic Policy," mimeo, Inst. for International Economics, Stockholm.

Poole, William. 1970. "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model," *Quart. J. Econ.*, 84, pp. 197–216.

Roberts, John M. 1997. "Is Inflation Sticky?" *J. Monet. Econ.*, 39, pp. 173–96.

Rogoff, Kenneth. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quart. J. Econ.*, 100:4, pp. 1169–89.

———. 1987. "Reputational Constraints on Monetary Policy," *Carnegie-Rochester Conf. Ser. Public Policy*, 26, 141–81.

Romer, Christina D. and David H. Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomics Annual*.

Rotemberg, Julio. 1996. "Prices, Hours and Output: An Empirical Analysis Based on a Sticky Price Model," *J. Monet. Econ.*, 37, pp. 505–34.

Rotemberg, Julio and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*. Ben Bernanke and J. Rotemberg, eds.

———. 1999. "Interest Rate Rules in an Estimated Sticky Price Model," in *Monetary Policy Rules*. John B. Taylor, ed. Forthcoming.

Rudebusch, Glenn D. 1995. "Federal Reserve Interest Rate Targeting, Rational Expectations and the Term Structure," *J. Monet. Econ.* 35: pp. 245–74.

Rudebusch, Glenn D. and Lars Svensson. 1998 forthcoming. "Policy Rules for Inflation Targeting," in *Monetary Policy Rules*. John Taylor, ed.

Sack, Brian. 1997a. "Uncertainty and Gradual Monetary Policy," mimeo, Fed. Res. Board.

———. 1997b. "Does the Fed Act Gradually? A VAR Analysis," mimeo, Fed. Res. Board, May.

Sargent, Thomas J. 1998. "The Conquest of American Inflation," mimeo, Econ. Dept., Stanford U.

———. 1999. "Comment on Ball," forthcoming in *Monetary Policy Rules*. John B. Taylor, ed.

Sargent, Thomas J. and Neil Wallace. 1975. "Rational Expectations, the Optimal Monetary Instrument and the Optimal Money Supply Rule," *J. Polit. Econ.*, 83, pp. 241–54.

Sbordone, Argia. 1998. "Prices and Unit Labor Costs: A New Test of Sticky Prices," mimeo, Rutgers U.

Shiller, Robert J. 1997. "Public Resistance to Indexation: A Puzzle," *Brookings Pap. Econ. Act.*, pp. 159–212.

Sims, Christopher. 1994. "A Simple Model for the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy," *Econ. Theory*, 4, pp. 381–99.

Svensson, Lars E. O. 1997a. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *Europ. Econ. Rev.*, 41, June, pp. 1111–47.

———. 1997b. "Inflation Targeting: Some Extensions," NBER Working Paper 5962, March, forthcoming in *Scand. J. Econ.*

———. 1997c. "Optimal Inflation Targets, Conservative Central Banks, and Linear Inflation Contracts," *Amer. Econ. Rev.*

———. 1998. "Inflation Targeting as a Monetary Policy Rule," NBER Working Paper 6790, forthcoming in *J. Monet. Econ.*

———. 1998. "Open-Economy Inflation Targeting," mimeo, forthcoming in *J. Int. Econ.*

Svensson, Lars E.O. and Leonardo Leiderman. 1995. *Inflation Targets*. London: CEPR.

Taylor, John B. 1979. "Estimation and Control of Macroeconomic Model with Rational Expectations," *Econometrica*, 47, Sept., pp. 1267–86.

———. 1993a. "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conf. Ser. Public Policy*, 39, pp. 195–214.

———. 1993b. *Macroeconomic Policy in a World Economy*. NY: W.W. Norton.

———, ed. 1999a. *Monetary Policy Rules*. Chicago: U. Chicago Press. Forthcoming.

———. 1998b. "An Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*. John B. Taylor, ed. Forthcoming

Theil, Henri. 1961. *Economic Forecasts and Policy*, 2nd ed. Amsterdam: North-Holland.

Tinbergen, Jan. 1952. *On the Theory of Economic Policy*, 2nd ed. Amsterdam: North-Holland.

Walsh. Carl. 1998. *Monetary Theory and Policy*. MIT Press.

Wieland, Volker. 1997. "Monetary Policy and Uncertainty about the Natural Unemployment Rate," mimeo, Fed. Res. Board.

Williams, John C. 1997. "Simple Rules for Monetary Policy," mimeo, Fed. Res. Board.

Wolman, Alexander. 1998. "Staggered Price Setting and the Zero Lower Bound on the Nominal Interest Rate," mimeo, Fed. Res. Bank Richmond.

Woodford, Michael. 1994a. "Nonstandard Indicators for Monetary Policy," in *Monetary Policy*. N. Gregory Mankiw ed. Chicago: U. Chicago Press, pp. 95–115.

———. 1994b. "Monetary Policy and Price Level Determinacy in a Cash-in-Advance Economy," *Econ. Theory*, 4, pp. 345–80.

———. 1996. "Control of the Public Debt: A Requirement for Price Stability?" NBER Working Paper 5684.

———. 1998. "Optimal Monetary Policy Inertia," mimeo, Princeton U.

Yun, Tack. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *J. Monet. Econ.*, 37, pp. 345–70.

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Implications of Inflation Dynamics for Monetary Policy Strategies

Hess Chung, Callum Jones, Antoine Lepetit, Fernando M. Martin

2025-072

Please cite this paper as:
Chung, Hess, Callum Jones, Antoine Lepetit, and Fernando M. Martin (2025). "Implications of Inflation Dynamics for Monetary Policy Strategies," Finance and Economics Discussion Series 2025-072. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.072.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.



Implications of Inflation Dynamics for Monetary Policy Strategies

Hess Chung, Callum Jones, Antoine Lepetit, Fernando M. Martin

The analysis in this paper was presented to the Federal Open Market Committee as background for its discussion of the Federal Reserve's 2025 review of its monetary policy strategy, tools, and communications.

Abstract: This paper considers robust monetary policy strategies both in situations of low demand and low inflation and when economic developments pose a tradeoff between inflation and output stabilization. We proceed in two parts. First, our quantitative analysis suggests that asymmetric average inflation targeting can provide modest benefits over other inflation-targeting strategies when the risks associated with the effective lower bound remain significant. Second, motivated by the recent experience of persistent supply shocks and rapid increases in inflation, we describe the main qualitative features of optimal policy in circumstances when the objectives of stabilizing inflation and economic activity conflict. We find that monetary policy may allow inflation to depart from the target in response to certain supply shocks or in cases when sectoral dynamics are relevant, but that it should be ready to respond forcefully and expeditiously to large inflationary shocks or if inflation expectations are at risk of becoming unanchored.

JEL Classification: E31, E52, E58.

Keywords: Alternative monetary policy strategies, monetary policy communication, effective lower bound, supply shocks, sectoral dynamics, inflation surges.

Note: Authors' affiliations are Board of Governors of the Federal Reserve System (Chung, Jones and Lepetit), and Federal Reserve Bank of St. Louis (Martin). The views expressed are those of the authors and do not necessarily reflect the views of anyone else affiliated with the Federal Reserve System. The authors benefited from the comments and suggestions of James Clouse, Marco Del Negro, Joshua Gallin, Benjamin Johannsen, Michael Kiley, Sylvain Leduc, David López-Salido, Thomas Mertens, Edward Nelson, Anna Paulson, Trevor Reeve, Robert Tetlow, and Stacey Tevlin. The authors would like to thank Karen Page for excellent research assistance.

1. Introduction and overview

Starting in the 1990s, an increasing number of central banks around the world adopted inflation-targeting strategies. Under these strategies, central banks set an explicit inflation target to help deliver low and stable inflation. Under "flexible" inflation targeting, central banks balance medium-run inflation stabilization against other objectives, including stabilizing economic activity or employment. The Federal Reserve adopted flexible inflation targeting in 2012, with the announcement of a 2 percent longer-run inflation objective in its first Statement on Longer-Run Goals and Monetary Policy Strategy.

Prior to the Global Financial Crisis, inflation targeting had generally performed well, although it was well understood that the effective lower bound (ELB) on nominal interest rates posed challenges to these strategies as conventionally pursued.[1] The difficulty of providing sufficient accommodation when the policy rate is constrained by the ELB was readily apparent during the protracted recovery from the Great Recession. In response, many central banks employed unconventional policy instruments, such as balance sheet policies and forward guidance to provide additional policy accommodation.

In the U.S., the decade following the Great Recession was characterized by below-target inflation, low nominal interest rates, concern about future binding ELB episodes, and the risk that longer-term inflation expectations might drift downwards. In light of these experiences and in line with the practices of other major central banks, the Federal Reserve conducted the first review of its monetary policy framework in 2019-20 to assess whether changes in its monetary policy strategy could improve outcomes when the federal funds rate is at risk of being constrained by the ELB. A key outcome of that review was the adoption of a new element in its 2020 Statement on Longer-Run Goals and Monetary Policy Strategy: In order to achieve 2 percent inflation on average over time, the Federal Open Market Committee (FOMC) would "likely aim to achieve inflation moderately above 2 percent for some time" in situations where inflation had been running persistently below 2 percent. This strategy shares some similarities with an asymmetric variant of flexible average inflation targeting, a class of strategies under which the central bank aims to stabilize average inflation over some time period under certain conditions. As widely discussed in the literature, strategies in this class help to move inflation towards target on average, better anchor inflation expectations, and improve economic outcomes when the policy rate is at the ELB.

The experience since 2020 presented different challenges to monetary policy than had been the primary focus during the previous framework review. As described further in the companion papers by Hajdini and others (2025) and Lipińska and others (2025), widespread

[1] See Bernanke and others (1999) for an early treatment of the benefits of inflation targeting. More recently, Borio and Chavaz (2025) and Mishkin and Kiley (2025) emphasize how inflation targeting became even more prevalent after the financial crisis. See Krugman (1998) and Reifschneider and Williams (2000) on challenges posed by the ELB.

shortages, supply-chain disruptions, and reductions in labor supply contributed to depressed economic activity and increased inflationary pressures during that period. The recent higher incidence of supply shocks, a more important role for sectoral dynamics, and the possibility of rapid increases in inflation underscore that the objectives of stabilizing inflation and economic activity can be in conflict and that monetary policy must balance them.

This paper considers robust monetary policy strategies both in situations of low demand and inflation, and when developments pose a tradeoff between inflation and output stabilization objectives. In the first part of the paper, we analyze the main costs and benefits of asymmetric flexible average inflation targeting compared with other inflation targeting strategies through simulations using the FRB/US model, focusing on the performance of these strategies at the ELB. In the second part, motivated by the experience since 2020, we review the tradeoffs that supply shocks, sectoral dynamics, and sudden inflation surges may pose for monetary policy. In addition, we describe the main qualitative features of optimal policy in such circumstances, as well as the possible implications for the relative benefits of flexible average inflation targeting.

For the purposes of this paper, we mainly focus on two inflation targeting strategies: (1) flexible inflation targeting (FIT) where monetary policy follows the prescriptions of a standard inertial Taylor rule at all times, and (2) asymmetric flexible average inflation targeting (FAIT), under which policy seeks to stabilize average inflation when that average has fallen persistently below 2 percent, and reverts to FIT at all other times. Where appropriate, we also consider modified versions of FIT in which the central bank uses forward guidance when the ELB binds and versions of both FIT and FAIT that adopt a "shortfalls" approach to economic slack.

The two main parts of this paper and the key conclusions are as follows:

Evaluating Inflation Targeting Strategies

- FAIT can improve economic outcomes relative to FIT when inflation runs persistently below target because it prescribes a lower expected federal funds rate path relative to FIT for some time. During recessions when the ELB binds, the decline in the inflation rate under FAIT is smaller than under FIT, and inflation converges back to the target faster. By design, following a period of low inflation, FAIT calls for inflation to run somewhat above target for some time; in our simulations, overshooting tends to be relatively modest.
- Some of the shortcomings of FIT at the ELB can be addressed with modifications that provide additional accommodation when the ELB binds. A leading example is the specification of thresholds as a form of forward guidance that delays exit from the ELB for an extended period beyond what would be called for in the absence of the thresholds. In our simulations, modifications to FIT must keep the funds rate at the ELB for a considerable period in order to deliver benefits comparable with FAIT. As with FAIT, FIT strategies modified in this way often imply inflation overshooting the target modestly.
- Fully realizing the benefits of either FAIT or modifications of FIT requires that the public views the strategies as credible and adequately understands their implications. Achieving

these prerequisites may be challenging in practice. Incomplete descriptions of these strategies may have the advantage of retaining flexibility, maintaining credibility, and simplifying communications. However, incomplete descriptions may also be less effective at shaping public expectations, reducing the effectiveness of the policies.

Some Principles of Optimal Monetary Policy in Response to Inflationary Shocks

- Provided inflation expectations remain well anchored, optimal monetary policy allows inflation to depart from the target in response to certain supply shocks or in cases when sectoral dynamics are relevant, for instance, when large shocks originate in specific sectors or when the disparate effect of aggregate shocks across sectors is significant. In such cases, fully and promptly stabilizing inflation may come at a substantial cost for economic activity.
- Large inflationary shocks may also cause the public to become highly sensitive to inflation dynamics, raising the risk that sudden inflation surges may develop. These nonlinear inflation dynamics may lead to a persistent upward spiral between inflation and inflation expectations if not countered by a strong and expeditious policy response.

2. Evaluating inflation targeting strategies

In this section, we explore the macroeconomic costs and benefits of flexible versions of two strategies: flexible inflation targeting (FIT) and flexible average inflation targeting (FAIT). In the case of FAIT, we consider an asymmetric variant, which seeks to stabilize average inflation only when that average has been running below target, as might occur when the effective lower bound (ELB) is binding.[2] In all other cases, FAIT reverts to FIT. For robustness, we also interact both FIT and FAIT with the shortfalls approach to the maximum employment objective.

2.1 FIT versus FAIT: A quantitative evaluation

While FAIT has been part of the FOMC's strategy since 2020, the historical record since then has not been particularly informative about the relative performance of FIT versus FAIT because the economy has been far from the ELB for most of this time and inflation has generally been above target. Accordingly, we focus on evaluating the relative performance of FIT and FAIT quantitatively in the FRB/US model. We represent these strategies in the form of simple monetary policy rules, calibrated to be broadly representative of the policy prescriptions that might be expected from these two strategies. These rules have been specified for illustrative purposes only and our results should be interpreted as indicative of general tendencies in outcomes, while the quantitative differences may vary depending on the exact rule specifications.

Concretely, as described in table 1 and in more detail in appendix A.1.1, we represent FIT by a standard inertial Taylor rule, according to which, when not constrained by the ELB, the federal funds rate is set as a function of its level over the past quarter, the output gap, and four-

[2] FAIT belongs to the class of "make-up strategies" that have been shown to be effective when the ELB constrains policy. Broadly speaking, make-up strategies commit monetary policy to follow deficits in certain variables (such as inflation below target) with surpluses in those variables later (such as inflation overshooting the target). Prominent examples of such strategies in the literature are price-level targeting, average inflation targeting, and, in the context of the ELB, the forward guidance strategy described in Reifschneider and Williams (2000). The literature showing the advantages of make-up rules in modern macroeconomic models dates back at least to the early 1980s, such as Bean (1983) on nominal income targeting. More recently, theoretical explorations of optimal policy at the ELB very generally lead to history-dependent rules. These rules function, in some ways, like make-up rules. For example, as shown in Eggertsson and Woodford (2003), the optimal monetary policy response to a binding ELB is a form of price-level targeting, in which the permanent increase in the price level is related to the severity of the ELB constraint. Make-up strategies, like FAIT, that do not allow the price level to rise permanently above its trend only approximate the optimal policy. See the discussion in Levin and others (2010) for a quantitative comparison of these policies versus the optimal policy.

A number of recent papers have evaluated the performance of these strategies, among them, Kiley and Roberts (2017), Hebden and López-Salido (2018), Bernanke, Kiley, and Roberts (2019), Reifschneider and Wilcox (2019), Amano and others (2020), Arias and others (2020), Budianto, Nakata, and Schmidt (2020), Coenen, Montes-Galdón, and Schmidt (2021), Erceg, Jakab, and Lindé (2021), Gerke and others (2021) and Coulter, Duncan, and Martinez-Garcia (2022). These studies typically find, as we do, that make-up strategies can improve on FIT, especially at the ELB, although the magnitude of the improvement, and the best performing make-up strategy, varies across papers.

quarter core personal consumption expenditures (PCE) price inflation.[3] As mentioned previously, standard Taylor rules are known to perform poorly at the ELB, so we will also consider modifications of FIT that use thresholds to delay exit from the ELB beyond what would be implied by the Taylor rule.

Table 1: Specifications of Monetary Policy Rules for FIT and FAIT

FIT: $R_t^{FIT} = 0.85R_{t-1} + 0.15(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t)$

FAIT: $R_t^{FAIT} = \begin{cases} R_t^{FIT} & \text{if } \bar{\pi}_t \geq \pi^{LR} \\ R_t^{FIT} + 0.15 * 8(\bar{\pi}_t - \pi^{LR}) & \text{if } \bar{\pi}_t < \pi^{LR} \end{cases}$

Notes: R_t is the annualized net federal funds rate, $\bar{\pi}_t$ is 32-quarter average inflation, $ygap_t$ is the output gap, π_t is the 4-quarter percent change in core PCE prices, π^{LR} is the inflation target, and r^{LR} is the long-run real funds rate.

A number of policy rules can be regarded as forms of FAIT, differing in the conditions under which policy starts and ceases to target average inflation.[4] In this paper, the FAIT rule targets average inflation whenever average inflation is below 2 percent and is identical to the calibrated Taylor rule used to represent FIT when average inflation is at or above 2 percent. Formally, we select a rule to represent FAIT that is identical to the standard Taylor rule but adds a term (in red in table 1) that reacts to the eight-year trailing average inflation rate, when this average runs below target.[5]

[3] We focus on rules with a symmetric output gap response, but, below, also consider rules that respond only to shortfalls in activity.

Allowing for inertia in the FIT rule blurs the conceptual distinction between FIT and make-up strategies. Just like certain make-up strategies, very inertial FIT rules react persistently to past inflation. For example, a FIT rule with a high degree of inertia could deliver outcomes very close to those obtained under a price-level target. In this paper, we choose a degree of inertia that is roughly in line with historical FOMC behavior. With this degree of inertia, a meaningful quantitative difference between FIT and make-up strategies still exists.

[4] Among alternative rules that could also be viewed as falling under FAIT are the proposed rule in Bernanke (2017), under which policy targets average inflation until such time as the rule departs from the ELB, and the temporary price-level and average-inflation targeting rules described in Hebden and López-Salido (2018) and Chung and others (2019), in which policy starts to target average inflation when the ELB binds and continues to target average inflation until the average is 2 percent. The Bernanke rule has been evaluated by Hebden and López-Salido (2018), who find that the rule performs well compared with FIT rules, and by Erceg, Jakab, and Lindé (2021), who find that the rule provides only relatively modest additional stimulus in an estimated model of the euro area. While we focus on FAIT rules that respond dynamically to the average inflation gap, as shown in Mertens and Williams (2019), the downward bias due to the ELB could also be corrected by adopting a FIT rule that aims to stabilize inflation at a constant level above the inflation target when the policy rate is not at the ELB.

[5] The length of the AIT horizon significantly affects the outcomes generated by the rule. In theoretical models where the public has complete understanding of the rule and its consequences, the optimal horizon is

We evaluate the performance of FIT and FAIT using stochastic simulations of the FRB/US model.[6] The simulations assume model-consistent expectations on the part of asset-market participants and wage- and price-setters, that is, the expectations of these agents are identical with the modal forecast of the model, in the absence of shocks. Figure 1 plots the median of the distribution of outcomes, conditional on shocks that would drive the economy into a recession and the federal funds rate to the ELB.[7] In these simulations, recessions severe enough to drive the funds rate to the ELB are largely driven by adverse shocks to demand. In the median recession, under FIT (the red solid line), the unemployment rate rises from roughly 4 percent in the year before the recession to around 6½ percent at the peak, while inflation drops to 1.3 percent at its trough and remains around 1½ percent for several years thereafter. Under FAIT (the blue dashed line), the median unemployment and inflation outcomes are appreciably better, with both variables closer to target.

FAIT achieves a faster convergence of inflation to target because of the additional accommodation it provides when average inflation is below target. While FAIT departs from the ELB at roughly the same time as FIT, it is substantially more accommodative after liftoff. This additional accommodation is clearly apparent in the lower trajectory for the expected 10-year average real funds rate compared with FIT. Because the eight-year average inflation rate returns to two percent very gradually, policy prescriptions under FAIT are persistently more accommodative than under FIT. Higher expected inflation than under FIT further depresses real rates. Wage- and price-setters, anticipating the higher level of demand (and hence inflation), raise nominal wages and prices earlier, even when the federal funds rate is constrained by the ELB and demand is weak.

typically very long. In models featuring imperfect understanding (Amano and others, (2020), Budianto, Nakata, and Schmidt (2020)), the horizon is typically shorter due to effects similar to those we discuss below in section 2.2. In the FRB/US model, because both output and inflation are quite inertial, a relatively long horizon is necessary to generate meaningful differences from FIT.

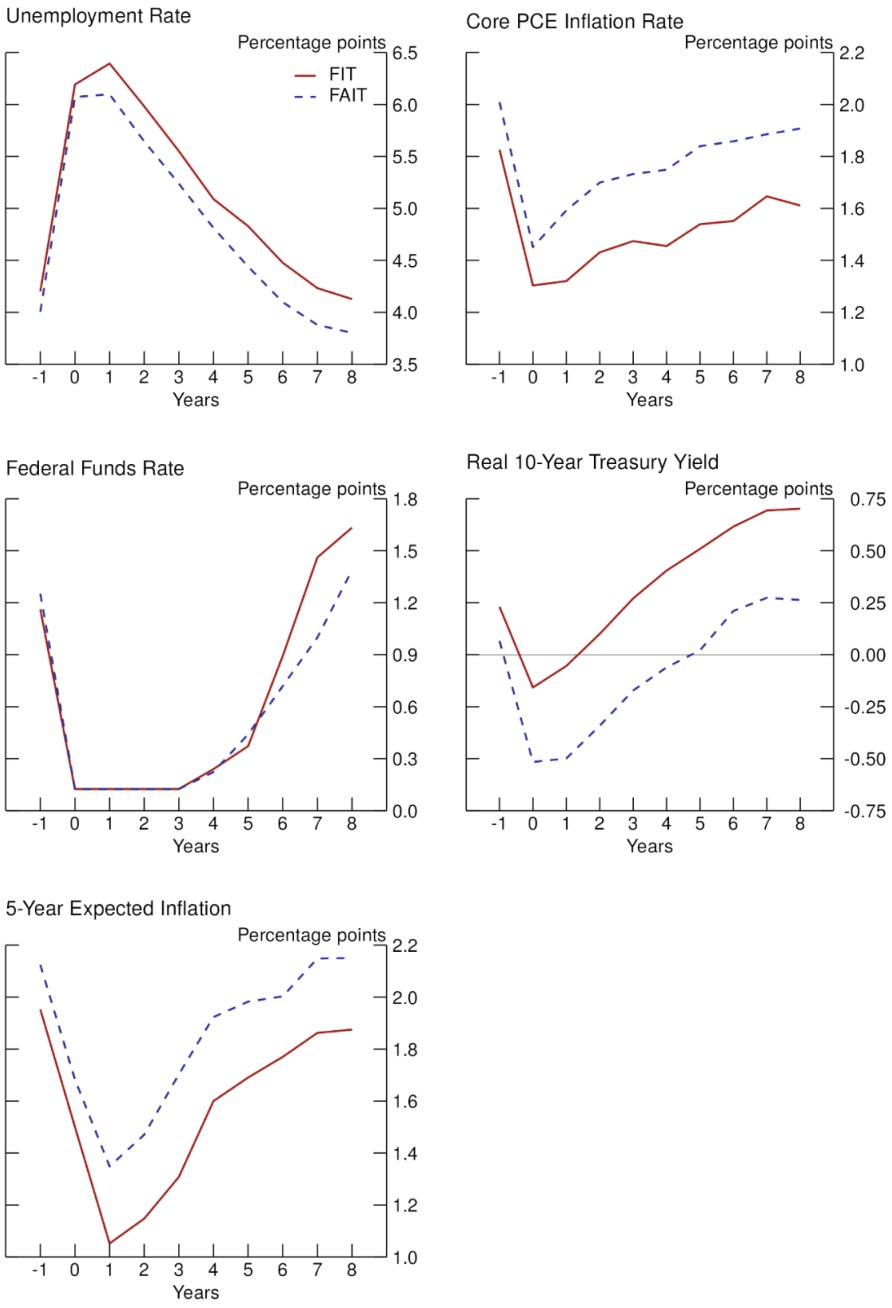
[6] The shocks for this exercise are drawn from FRB/US model residuals from 1969q1 to 2024q4, excluding the pandemic (2020q1 to 2021q2). The baseline trajectories for these simulations are taken from the public FRB/US dataset, starting in 2030q1 (near the steady-state, in which the long-run real funds rate is 1 percent).

Unless otherwise noted, the rules are simulated in the absence of forward guidance strategies at the ELB; that is, after hitting the ELB, the federal funds rate is raised in the quarter when the rule prescription is above the ELB. Appendix A.1 describes further technical details behind the simulations, as well as simulations of recessions under VAR and fully model-consistent expectations comparable to those described in this section.

We abstract from active balance-sheet policies, which could be incorporated into either strategy.

[7] For the purposes of this figure, a recession is defined as in González-Astudillo and Vilán (2019): an event in which the unemployment rate rises in four consecutive quarters along with at least two—not necessarily consecutive—quarters of negative gross domestic product (GDP) growth. The median recession in our simulations does not imply a binding ELB and we therefore focus here on events that also cause the ELB to bind. All rules are simulated using the same randomly drawn shock paths. Median outcomes are conditional on entering a recession along with a binding ELB under FIT in the fortieth quarter of the simulation; this period is labelled as period 0 in the figures.

Figure 1: Median outcomes at the ELB in a recession, alternative policy rules



Notes: Authors' estimates. Median outcomes conditional on entering a recession along with a binding ELB under FIT in the fortieth quarter of the simulation; this period is labelled as period 0 in the figures. All rules are simulated using the same randomly drawn shock paths.

By design, following a period of low inflation, FAIT calls for inflation to run somewhat above target for some time. In these simulations, monetary policy affects inflation only with a considerable lag. As a result, the expected overshooting occurs many years after the ELB binds

and is not visible in figure 1.[8] Moreover, by the time the overshooting occurs, the eight-year average inflation rate is only slightly above the target. Consequently, the overshooting tends to be relatively modest. As shown in figure 1, under FAIT, the median expected five-year-ahead average inflation rate rises a bit above 2 percent six years after the trough of the recession.[9] If policy lags were shorter, the overshoot of inflation would occur sooner.

Both an extensive research literature and the record of central bank practice after the Great Recession suggest many different modifications to the standard Taylor rule approach that might improve the performance of FIT at the ELB.[10] A prominent example of such a modification is the specification of threshold conditions on economic outcomes that must be met before lifting rates off the ELB, delaying exit beyond what would naturally be called for by the Taylor rule and thus providing a "lower for longer" policy when the ELB binds. Results in the literature indicate that the use of thresholds may be able to effectively address the weaknesses of FIT at the ELB, if exit from the ELB is delayed sufficiently beyond what would be called for by FIT alone.[11] Specifically, the literature suggests that thresholds may have to delay lift-off until either economic activity is above potential or inflation is above target in order to provide sufficient additional accommodation.[12] With such thresholds, the additional accommodation is provided by a longer period of time at the ELB rather than the lower policy path after liftoff that

[8] The onset of overshooting is also delayed by a technical issue with the perfect foresight methodology used to simulate the model. Under the assumption of perfect foresight, households and firms do not take into account the downward bias on inflation as a result of the ELB and hence do not anticipate real interest rates low enough to compensate for the bias. Consequently, the median inflation path does not overshoot as quickly as it would under fully rational expectations. Related issues arise below in our discussion of threshold-based forward guidance, as further elaborated in footnote 12.

[9] The long delay before inflation overshoots under FAIT, as well as its modest magnitude, in FRB/US simulations has been previously noted and discussed by Reifschneider and Wilcox (2019) and Arias and others (2020). Both features stem from the very gradual recovery of inflation in the model.

[10] In addition to using forward guidance to ameliorate the shortcomings of FIT, suitably chosen balance-sheet policies could also diminish the gap between FIT and FAIT.

[11] See Boneva, Harrison, and Waldron (2018) and Chung and others (2019).

[12] Appendix A.1.4 reports outcomes under an inflation threshold, which delivers roughly as much accommodation as FAIT and achieves similar outcomes for inflation. However, as further discussed in that appendix, inflation can be volatile, and, since inflation in these simulations is only a few tenths of 1 percent below the threshold, the probability of a short-lived shock causing inflation to breach the threshold is appreciable. In the simulation methodology used in generating the FRB/US simulations, agents do not account for the possible effects of these shocks when forming expectations and, as a result, significantly over-estimate the duration of the ELB episode under an inflation threshold. Boneva, Harrison, and Waldron (2018), who take account of this effect, find that an unemployment rate threshold (without an inflation escape clause) works better than an inflation threshold when demand shocks dominate the distribution of risks, but that an inflation threshold can work better in the presence of volatile supply shocks.

is prescribed by FAIT.[13] Thresholds that permit exit from the ELB earlier often fail to improve meaningfully over FIT.[14]

In typical models in the literature, the abnormally sharp drop in economic activity and, hence, inflation at the ELB causes average inflation under FIT to run below the inflation target. The magnitude of the downward bias in the average inflation rate depends on several factors, including the strength of monetary policy effects on inflation and the distribution of shocks. In addition, the probability of hitting the ELB—and hence the magnitude of the downward bias in inflation—depends importantly on the long-run level of the federal funds rate, with a higher level reducing the probability of a binding ELB. Quantitative evaluations of the downward bias in inflation differ significantly along these dimensions. For example, for an inflation target of 2 percent and a long-run real equilibrium federal funds rate of 1 percent, prominent estimates in the literature of the average inflation rate lie between 1 and 2 percent.[15] In the simulations reported in Appendix A.1.3, FAIT strategies effectively address this downward bias, even in environments where the downward bias is large under FIT.[16]

2.2 FIT versus FAIT: Robustness

Interactions with the "shortfalls" approach. The relative costs and benefits of the two strategies, FIT and FAIT, also depend on the degree to which policy responds asymmetrically to economic activity. As discussed in Bundick, Cairó and Petrosky-Nadeau (2025), a rule that does not respond to positive output gaps or to unemployment below the unemployment rate consistent with maximum employment (a "shortfalls" approach) could result in a higher average inflation rate than a similar rule that has a symmetric response. A shortfalls approach may therefore

[13] For more discussion of the tradeoffs between providing accommodation through a longer period of time at the ELB versus a slower pace of tightening after liftoff, see Erceg, Kiley, and Tetlow (2014).

[14] For example, as discussed in more detail in appendix A.1.4, in these FRB/US simulations, an unemployment rate threshold set at the natural rate of unemployment increases the expected duration of the ELB episode only an extra year or so at the trough of the recession in the median case, compared with the unmodified FIT rule, and hence it provides only very modest additional stimulus.

[15] Kiley and Roberts (2017) and Bernanke and others (2019) report average inflation around 1 percent under these conditions. A survey of model results presented as part of the European Central Bank's price stability framework review in 2021 finds that, in most models, inflation averages above 1½ percent, with a median result about mid-way between 1½ and 2 percent. Appendix A.1.3 provides further discussion of the determinants of the average inflation rate in FRB/US simulations.

[16] In theory, the probability that the funds rate is at the ELB could be either higher or lower under FAIT than under FIT, depending mainly on how strong the effects of monetary policy on inflation and output are. Because inflation is higher, on average, under FAIT than under FIT, it is possible that the funds rate could be at the ELB less frequently under FAIT than under FIT, as suggested by some results in the literature—see, for example, Mertens and Williams (2019). As shown in the FRB/US simulations in appendix A.1.3, under the assumption of model-consistent expectations for financial market participants and wage and price setters, the fraction of quarters at the ELB is slightly higher under FAIT than FIT. By contrast, assuming fully model-consistent expectations and a low degree of fiscal stabilization—and thus a more adverse effect of being constrained by the ELB—implies that the fraction of quarters at the ELB is much higher under FIT than FAIT.

partly substitute for FAIT as a strategy for supporting average inflation against downside risks. Furthermore, such an asymmetric approach may imply a more gradual tightening path after exiting the ELB, helping to provide additional accommodation when the ELB is binding. As shown in Figure 2, when FIT is combined with a shortfalls approach, inflation outcomes during a recession are closer to those seen under FAIT. When FAIT is combined with a shortfalls approach, inflation now overshoots 2 percent within seven years of the start of the episode, much earlier than under a FAIT rule that responds symmetrically to the output gap.

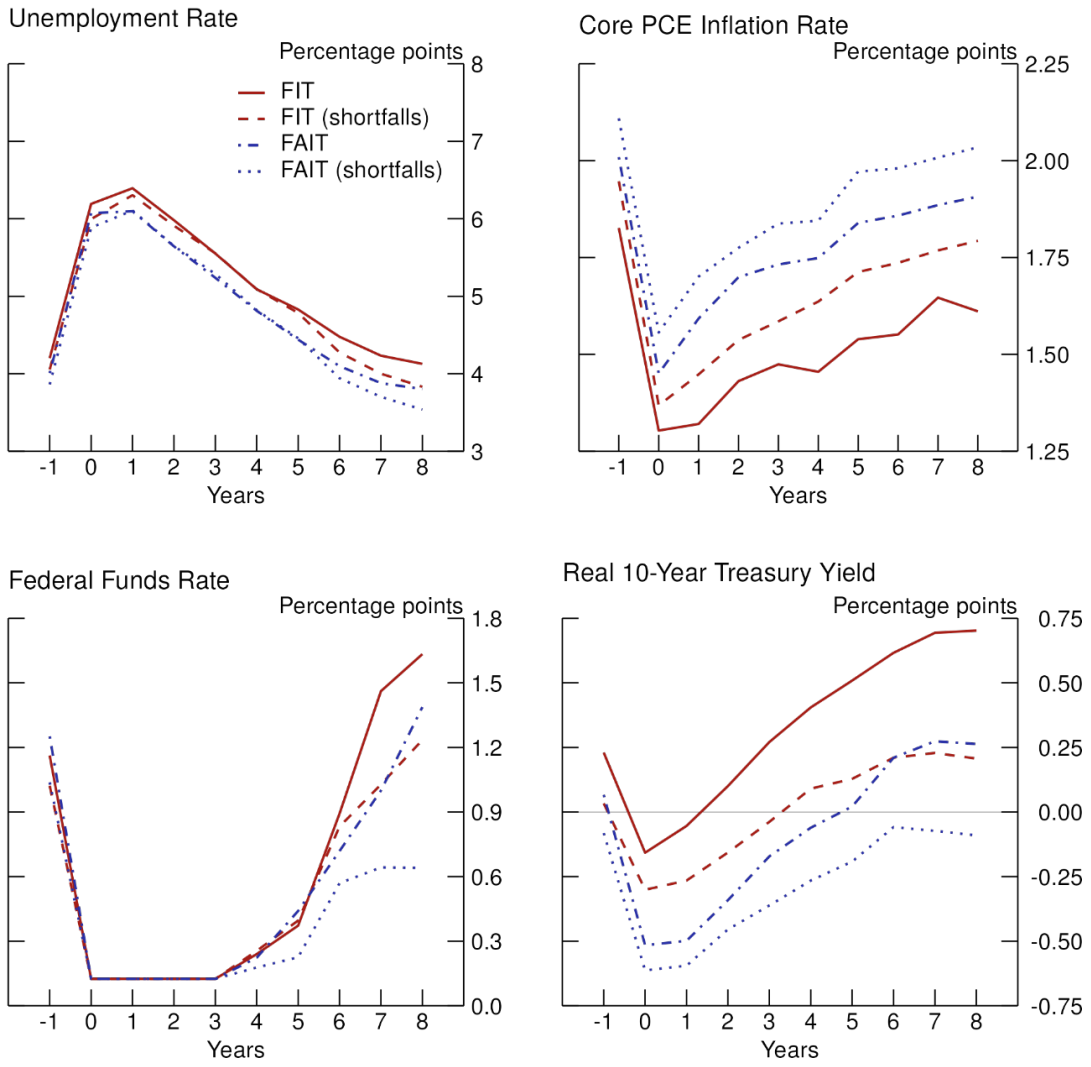
The importance of public understanding and credibility. The full benefits of FAIT materialize only if the public understands the strategy adequately and views it as credible.

First, the benefits of FAIT shown in the simulations require that, at the ELB, the public correctly anticipates that monetary policy will remain more accommodative than FIT for a number of years into the future.[17] FAIT is a more complex policy than FIT, and although describing the desired policy in a particular scenario is straightforward, communicating how the policy would work out under a variety of scenarios that may materialize over the implementation horizon is more challenging. For example, a policy that has a substantial backward-looking element, such as a long averaging horizon, would have to be qualified to allow policymakers to respond appropriately to rapid developments, such as a surge in inflation or other unforeseen events. Policymakers may wish to specify such exit conditions from FAIT. Incomplete descriptions of the policy strategy may have the advantage of retaining flexibility and simplifying communications with the public.[18] However, incomplete descriptions may also be less effective at shaping public expectations, reducing the effectiveness of the policy. Nonlinearities in inflation dynamics—discussed in more detail in the next section—could complicate this communication further, as accommodation may have to be withdrawn more rapidly than previously anticipated.

[17] Formally, the simulations assume that financial market participants form model-consistent expectations, which are then reflected in asset prices.

[18] See Bernanke (2003) for a discussion of the merits of what he terms "constrained discretion," in which "the central bank is free to do its best to stabilize output and employment in the face of short-run disturbances", (paragraph 11) subject to maintaining a strong commitment "to keeping inflation low and stable" (paragraph 6).

Figure 2: Median outcomes at the ELB in a recession, alternative shortfall rules



Notes: Authors' estimates. Median conditional on being in a recession along with a binding ELB under FIT in the fortieth quarter of the simulation.

Second, outcomes under the two strategies also depend on the extent to which monetary policy strategies can shape expectations of future activity and inflation. For example, simulations reported in appendix A.1.5 suggest that, if the public's expectations of inflation are sufficiently backward looking, FAIT does not succeed in raising inflation expectations when demand is low.[19] Therefore, inflation is only a little higher under FAIT compared with FIT.

[19] To get a sense of how the benefits and costs of the two strategies depend on how the public may form expectations, we re-run the simulations assuming that the public correctly anticipates the path of the federal funds rate but does not fully understand the implications for inflation—instead relying on a simple (backward looking)

Nonetheless, if longer-term inflation expectations are linked to average realized inflation, FAIT may still help stabilize longer-term inflation expectations by stabilizing average inflation.

Finally, reaping the benefits of FAIT requires that the rule is perceived to be credible. Under FAIT, following a period of low inflation, inflation will often overshoot the target, potentially over a prolonged period, and this overshooting must occur without impairing the public's confidence in a continued strong commitment to price stability. On account of the tension between the commitment to allow overshooting and the mandate to achieve price stability, the public may not be certain that past promises will be honored once the benefits are realized, possibly diminishing the effectiveness of the strategy. Incomplete communication about the policy reaction function, along with escape clauses, could reduce the burden of commitment, but, as previously mentioned, could also temper the effectiveness of FAIT at guiding expectations.[20]

These considerations apply with similar force to threshold-based forward guidance under FIT. For example, credibility concerns and communication challenges similar to those with FAIT would also arise for a threshold strategy under which the minimum inflation threshold for exiting the ELB was chosen at or above the target.

Summing up.

In this section, we focused on alternative monetary policy strategies that may be appropriate under the mix of risks that were prominent before the pandemic—in particular, the challenges posed by the ELB and persistently below-target inflation. These risks may continue to be relevant. However, the post-pandemic experience has highlighted some additional risks—such as a heightened prevalence of supply shocks, sectoral developments, and nonlinearities in inflation dynamics. In the years ahead, these risks may remain salient. We next turn to an examination of these recently prominent risks and their possible implications for monetary policy strategies.

3. Some principles of optimal monetary policy in response to inflationary shocks

As described in the companion papers by Hajdini and others (2025) and Lipińska and others (2025), widespread shortages, supply-chain disruptions and reductions in labor supply likely contributed to depressed economic activity and increased inflationary pressures during and after the pandemic. This experience has brought renewed attention to the risks and challenges posed by supply shocks and sectoral dynamics, by which we mean both shocks that originate in

statistical forecasting model that does not change across the different policy rules. As shown by appendix figure A.1.5.1, inflation remains well below 2 percent for the first five years after the recession under both rules, implying a significant loss in effectiveness for FAIT.

[20] See Jia and Wu (2023).

specific sectors and the disparate effect of aggregate shocks across sectors. When faced with developments of this kind, monetary policy cannot generally stabilize both output and inflation. As also discussed by Hajdini and others (2025), the post-pandemic rise in inflation was much larger than would have been predicted based on the patterns in the pre-pandemic data. The risk of rapid increases in inflation (inflation surges) further complicates the task of monetary policy.

In the following sub-sections, we discuss each of the risks enumerated above and illustrate the appropriate monetary policy response in stylized models. The discussion is complemented by an appendix describing simulations of models that have been used in the literature. The models and the simulations presented in the appendix are intentionally simplified versions of more complex settings. These simplified settings abstract from uncertainty about the shocks and transmission mechanisms, and thus, this section does not consider implications for risk-management. Even so, the discussion presents general principles likely to be applicable in more complex scenarios.

In addition, we briefly discuss possible implications of these shocks for the costs and benefits of FIT and FAIT. In particular, the frequency and size of the inflationary shocks mentioned above can affect the distribution of aggregate inflation as well as the likelihood and duration of ELB episodes—all relevant to the relative merits of FIT versus FAIT. While the probability of hitting the ELB may remain appreciable going forward, if the distribution of inflationary shocks were such that either the probability of hitting the ELB or its severity is noticeably lower, the expected benefit of FAIT may be smaller than it would have been had the distribution of shocks remained as it was.

3.1 Aggregate supply shocks

Aggregate supply shocks—which move inflation and economic activity in opposite directions—have been especially notable in recent years. For example, the post-pandemic environment featured both significant movements in labor supply, and higher labor productivity growth, on average, than in the previous decade. In this section, we discuss some of the considerations that shape the appropriate policy response to supply shocks. Broadly speaking, provided inflation expectations remain well anchored, the central bank may wish to "look through" a *temporary* deviation of inflation from its longer-run goal rather than adjust policy in ways that reduce that deviation but magnify the adverse short-run effect on output and employment.[21]

The term "supply shocks" is broad and encompasses different types of disturbances. In general, it is useful to distinguish between two classes of aggregate supply shocks:

[21] For example, in the context of a New Keynesian model, Bodenstein, Erceg, and Guerrieri (2008) note that the optimal policy should not be overly aggressive at stabilizing inflation in response to *transitory* energy supply shocks. Similarly, if short-lived supply disruptions may potentially result in permanent damage to the productive capacity of the economy, the optimal policy should focus on stabilizing real activity rather than inflation; see Galí (2022).

- Supply shocks in the first class do not affect the productive capacity of the economy (that is, potential output and its determinants, such as the natural rate of unemployment). These shocks are referred to as "cost-push" shocks.[22] Given the link between inflation and economic activity, offsetting an inflationary cost-push shock comes at the cost of a contraction in aggregate demand and a negative output gap. Consequently, optimal monetary policy might allow some inflation, as long as inflation expectations remain well anchored.[23] When the link between economic activity and inflation is weak, reducing inflation would require large declines in output.[24] In such circumstances, optimal policy may lean against the shock only modestly. In this case, the funds rate may rise a little more than inflation, raising real interest rates slightly to maintain a somewhat contractionary monetary policy stance.[25]
- A second class of supply shocks encompasses those that affect potential output, such as productivity shocks.[26] Following these shocks, the central bank can in principle intervene to align aggregate demand with aggregate supply, thereby largely closing the output gap and stabilizing inflation.[27] In appendix A.2.2, we use a simulation from a stylized macroeconomic model to illustrate the optimal monetary policy response to a positive and persistent increase in the level of aggregate productivity. We show that, for this shock, there is no conflict between stabilizing inflation and the output gap under optimal policy.

3.2 Sectoral dynamics

The previous section described some broad principles for monetary policy following aggregate shocks. However, some shocks may originate at the sectoral level. In addition, some aggregate shocks may particularly affect certain sectors of the economy (for example, through sector-specific capacity constraints). Recent examples include the post-pandemic supply-chain disruptions and the shift in the composition of demand from services to goods described in Hajdini and others (2025). When sectoral dynamics are important, aggregate inflation is not tightly linked to overall aggregate demand (the output gap), but can be affected by other factors,

[22] Examples of cost-push shocks in simple models are somewhat artificial. We discuss below realistic examples of shocks that have cost-push properties in the context of richer models.

[23] The optimal policy response to cost-push shocks has been the subject of an extensive discussion in the literature. Galí (2008) for example, lays out the optimal policy response to cost-push shocks in the canonical three-equation New Keynesian model.

[24] Recent surveys of the evidence on this relation—the Phillips curve—include Furlanetto and Lepetit (2024) and Tetlow (2024).

[25] The optimal policy response described here differs from the case in which the funds rate does not respond to a shock. Not responding may be warranted when the effects of monetary policy are felt well after the effects of a shock have fully dissipated, provided that longer-term inflation expectations remain well anchored.

[26] Shocks to productivity have been a significant feature of this decade and may continue to be important going forward, with advances in artificial intelligence technology as a leading example. Qualitatively similar outcomes would result from other shocks to potential output and its determinants, such as an increase in population stemming from immigration.

[27] In practice, a policymaker uncertain about how much potential output has shifted may be reluctant to induce large movements in activity. In that case, policy might stabilize inflation only to a limited extent.

such as imbalances between supply and demand across sectors. As a result, monetary policy will generally not be able to stabilize both aggregate inflation and the output gap.[28]

In models that focus on sectoral dynamics, a common finding is that a heavy emphasis on stabilizing aggregate inflation can result in undesirable declines in aggregate output and misallocation across sectors, while prioritizing output gap stabilization would be close to optimal.[29] When weighing the degree to which monetary policy should intervene, the literature emphasizes that policymakers should especially seek to lean against inflation when it arises in sectors with stickier prices, as inflation in those sectors can cause persistent movements in aggregate inflation. By contrast, optimal policy should tolerate more inflation arising in sectors with more flexible prices.[30] In appendices A.2.3 and A.2.4, using a stylized macroeconomic model, we illustrate that optimal monetary policy may allow some inflation when sectoral dynamics are consequential, provided that inflation expectations remain well anchored.

The literature also highlights the relevance for optimal monetary policy of the input-output network—the fact that outputs of some sectors are inputs to other sectors. Some shocks may affect the price of certain inputs more than others and may therefore affect the costs of some sectors more than others. The result, again, is that monetary policy cannot fully stabilize both inflation and output. While quantitative results on optimal policy in network models are relatively few and recent, some results suggest that, as in the earlier literature on optimal policy in multi-sector models, a heavy focus on output gap stabilization is close to optimal.[31] Taking supply chains into account, the literature emphasizes the advantages of focusing on mitigating

[28] These issues are not new. For example, the contribution of shifts in the composition of demand to the inflation of the late 1950s was extensively discussed by economists at the time (Schulze, 1959, among others) and oil price shocks were key drivers of inflation in the 1970s. Consequently, there is a rich body of literature that studies monetary policy in response to sectoral shocks.

Interest in sectoral shocks had been growing before the pandemic—see Baqaee and Farhi (2020)—and the pandemic experience has further increased interest in such shocks in the literature; see the discussion in Hajdini and others (2025). Guerrieri and others (2021) study optimal monetary policy following a shock that shifts demand from one sector to another and argues that, when labor is mobile across sectors, accommodative monetary policy, by stimulating wages in the growing sector, helps reallocate activity towards it.

[29] In the context of New Keynesian models, an important reference is Erceg, Henderson and Levin (2006).

[30] See Aoki (2001), who studies optimal monetary policy in a multi-sector model following a shock that affects relative prices across sectors and argues that optimal policy should aim to stabilize inflation in sticky-price sectors. This prescription is akin to a policy that focuses on core inflation as a guide to implement policy (while still targeting headline inflation). See also the detailed discussion of this case, as well as the case of sticky wages, in Erceg, Levin and Henderson (2000) and in Woodford (2003).

[31] A recent example in this literature is La'O and Tahbaz-Salehi (2022). These authors characterize optimal monetary policy in an environment in which price stickiness arises because of information frictions and use simulations of a calibrated model to find that stabilizing a consumer price index leads to quantitatively relevant welfare losses relative to the optimal policy. In addition, Rubbo (2023) uses a multisector model with a realistic input-output network to argue that optimal monetary policy should not target a consumer price index, but rather a "divine coincidence index" that weights sectors according to their share of sales in the input-output network as well as their degree of price stickiness.

fluctuations in the prices of important suppliers. Price movements in sectors that are important suppliers set off a process of adjustment as each sector reacts to changes in its suppliers' prices. Because this process is uncoordinated and likely to involve significant inefficiency, research suggests that monetary policy should especially lean against persistent inflation in sectors that are important to the supply chains of the rest of the economy.[32]

3.3 Nonlinearities

The pandemic experience has highlighted that the relationship between inflationary pressures and the strength of the economy may be weak in normal times but strengthen sharply in certain circumstances.[33] For example, shocks to demand that affect sectors that are at or close to full capacity may trigger a sudden surge in inflation.[34] In addition, elevated inflation may induce a situation in which price-setting becomes highly sensitive to current macroeconomic conditions, perhaps because firms concurrently adjust their prices.[35] Furthermore, a large increase in realized inflation may raise the public's attention to inflation dynamics, causing inflation developments to become very salient for the public.[36] These nonlinear inflation dynamics may lead to a persistent upward spiral between inflation and inflation expectations, if not countered by a strong policy response and a credible commitment to return inflation to target.

In Appendix A.3, we illustrate the optimal policy response in the presence of nonlinearities in several macroeconomic models featuring empirically plausible nonlinear inflation dynamics.[37] The simulations presented there illustrate an important point made in the literature: in the presence of nonlinearities, policymakers should react with more force to large inflationary shocks than would be optimal in normal times, in order to forestall spikes in inflation.[38]

---

[32] Examples of such sectors include electronics and machinery sectors. These sectors have correspondingly substantial weights in the divine coincidence index of Rubbo (2023), cited in footnote 31.

[33] See the discussion in Hajdini and others (2025). Furlanetto and Lepetit (2024) provide a summary of the state of knowledge on Phillips curve nonlinearities.

[34] See for example Comin, Johnson and Jones (2023).

[35] Blanco and others (2024) provide a model in which nonlinear inflation dynamics arise because firms can choose to readjust their prices more frequently following a persistent inflationary shock. Thus, periods of high aggregate inflation would induce firms to adjust prices more frequently, which, in turn, amplifies the initial inflation response to the negative supply shock.

[36] Several recent papers (Bracha and Tang (2024), Pfäuti (2025), Weber and others (2025)) emphasize that the public's attention to inflation increases as inflation increases, which raises the possibility that inflation expectations may become more sensitive to realized inflation when inflation is high.

[37] These simulations are based on the model of Blanco and others (2024). We examine the robustness of our results in a version of the FRB/US model that incorporates nonlinear inflation dynamics. This analysis confirms that policymakers should react more strongly to inflation deviations in the presence of nonlinearities.

[38] Karadi and others (2024) examine a menu cost model in which the sensitivity of inflation to economic activity rises after large shocks due to an endogenous increase in the frequency of price changes and conclude that

3.4 Implications for FIT versus FAIT

The previous five years featured a large role for supply shocks and sectoral dynamics and highlighted the significant risk of inflation surges. The degree to which these events continue to be important has some bearing on the relative benefits of FIT and FAIT.

If the future balance of risks is more like the past five years, with a higher incidence of inflationary shocks, the probability that inflation will be persistently low enough to trigger FAIT may decline, if either the ELB binds less often or, when it does bind, the severity of the episode is reduced. In addition, the possibility of a higher long-run natural rate of interest may also diminish the expected benefits of FAIT.

The simulation results shown in appendix 1.3 suggest that the probability of hitting the ELB remains appreciable as long as the long-run natural rate of interest remains comparatively low. The risks in the years ahead may thus be marked by low inflation and more closely resemble the pre-pandemic period. In this case, FAIT, or similar make-up strategies (for instance, the modified FIT strategy discussed above), can offer benefits compared with FIT in ELB episodes and provide a degree of insurance against such risks.[39] These benefits need to be weighed against the credibility concerns and communication challenges outlined previously.

4. Conclusion: Effective monetary policy in the post-pandemic world

The period between the Great Recession and the COVID-19 pandemic highlighted the risk that inflation might run persistently below target, with attendant risks to the stability of longer-term inflation expectations. The experience since the onset of the pandemic has brought back into focus the risks of persistent supply shocks and sectoral dynamics, as well as sudden inflation surges. These risks may well continue to be relevant in the future.

In this paper, we have explored the costs and benefits of FIT versus FAIT, focusing on outcomes during recessions with a binding ELB, and reviewed some principles of optimal monetary policy in response to supply shocks, sectoral dynamics and the risk of inflation surges.

Regarding the costs and benefits of FIT versus FAIT, in circumstances when the risks associated with the ELB remain significant, alternative strategies that support well-anchored inflation expectations, such as FAIT, might be warranted. FAIT can provide moderate support when inflation is persistently below target. FIT may be able to achieve broadly similar outcomes, if modified with the use of thresholds delaying ELB departure. Importantly, realizing

optimal monetary policy should lean more forcefully against large cost-push shocks than small cost-push shocks. Beaudry, Carter, and Lahiri (2023) find that it is optimal for a central bank to adopt a hawkish anti-inflationary stance in response to shocks if inflation expectations are at risk of de-anchoring. Harding and others (2025) study optimal monetary policy in a model with a nonlinear Phillips curve and find that optimal policy takes a more aggressive approach to curbing inflationary pressures associated with larger shocks.

[39] As shown in tables A.1.3.1 and A.1.3.3, the fraction of time spent at the ELB in FRB/US stochastic simulations ranges from around 15 percent to around 30 percent of all quarters. See Nakata (2017) for a discussion of the effect on risk premiums from the likelihood that policy may be constrained by the ELB in the future.

the benefits under either approach requires that the public perceives the strategies to be credible and understands their implications adequately.

Supply shocks and sectoral dynamics may pose a tradeoff between inflation and maximum employment for monetary policy. The weak link between inflation and economic activity in normal times implies that offsetting inflationary supply shocks could require substantial declines in economic activity. Such declines might be justified if the supply shock persistently depresses potential output. In the case of a shock that affects sectors of the economy differently, monetary policy might tolerate some aggregate inflation, especially if it occurs mostly in flexible-price sectors, but should avoid inflation in sectors with sticky prices, or sectors that have a high importance in the supply chains of other sectors. In the event of large shocks, policymakers should be vigilant about the risk that elevated inflation may cause inflation dynamics to become very sensitive to economic conditions, possibly triggering an inflation surge, and should forcefully intervene to avoid such transitions.

References

Amano, Robert, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner (2020). "Average is Good Enough: Average-Inflation Targeting and the ELB," Working Paper 2020-21. San Francisco: Federal Reserve Bank of San Francisco, June, https://www.frbsf.org/research-and-insights/publications/working-papers/2020/06/average-is-good-enough-average-inflation-targeting-and-the-elb/.

Aoki, Kosuke (2001). "Optimal Monetary Policy Responses to Relative-Price Changes," *Journal of Monetary Economics*, vol. 48 (August), pp. 55–80, https://doi.org/10.1016/S0304-3932(01)00069-1.

Arias, Jonas, Martin Bodenstein, Hess Chung, Thorsten Drautzburg, and Andrea Raffo (2020). "Alternative Strategies: How Do They Work? How Might They Help?" Finance and Economics Discussion Series 2020-068. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2020.068.

Baqaee, David Rezza, and Emmanuel Farhi (2020). "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, vol. 135 (February), pp. 105–63, https://doi.org/10.1093/qje/qjz030.

Bean, Charles R. (1983). "Targeting Nominal Income: An Appraisal," *The Economic Journal*, vol. 93 (December), pp. 806–19, https://doi.org/10.2307/2232747.

Beaudry, Paul, Thomas J. Carter, and Amartya Lahiri (2023). "The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?" NBER Working Paper Series 31741. Cambridge, Mass.: National Bureau of Economic Research, September, https://www.doi.org/10.3386/w31741.

Bernanke, Ben (2003). "Constrained Discretion and Monetary Policy," remarks delivered before the Money Marketeers of New York University, New York, February 3, https://www.federalreserve.gov/boarddocs/Speeches/2003/20030203/default.htm.

——— (2017). "Monetary Policy in a New Era," in Olivier Blanchard and Lawrence H. Summers, eds., *Evolution or Revolution? Rethinking Macroeconomic Policy After the Great* Recession. Cambridge, Mass.: MIT Press, pp. 3–48.

Bernanke, Ben, Michael T. Kiley, and John M. Roberts (2019). "Monetary Policy Strategies for a Low-Rate Environment," *AEA Papers and Proceedings*, vol. 109 (May), pp. 421–26, https://doi.org/10.1257/pandp.20191082.

Bernanke, Ben, Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen (1999). *Inflation Targeting: Lessons from the International Experience*, Princeton University Press, https://doi.org/10.2307/j.ctv301gdr.

Blanco, Andrés, Corina Boar, Callum Jones, and Virgiliu Midrigan (2024). "The Inflation Accelerator," NBER Working Paper Series 32531. Cambridge, Mass.: National Bureau of Economic Research, May (revised June 2024), https://doi.org/10.3386/w32531.

Blanchard, Olivier, and Galí, Jordi (2007). "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, vol. 39 (January), pp. 35–65, https://www.doi.org/10.1111/j.1538-4616.2007.00015.x.

Bodenstein, M., Cristopher J. Erceg, and Luca Guerrieri (2008). "Optimal Monetary Policy with Distinct Core and Headline Inflation Rates," *Journal of Monetary Economics*, vol. 55 (October), pp. S18–33, https://doi.org/10.1016/j.jmoneco.2008.07.010.

Boneva, Lena, Richard Harrison, and Matt Waldron (2018). "Threshold-based Forward Guidance," *Journal of Economic Dynamics and Control*, vol. 90 (May), pp. 138–55, https://doi.org/10.1016/j.jedc.2018.02.003.

Borio, Claudio, and Matthieu Chavaz (2025). "Moving Targets? Inflation Targeting Frameworks, 1990–2025," *BIS Quarterly Review,* March, pp. 31–52, https://www.bis.org/publ/qtrpdf/r_qt2503c.htm.

Bracha, Anat, and Jenny Tang (2024). "Inflation Levels and (In)Attention," *The Review of Economic Studies*, rdae063 (June), https://doi.org/10.1093/restud/rdae063.

Bundick, Brent, Isabel Cairó, and Nicolas Petrosky-Nadeau (2025). "Labor Market Dynamics, Monetary Policy Tradeoffs, and a Shortfalls Approach to Pursuing Maximum Employment," Finance and Economics Discussion Series 2025-068. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.068.

Budianto, Flora, Taisuke Nakata, and Sebastian Schmidt (2023). "Average Inflation Targeting and the Interest Rate Lower Bound," *European Economic Review*, vol. 152 (February), https://doi.org/10.1016/j.euroecorev.2023.104384.

Comin, Diego A., Robert C. Johnson, and Callum Jones (2023). "Supply Chain Constraints and Inflation," NBER Working Paper Series 31179. Cambridge, Mass.: National Bureau of Economic Research, April (revised August 2024), https://doi.org/10.3386/w31179.

——— (2025). "What Was the Fed Thinking? Monetary Policy with Supply Constraints," working paper, Dartmouth College, April.

Chung, Hess, Etienne Gagnon, Taisuke Nakata, Matthias Paustian, Bernd Schlusche, James Trevino, Diego Vilán, and Wei Zhang (2019). "Monetary Policy Options at the Effective Lower Bound: Assessing the Federal Reserve's Current Policy Toolkit," Finance and Economics Discussion Series 2019-003, Washington: Board of Governors of the Federal Reserve System, February, https://doi.org/10.17016/FEDS.2019.003.

Coenen, Gunter, Carlos Montes-Galdón, and Sebastian Schmidt (2021). "Macroeconomic Stabilisation and Monetary Policy Effectiveness in a Low-Interest-Rate Environment," *Journal of Economic Dynamics and Control*, vol. 132 (November), https://doi.org/10.1016/j.jedc.2021.104205.

Coibion, Olivier, Yuriy Gorodnichenko, Edward S. Knotek II, and Raphael Schoenle (2023a). "Average Inflation Targeting and Household Expectations," *Journal of Policy Economy Macroeconomics*, vol. 1 (May), pp. 403–46, https://doi.org/10.1086/722962.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber (2023b). "Forward Guidance and Household Expectations," *Journal of the European Economic Association*, vol. 21 (October), pp. 2131–71, https://doi.org/10.1093/jeea/jvad003.

Coulter, Jarod, Roberto Duncan, and Enrique Martinez-Garcia (2022). "Flexible Average Inflation Targeting: How Much is U.S. Monetary Policy Changing?" Globalization Institute Working Paper 417. Dallas: Federal Reserve Bank of Dallas, July, https://doi.org/10.24149/gwp417.

Eggertsson, Gauti B., and Marc P. Giannoni (2021). "Medium-Term Money Neutrality and the Effective Lower Bound," *Journal of Money, Credit and Banking*, vol. 52 (February), pp. 561–600, https://doi.org/10.1111/jmcb.12762.

Eggertsson, Gauti B., and Michael Woodford (2003). "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, vol. 44, August, pp. 139–233, https://www.brookings.edu/wp-content/uploads/2003/01/2003a_bpea_eggertsson.pdf.

Erceg, Christopher J., Michael Kiley, and Robert Tetlow (2014). "Potential Implications of Alternative Approaches to the Timing and Pace of Tightening," memorandum to the Federal Open Market Committee, Board of Governors of the Federal Reserve System, September 5, https://www.federalreserve.gov/monetarypolicy/files/FOMC20140905memo04.pdf.

Erceg, Christopher J., Zoltan Jakab, and Jesper Lindé (2021). "Monetary Policy Strategies for the European Central Bank," *Journal of Economic Dynamics and Control*, vol. 132 (November), https://doi.org/10.1016/j.jedc.2021.104211.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000). "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, vol. 46 (October), pp. 281–313, https://doi.org/10.1016/S0304-3932(00)00028-3.

——— (2006). "Optimal Monetary Policy with Durable Consumption Goods," *Journal of Monetary Economics*, vol. 53 (October), pp. 1341–59, https://doi.org/10.1016/j.jmoneco.2005.05.005.

Furlanetto, Francesco, and Antoine Lepetit (2024). "The Slope of the Phillips Curve," Finance and Economics Discussion Series 2024-043. Washington: Board of Governors of the Federal Reserve System, June, https://doi.org/10.17016/FEDS.2024.043.

Galí, Jordi (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

——— (2022). "Insider-Outsider Labor Markets, Hysteresis, and Monetary Policy," *Journal of Money, Credit and Banking*, vol. 54 (February), pp. 53–88, https://doi.org/10.1111/jmcb.12898.

Gerke, Rafael, Sebastian Giesen, Daniel Kienzler, Joost Röttger, and Alexander Scheer (2021). "A Comparison of Monetary Policy Rules in an Estimated TANK Model," Technical Paper No. 05/2021. Frankfurt: Deustche Bundesbank.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning (2021). "Monetary Policy in Times of Structural Reallocation," prepared for a symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyo., August.

González-Astudillo, Manuel, and Diego Vilán (2019). "A New Procedure for Generating the Stochastic Simulations in FRB/US," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, March 7, https://doi.org/10.17016/2380-7172.2314.

Hajdini, Ina, Adam Shapiro, A. Lee Smith, and Daniel Villar (2025). "Inflation since the Pandemic: Lessons and Challenges," Finance and Economics Discussion Series 2025-070. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.070.

Harding, Martin, Jesper Lindé, David López-Salido, and Mathias Trabandt (2025). "Optimal Monetary Policy with Large Shocks," working paper, Bank of Canada.

Hebden, James, Edward P. Herbst, Jenny Tang, Giorgio Topa, and Fabian Winkler, (2020). "How Robust Are Makeup Strategies to Key Alternative Assumptions?" Finance and Economics Discussion Series 2020-069. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2020.069.

Hebden, James, and David López-Salido (2018). "From Taylor's Rule to Bernanke's Temporary Price Level Targeting," Finance and Economics Discussion Series 2018-051. Washington: Board of Governors of the Federal Reserve System, July, https://doi.org/10.17016/FEDS.2018.051.

Jia, Chengcheng, and Jing Cynthia Wu (2023). "Average Inflation Targeting, Time Inconsistency and Ambiguous Communication," *Journal of Monetary Economics*, vol. 138 (September), pp. 69–86, https://doi.org/10.1016/j.jmoneco.2023.05.010.

Karadi, Peter, Anton Nakov, Galo Nuño, Ernesto Pasten, and Dominik Thaler (2024). "Strike while the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve," CEPR Discussion Paper No. 19339, August. CEPR Press, Paris & London, https://cepr.org/publications/dp19339.

Kiley, Michael T., and John M. Roberts (2017). "Monetary Policy in a Low Interest Rate World," *Brookings Paper on Economic Activity*, Spring, pp. 317–96, https://www.brookings.edu/wp-content/uploads/2017/08/kileytextsp17bpea.pdf.

Krugman, Paul (1998). "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, vol. 29, February, pp. 137–205, https://www.brookings.edu/wp-content/uploads/1998/06/1998b_bpea_krugman_dominquez_rogoff.pdf.

La'O, Jennifer and Alireza Tahbaz-Salehi (2022). "Optimal Monetary Policy in Production Networks," *Econometrica*, vol. 90 (May), pp. 1295–1336, https://doi.org/10.3982/ECTA18627.

Lipińska, Anna, Enrique Martinez-Garcia, and Felipe Schwartzman (2025). "Pandemic and War Inflation: Lessons from the International Experience," Finance and Economics Discussion Series 2025-071. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.071.

Levin, Andrew, David López-Salido, Edward Nelson, and Tack Yun (2010). "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound," *International Journal of Central Banking*, vol. 6 (March), pp. 143–189, https://www.ijcb.org/journal/ijcb10q1a8.htm.

McClure, Ethan M.L., Vitalia Yaremko, Olivier Coibion, and Yuriy Gorodnichenko (2024). "The Macroeconomic Expectations of U.S. Managers," *Journal of Money, Credit and Banking*, https://doi.org/10.1111/jmcb.13163.

Mertens, Thomas M., and John C. Williams (2019). "Monetary Policy Frameworks and the Effective Lower Bound on Interest Rates," *AEA Papers and Proceedings*, vol. 109 (May), pp. 427–32, https://doi.org/10.1257/pandp.20191083.

Mishkin, Frederic S., and Michael Kiley (2025). "The Evolution of Inflation Targeting from the 1990s to the 2020s: Developments and Challenges", NBER Working Paper Series 33585, Cambridge, Mass.: National Bureau of Economic Research, March, https://doi.org/10.3386/w33585.

Nakata, Taisuke (2017). "Uncertainty at the Zero Lower Bound," *American Economic Journal: Macroeconomics*, vol. 9 (July), pp. 186–221, https://doi.org/10.1257/mac.20140253.

Pfäuti, Oliver (2025). "The Inflation Attention Threshold and Inflation Surges," working paper, University of Texas at Austin, Department of Economics.

Reifschneider, David, and David Wilcox (2019). "Average Inflation Targeting Would Be a Weak Tool for the Fed to Deal with Recession and Chronic Low Inflation," Peterson Institute for International Economics 19-16, November, https://www.piie.com/publications/policy-briefs/average-inflation-targeting-would-be-weak-tool-fed-deal-recession-and.

Reifschneider, David, and John C. Williams (2000). "Three Lessons for Monetary Policy in a Low Inflation Era," *Journal of Money, Credit and Banking*, vol. 32 (November), pp. 936–66, https://doi.org/10.2307/2601151.

Rubbo, Elisa (2023). "Networks, Phillips Curves, and Monetary Policy," *Econometrica*, vol. 91 (August), pp. 1417–55, https://doi.org/10.3982/ECTA18654.

Schulze, Charles L. (1959). "Recent Inflation in the United States," Study Paper 1, Joint Economic Committee, Congress of the United States.

Smets, Frank, and Rafael Wouters (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, vol. 97 (June), pp. 586–606, https://doi.org/10.1257/aer.97.3.586.

Tetlow, Robert (2024). "How Large is the Output Cost of Disinflation," *International Finance,* https://doi.org/10.1111/infi.12451.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce (2025). "Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings," *Econometrica*, vol. 93 (January), 229–64, https://doi.org/10.3982/ECTA22764.

Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Various authors (2021). "The ECB's Price Stability Framework: Past Experience, and Current and Future Challenges," European Central Bank Occasional Paper Series No. 269, September, https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op269~3f2619ac7a.en.pdf.

Appendixes

A.1 Supplementary material on FIT versus FAIT: Average inflation, thresholds rules, and VAR expectations

A.1.1 Rules specifications

FIT: $R_t = 0.85R_{t-1} + 0.15(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t) - 0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$

FAIT: $R_t = 0.85R_{t-1} + 0.15(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t + 8(\bar{\pi}_t - \pi^{LR}) * \mathbb{1}(\bar{\pi}_t - \pi^{LR} < 0)(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t + 8(\bar{\pi}_t - \pi^{LR}) * \mathbb{1}(\bar{\pi}_t - \pi^{LR} < 0)) - 0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$

where $\bar{\pi}_t$ is 32-quarter average inflation and $\mathbb{1}(\phi)$ is the indicator function equaling 1 when ϕ is true and zero otherwise. Relative to the simplified description in the main text, the additional terms reacting the two-quarter change in the unemployment rate are designed to mimic the rapid decline in the funds rate typical in a recession.

In addition to these rules, we also report outcomes under rules that react only to shortfalls in economic activity.

FIT (shortfalls): $R_t = 0.85R_{t-1} + 0.15(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t * \mathbb{1}(ygap_t < 0)\ ) - 0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$

FAIT (shortfalls): $R_t = 0.85R_{t-1} + 0.15(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t * \mathbb{1}(ygap_t < 0)\ + 8(\bar{\pi}_t - \pi^{LR}) * \mathbb{1}(\bar{\pi}_t - \pi^{LR} < 0)(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + ygap_t * \mathbb{1}(ygap_t < 0)\ + 8(\bar{\pi}_t - \pi^{LR}) * \mathbb{1}(\bar{\pi}_t - \pi^{LR} < 0)) - 0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$

Finally, we consider rules that respond to the change in the output gap rather than the level of the gap. Optimized rules typically contain a large weight on the change, compared with the level of the gap, and, importantly for the current setting, can moderate the upward inflation bias as a result of FAIT in conjunction with a shortfalls-based approach to the output gap. Accordingly, for our simulations, these simulations only react to the change in the output gap when the level of the gap is negative.

FIT (change): $R_t = 0.85R_{t-1} + 0.15\big(r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + (ygap_t - ygap_{t-1}) * \mathbb{1}(ygap_t < 0)\big) - 0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$

FAIT (change): $R_t = 0.85R_{t-1} +$

$$0.15\begin{pmatrix} r^{LR} + \pi^{LR} + 2.5(\pi_t - \pi^{LR}) + (gap_t - gap_{t-1}) * \mathbb{1}(gap_t < 0) \\ +8(\bar{\pi}_t - \pi^{LR}) * \mathbb{1}(\bar{\pi}_t - \pi^{LR} < 0) \end{pmatrix}$$

$$-0.85 * (UR_t - UR_{t-2}) * \mathbb{1}(UR_t - UR_{t-2} > 0.25)$$

A.1.2 Simulation methodology

Simulations are conducted using a linearized version of the FRB/US model under the assumption of model-consistent expectations on the part of asset-market participants and wage- and price-setters. The baseline is taken from the public FRB/US dataset, starting in 2030, when the economy is essentially at its long-run steady-state, in which the federal funds rate stands at 3 percent. Shocks are drawn according to the block bootstrap procedure outlined by González-Astudillo and Vilán (2019), using residuals between 1969:Q1 and 2024:Q4, with the exclusion of the exceptionally large residuals associated with the pandemic (2020:Q1-2021:Q2). Simulations run for 20 years and the distributions are taken over 20,000 draws, which are held the same for all the rules simulated.

Long-term inflation expectations (PTR) are endogenous to actual inflation.

$$PTR_t = 0.9\, PTR_{t-1} + 0.05\, \pi_t + 0.05\, \pi^{LR}.$$

A.1.3 Average inflation in FRB/US

As discussed in section 2.1, estimates of the downward bias in average inflation due to the ELB cover a wide range, even across studies using the FRB/US model. Results in Kiley and Roberts (2017) and appendix II of Chung and others (2019) highlight two important sources of divergence in the context of the FRB/US model. In many New Keynesian models, sufficiently large adverse shocks can cause the economy to enter a regime in which activity contracts explosively.[40] In the FRB/US model, this threshold is significantly more likely to be reached under the assumption that all agents in the model form expectations in a model-consistent way. Consequently, the effects of hitting the ELB are, on average, much worse under full model-consistent expectations than under mixed assumptions for expectations formation. Relatedly, previous studies with the FRB/US model have moderated the explosive regime by assuming some degree of active fiscal stabilization. Because the economy deteriorates exponentially in the explosive regime, the details of the assumed fiscal stabilization are important, as slightly weaker fiscal support can translate into large differences in outcomes.

These points are illustrated in table A.1.3.1, which reports features of the distribution of outcomes in stochastic simulations with the FRB/US model, under different assumptions about expectations formation and fiscal stabilization. In these exercises, the degree of fiscal stabilization is modeled by a parameter that determines the responsiveness of the desired surplus-GDP (gross domestic product) ratio target as a function of the output gap. Setting this parameter so that a 1 percent increase in the output gap raises the target surplus-GDP ratio by 1 percent (*gfsrt* = 0.01) is sufficient to avoid hitting the explosive regime, even for very long horizon

[40] See appendix II of Chung and others (2019) and Eggertsson and Giannoni (2021).

stochastic simulations (20 to 30 years). Reducing this parameter to 0.005 is enough to cause a non-negligible fraction of draws to enter into the explosive region within relatively short time horizons (10 years or so).

The main takeaway from the table is that, with a significant degree of active fiscal stabilization, the average inflation bias under FIT is relatively modest, with average inflation around 1.8 percent in the worst case (fully model-consistent expectations). However, without this degree of fiscal stabilization, the inflation bias under FIT can be more severe. Under fully model-consistent expectations (MCE) and reduced fiscal support, average inflation is only 1.5 percent. Further reductions in the degree of fiscal support could produce even lower average inflation rates, depending on how often the explosively deteriorating regime is encountered and how long it is permitted to continue.

Table A.1.3.1: Distributions of outcomes under alternative assumptions about expectations formation and fiscal support

			Inflation			Unemployment Rate			ELB Statistics	
			mean	5%	95%	mean	5%	95%	P(entry ELB)	Fraction of Qtrs at ELB
gfsrt = 0.01										
	MCAP+WP									
		FIT	1.91	0.53	3.17	4.37	2.30	7.58	0.14	0.23
		FAIT	2.07	0.96	3.21	4.19	2.12	7.18	0.14	0.25
		FIT (UR threshold)	1.94	0.59	3.18	4.33	2.31	7.47	0.14	0.25
	MCAP									
		FIT	1.94	0.87	2.82	4.35	2.18	7.64	0.14	0.23
		FAIT	2.04	1.11	2.85	4.14	1.90	7.24	0.14	0.26
		FIT (UR threshold)	1.95	0.90	2.82	4.31	2.18	7.54	0.14	0.25
	MCE									
		FIT	1.83	-0.08	3.28	4.64	1.8	9.76	0.14	0.27
		FAIT	2.07	0.79	3.31	4.29	1.61	8.44	0.14	0.27
gfsrt = 0.005										
	MCE									
		FIT	1.50	-2.14	3.34	5.36	1.22	15.32	0.14	0.28
		FAIT	1.98	0.24	3.39	4.55	0.96	10.61	0.14	0.28

Notes: Authors' estimates.

Table A.1.3.2: Distribution of inflation (π) under FIT and FAIT

		$0.5 \leq \pi$	$0.5 < \pi \leq 1.0$	$1.0 < \pi \leq 1.5$	$1.5 < \pi \leq 2.5$	$2.5 < \pi \leq 3.0$	$3.0 < \pi$
MCAP+WP	FIT	0.05	0.07	0.16	0.50	0.15	0.08
	FIT (UR threshold)	0.04	0.07	0.15	0.51	0.15	0.08
	FAIT	0.02	0.04	0.13	0.57	0.17	0.08
MCE	FIT	0.09	0.07	0.14	0.45	0.15	0.10
	FAIT	0.03	0.04	0.12	0.53	0.18	0.10

Notes: Authors' estimates.

As noted in the main text, adopting a shortfalls approach to the maximum employment objective will imply asymmetrically low probability mass on low inflation outcomes, compared with a symmetric rule. The resulting bias will depend on many of the same factors as the ELB

bias reported in table A.1.3.1. To provide an illustration of the magnitude of the bias, we present results under several different shortfalls rules in table A.1.3.3, under the assumption of model-consistent asset pricing and wage- and price-setting.

Table A.1.3.3: Distribution of outcomes under alternative shortfall policy rules

	Inflation			Unemployment Rate			ELB Statistics	
	mean	5%	95%	mean	5%	95%	P(entry ELB)	Fraction of Qtrs at ELB
MCAP+WP								
FIT (shortfalls)	2.09	0.69	3.43	4.01	1.01	7.58	0.14	0.26
FAIT (shortfalls)	2.21	1.01	3.44	3.89	0.94	7.25	0.14	0.29
FIT (shortfalls,change)	1.97	0.56	3.33	4.18	0.91	8.03	0.15	0.15
FAIT (shortfalls, change)	2.13	0.98	3.35	3.99	0.86	7.60	0.15	0.21

Notes: Authors' estimates; model-consistent wage and price setting; baseline assumptions regarding fiscal support.

A.1.4 Threshold rules

This section presents results on the performance of FIT with threshold conditions. We consider two thresholds: first, a threshold that delays lift-off until either the unemployment rate has reached its natural rate ("UR thresh") or four-quarter PCE price inflation is above 2¼ percent and, second, a threshold that delays lift-off until the four-quarter change in core PCE prices has reached 2 percent ("infl. thresh"). Figure A.1.4.1 presents median outcomes under the same recessionary conditions as described in the main text.

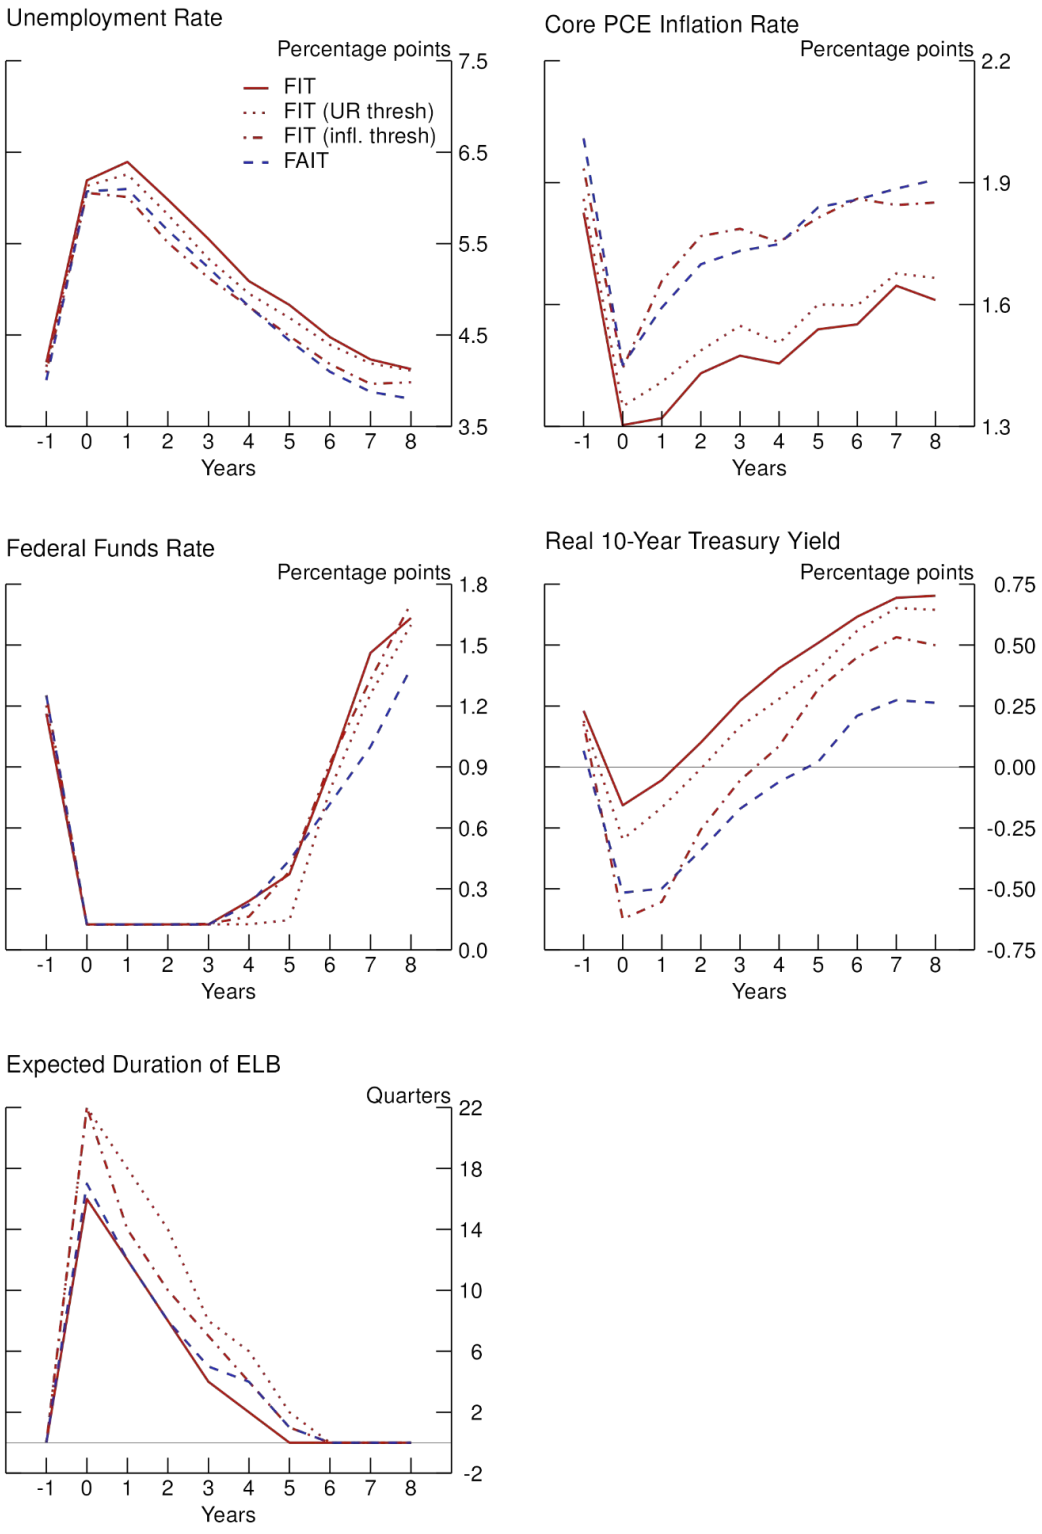
Outcomes under the unemployment rate threshold are very similar to those under the Taylor rule, largely because the expected duration of the ELB under the unemployment threshold is only about six quarters more than under FIT. Perhaps surprisingly, the expected duration of the ELB episode under FAIT is only a little bit longer than under FIT; most of the additional accommodation implied by FAIT comes from the lower path of the funds rate *after* exiting the ELB.

Outcomes under the inflation threshold are quite similar to FAIT, but come with a significant caveat. These simulations are conducted assuming that the model-consistent expectations can be modeled as modal forecasts—that is, in the absence of shocks (a so-called perfect foresight assumption). As is apparent, however, modeling expectations in this way leads to an "expected" duration of the ELB that is systematically larger than the actual outcomes of the stochastic simulations—a violation of the rational expectations hypothesis.[41] In the case of FIT,

[41] Similar considerations also explain why the median level of the funds rate rises above the ELB within four years after hitting it, at a time when the median inflation is still well below two percent. These outcomes are compatible because temporary inflation shocks can cause the ELB to cease binding, even though the inflation rate then returns to its previous (low) level. By year four, most recession draws have experienced at least one of these temporary shocks, as a result of which the ELB has ceased to bind, while inflation remains persistently below two percent.

FIT (UR thresh) and FAIT, the discrepancy is not especially consequential, but the discrepancy is large in the case of the inflation threshold (an expected duration of around five and a half years versus an actual duration of around three years). The large discrepancy in this case is a result of two factors: a long modal duration, which allows the effects of stochastic shocks to cumulate, and the presence of large inflationary shocks in the shock distribution, which implies a significant probability that shocks will cause the threshold to be breached. In light of this discrepancy, it is likely that obtaining a quantitatively reliable evaluation of the inflation threshold case will require a more sophisticated solution methodology, as in Boneva, Harrison, and Waldron (2018).

Figure A.1.4.1: Median outcomes under alternative threshold policy rules at the ELB in a recession



Notes: Authors' estimates.

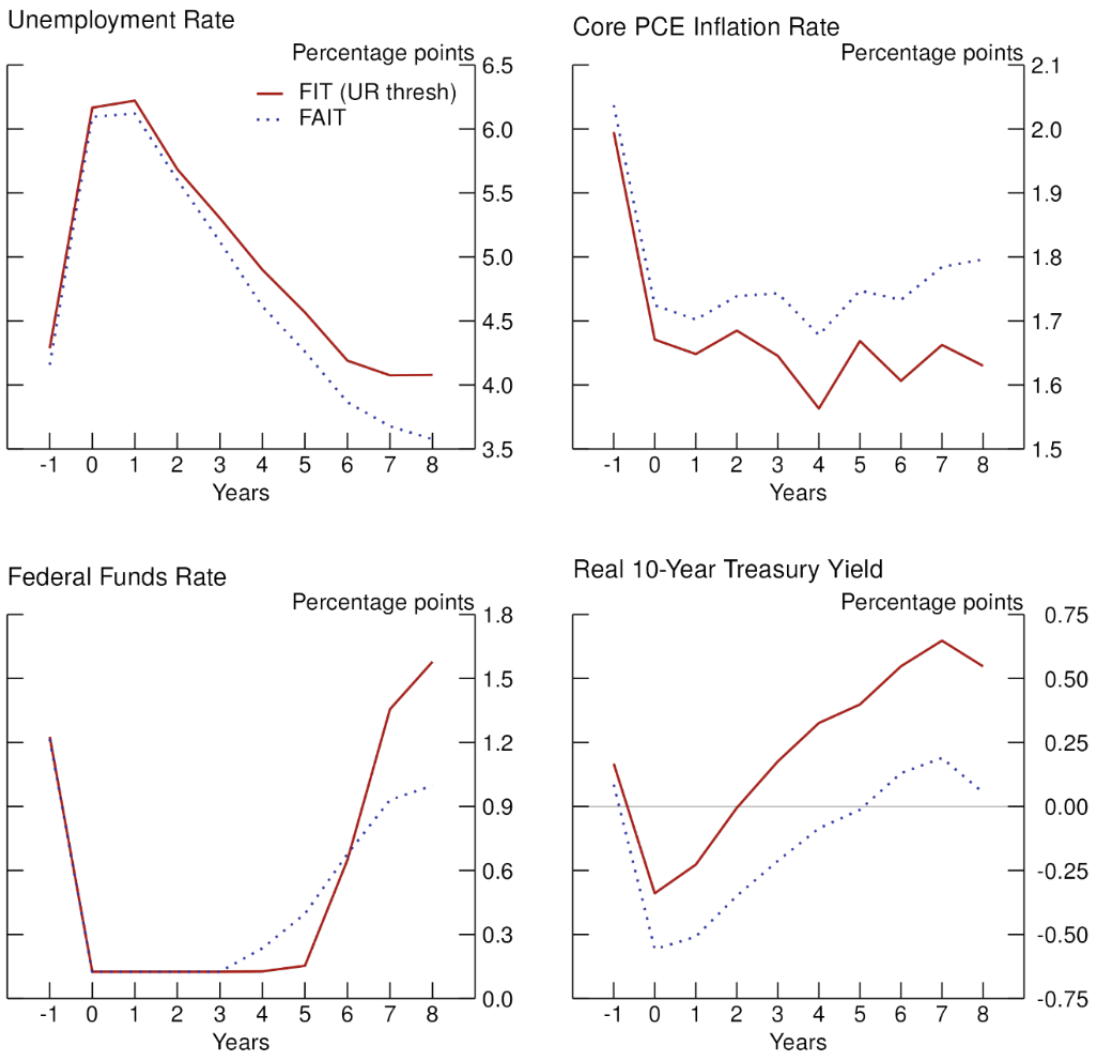
A.1.5 VAR expectations

Figure A.1.5.1 shows median outcomes in a recession under both FIT and FAIT, under the assumption that, while asset market participants form model-consistent (MCAP) expectations, all other agents forecast using simple time-series models.[42] The response of the model to shocks is quite different under MCAP than it is under the version of the model used in the main text, which assumes model-consistent expectations for both asset market participants and wage- and price-setters.[43] For this reason, the distribution of outcomes used for figure A.1.5.1 is the set of draws that result in a recession and a binding ELB under FIT, assuming MCAP expectations–a different set of draws than used in other comparable figures elsewhere in this paper.

[42] These assumptions might be appropriate for an economy in which financial market participants understand the implications of the monetary policy strategy for current and future asset prices, but households and firms, while forward-looking, form expectations on the basis of general historical regularities and are not very attentive to or knowledgeable about the strategy and its implications for future economic performance. Evidence pointing in this direction includes Coibion and others (2023a), who find that the 2020 revision to the Statement had little effect on household expectations a year after its introduction, Coibion and others (2023b), who find experimentally that forward guidance about the future policy rate path has only small effects on expected real interest rates and McClure and others (2025), who find that inflation expectations of managers behave similarly to those of survey respondents in general.

[43] See Hebden and others (2020) for a more complete discussion of the role of expectations for outcomes under FAIT.

Figure A.1.5.1: Median outcomes under alternative policy rules at the ELB in a recession (MCAP expectations)



Notes: Authors' estimates.

A.2 Sectoral shocks and supply constraints

A.2.1 Model

In this section, we present optimal policy results from a model with sectoral shocks and supply constraints. The model is a simplified version of the one used in Comin, Johnson and Jones (2023) and described in Comin, Johnson and Jones (2025). The model features two sectors—goods and services—whose prices are sticky. The occasionally binding constraint is on goods production.

A.2.2 Impulse response to an aggregate productivity shock

We use a simulation from the stylized macroeconomic model to illustrate the optimal monetary policy response to a positive and persistent increase in the level of aggregate productivity.[44] We assume in this subsection that both sectors are the same. In this case, the model aggregates to the standard and stylized three-equation New Keynesian model described in, for example, Galí (2008). In the next subsection, we will adjust the model and introduce an occasionally binding sectoral capacity constraint on goods production.

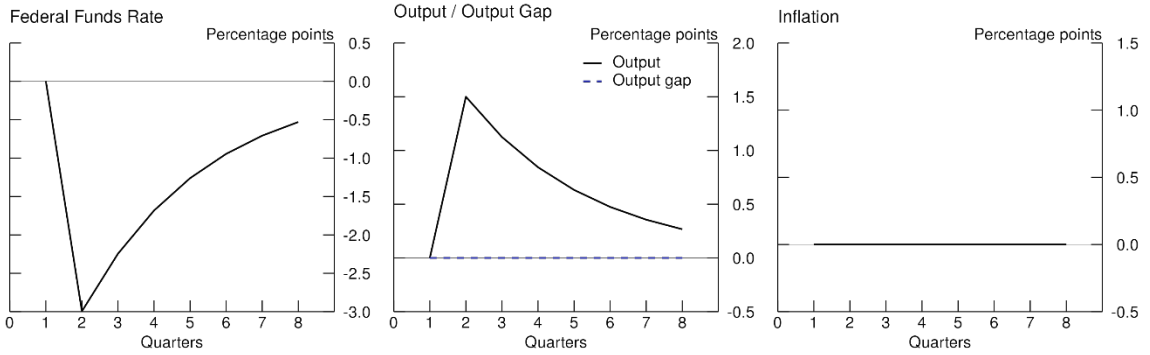
We plot the response of the federal funds rate, output and the output gap, and inflation in figure A.2.2.1 following this shock. All else being equal, in response to a positive productivity shock, output would increase and inflation would drop. Under optimal monetary policy, the federal funds rate (left panel) declines to support aggregate demand, and actual output rises to its new potential level. The interest rate declines under optimal policy because the shock is temporary and households would like to save some portion of the additional income, raising the aggregate savings curve.[45] The output gap closes (middle panel). With the output gap closed, inflation remains at target (right panel). For this type of shock, under optimal policy, there is no conflict between stabilizing inflation and the output gap.[46]

[44] In all simulations discussed in this section, inflation expectations are well anchored at the target, and optimal monetary policy is computed under discretion and weighs aggregate inflation and output gap deviations.

[45] If the shock had instead permanently raised the level of productivity—and, over time, the level of output—the optimal response would have been to raise the interest rate.

[46] The modern literature on New Keynesian models refers to this situation as "divine coincidence." See Blanchard and Galí (2007) and Galí (2008).

Figure A.2.2.1: Impulse responses: Productivity shock under optimal policy



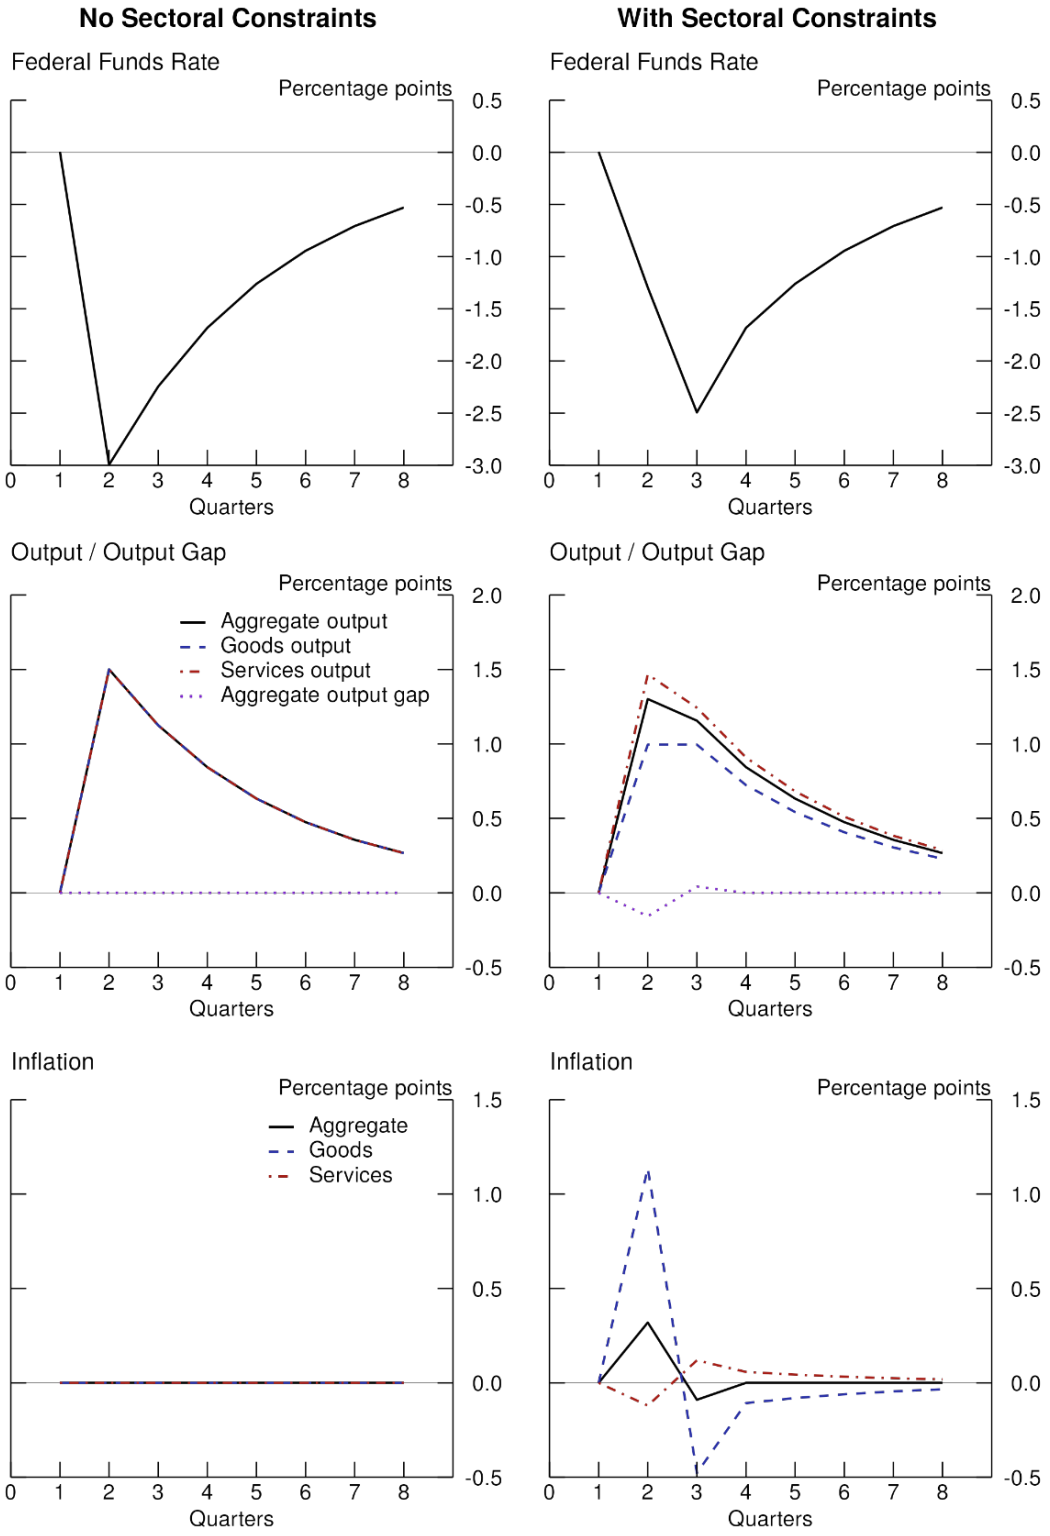
Notes: Authors' estimates. The figure plots impulse responses of model variables relative to their steady states.

A.2.3 Impulse response to an aggregate productivity shock with occasionally binding sectoral capacity constraints

We illustrate how sectoral dynamics may change prescriptions for optimal monetary policy with a scenario in which an aggregate productivity shock causes the capacity constraint on the production of one of the sectors to bind. In contrast to the previous section, the aggregate productivity shock in this case affects the production and productivity of the sectors differently. A heterogenous response of output and prices across sectors prevents the central bank from stabilizing inflation and closing the aggregate output gap.

Figure A.2.3.1 compares the optimal response presented in the previous subsection (the left column) with the optimal policy when the aggregate shock has different consequences across sectors because of the presence of sector-specific capacity constraints (the right panel). In particular, in the scenario shown in the right-hand panels, the aggregate productivity shock causes output in the goods sector to run up against its capacity constraint. For simplicity, in the simulation, we assume that the capacity constraint binds for two quarters after the shock. When goods production hits the constraint (dashed blue line in the middle right panel), that sector is unable to immediately benefit from higher aggregate productivity and the supply of goods cannot expand to meet any additional demand. The price of goods must then rise to balance supply and demand for goods (bottom right panel). Under optimal monetary policy, the funds rate declines, just as in the previous example (top right panel). However, if policy were to ease as much as in the previous example, given the increase in prices in the goods sector, inflation would be much higher. Accordingly, policy does not ease as much as it would in the unconstrained case and the output gap becomes slightly negative. Output and prices in the services sector fall as well. The reduction in services prices is not enough to compensate for the increase in goods prices and, all told, aggregate inflation rises slightly.

Figure A.2.3.1: Impulse response: Aggregate productivity shock under optimal policy with and without sectoral constraints



Notes: Authors' estimates. The figure plots impulse responses of model variables relative to their steady states.

A.2.4 Impulse response to a sectoral capacity shock

Here, we consider outcomes in response to a negative shock to the capacity of goods production that lasts for one period. We compare two cases, one in which monetary policy suboptimally fully stabilizes aggregate inflation, and a second in which policy is set optimally. The key result of this exercise is that, if a supply shock originates in one sector of the economy, fully stabilizing aggregate inflation may entail unwelcome volatility of relative prices and output, misallocation of resources, and a costly decline in aggregate output.

The impulse responses shown in the left column panels of figure A.2.4.1 illustrate outcomes where the funds rate is suboptimally set to keep inflation at target ("Inflation Stabilization").[47] Goods production declines, as shown in the middle left panel in dashed blue. When the capacity constraint in the goods sector binds, goods output falls to the constrained level and goods inflation spikes (bottom left panel). In order for aggregate inflation to be stable (the sole objective of the central bank in this simulation), services inflation must fall to offset higher goods prices, so monetary policy tightens notably to induce a contraction in services output (top left panel). As a result, aggregate output falls, with production in both sectors declining. Once the shock abates, aggregate output returns to its baseline value, while monetary policy returns to a neutral stance, although sectoral dynamics persist. Specifically, goods output continues to undershoot its long-run level and services output overshoots, while inflation in the goods sector turns negative and inflation in services is positive.[48]

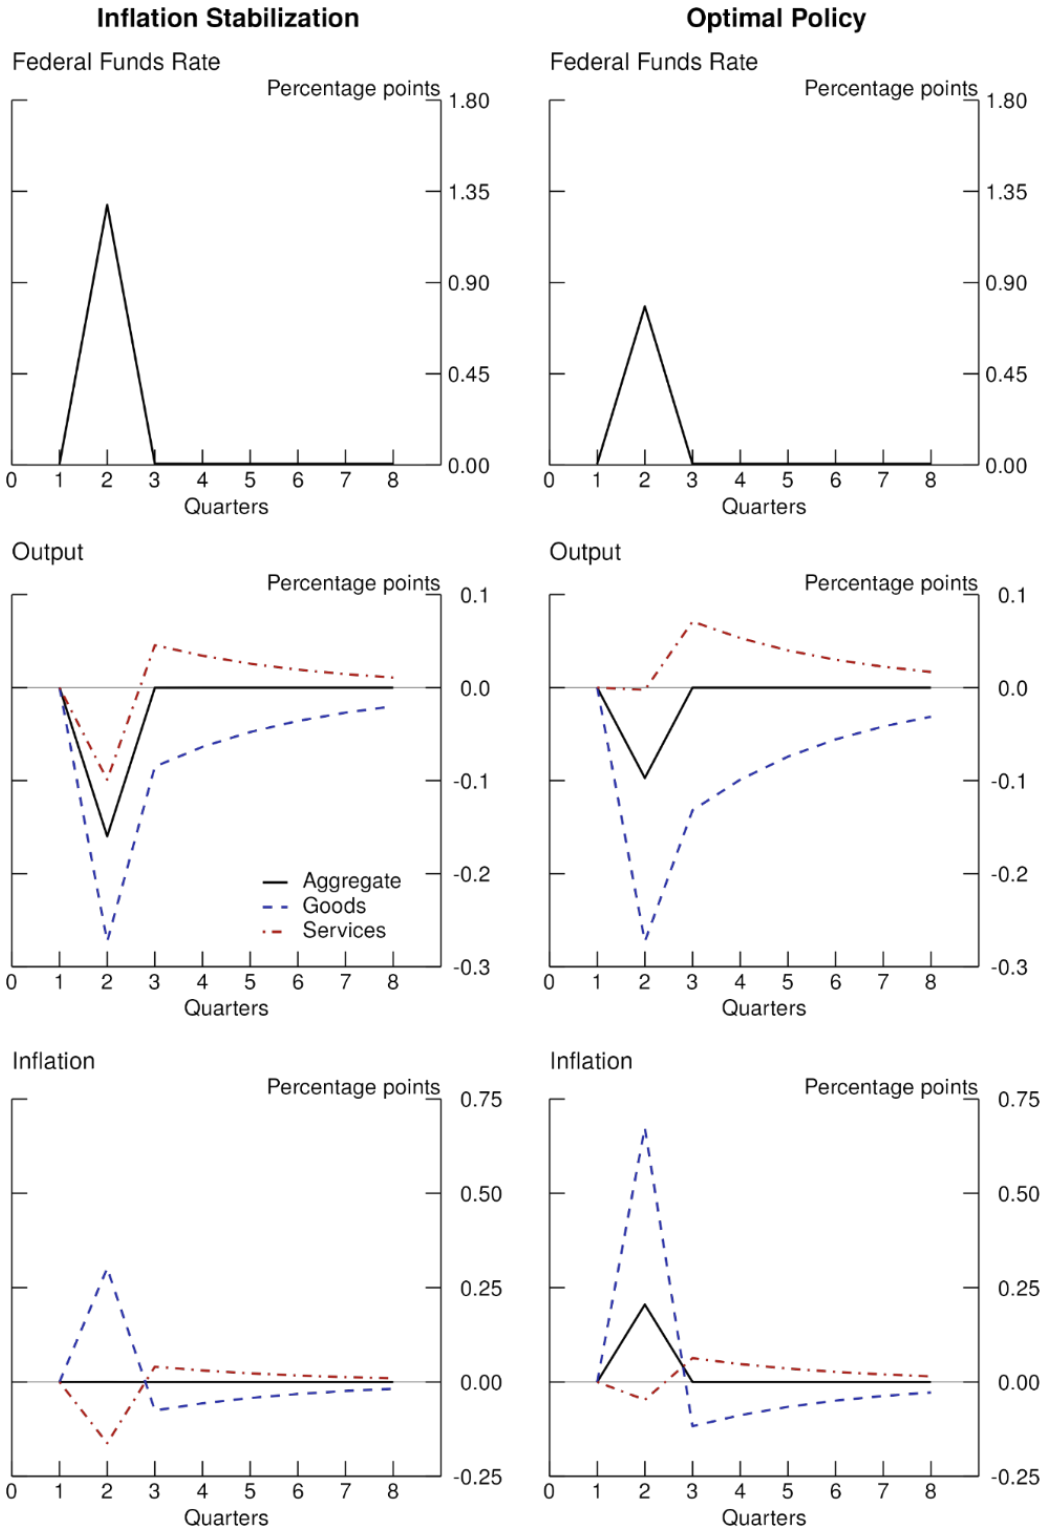
The impulse responses shown in the right column panels of figure A.2.4.1 present outcomes following the goods sectoral supply shock when policy is set optimally. In contrast to the aggregate inflation stabilization policy, optimal policy tolerates some aggregate inflation in the period of the shock (quarter 2), thereby mitigating the decline in aggregate output.[49] A comparison of outcomes under these two policies illustrates the inflation-output tradeoff monetary policy may face following a sectoral supply shock.

[47] In figure A.2.4.1, we plot aggregate variables using solid black lines, and sectoral variables using dashed blue (for goods) and dot-dashed red (for services).

[48] This exercise is only an illustration, as the policy lags in this stylized model are very short, and in more complex models aggregate inflation dynamics might be more persistent.

[49] In addition, in this simulation, goods inflation is higher under optimal policy compared with the full inflation stabilization policy. In part, this reflects how, under optimal policy, aggregate output does not fall as much as it does under the inflation stabilization policy, so that the demand for goods relative to services is also higher.

Figure A.2.4.1: Impulse response: Sectoral production shock under optimal and inflation stabilization policies



Notes: Authors' estimates. The figure plots impulse responses of model variables relative to their steady states.

A.3 Nonlinearities in the slope of the Phillips curve

A.3.1 Model

In this section, we use the model with nonlinearities from Blanco and others (2024) to compute the impulse response below. Details of the equations of the model can be found in Blanco and others (2024).

In computing the optimal policy, we maximize household preferences, given below, subject to the model's equilibrium constraints:

$$\sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\sigma}}{1-\sigma} - \frac{l_t^{1+\gamma}}{1+\gamma} \right),$$

where β is the household discount factor, c_t is household consumption, l_t is household labor supply, and σ, γ are parameters.

In simulations in which monetary policy is following a Taylor rule, we use a version of the inertial Taylor rule: $R_t = 0.81R_{t-1} + 0.19(R^{LR} + \pi^{LR} + 2.04(\pi_t - \pi^{LR}) + 0.08ygap_t)$, where the values for the Taylor rule parameters are taken from Smets and Wouters (2007).

A.3.2 Impulse response to supply shock with nonlinearities

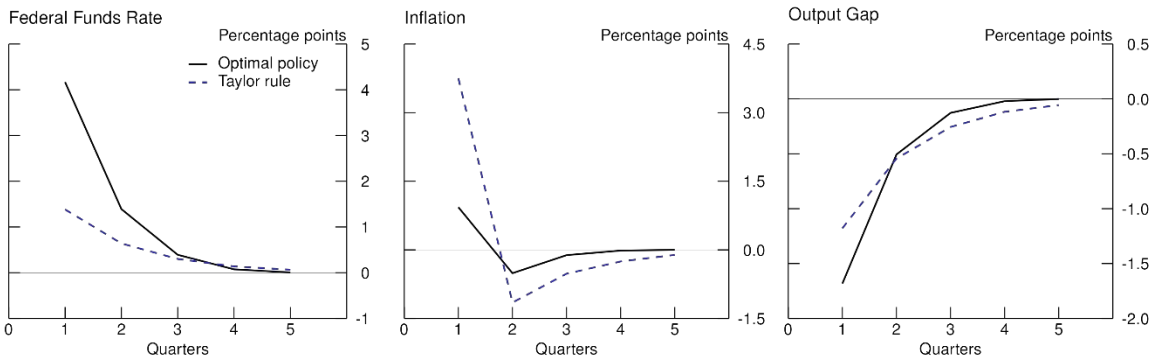
In figure A.3.2.1, we illustrate the optimal monetary response to supply shocks in the presence of nonlinear inflation dynamics.[50] Panel A compares outcomes under optimal policy (the solid black line) with those arising under a Taylor rule (the dashed blue line) that is intended to describe the behavior of policymakers in normal times. As can be seen in the middle panel, reacting to shocks as in normal times may lead to large deviations of inflation from target in the presence of nonlinearities. Instead, optimal monetary policy reacts expeditiously and forcefully, keeping inflation under control, at the cost of a somewhat weaker output gap.[51]

[50] We consider a cost-push shock that temporarily raises the cost of labor. Expectations of all agents in the model are model-consistent and long-run inflation expectations are anchored at the target rate.

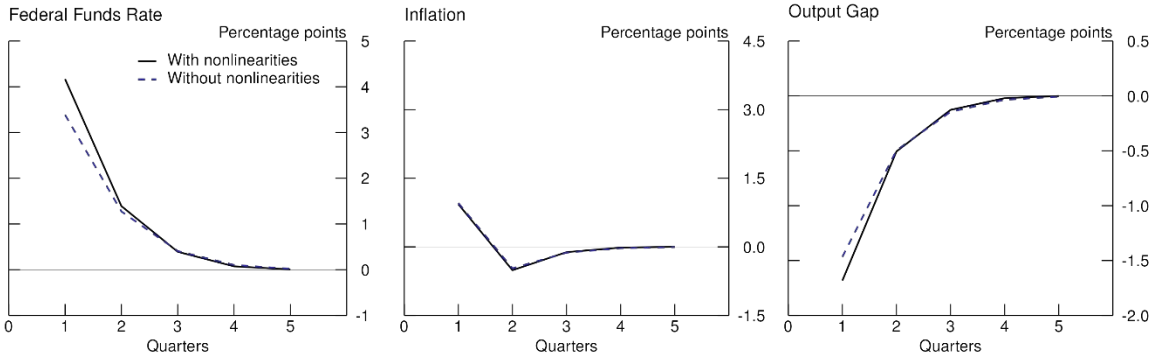
[51] The example considered here illustrates the optimal policy starting from at-target inflation. If, instead, inflation is already notably above target, the nonlinearities may imply that the output loss associated with a reduction in inflation is lower than in normal times, assuming that inflation expectations remain well anchored.

Figure A.3.2.1: Impulse responses in the Three-Equation New Keynesian Model with nonlinearities

Panel A: Optimal policy and Taylor rule policy in the case of a supply shock, model with nonlinearities



Panel B: Optimal policy in the case of a supply shock, models with nonlinearities and without nonlinearities



Notes: Authors' estimates. The figure plots impulse responses of model variables relative to their steady states.

To complement these results, panel B conducts an alternative exercise: It contrasts outcomes under optimal policy in the baseline model with nonlinearities (the same solid black line as in panel A) with those in a counterfactual model where nonlinearities are not present (the dashed blue lines). The simulations illustrate that, in the presence of nonlinearities, the optimal prescription is for policymakers to react with more force to increases in inflation.[52]

[52] Optimal policy calls for a persistently negative output gap, and inflation that falls below baseline after the first period, because such a path lowers the incentive to raise prices initially, helping to limit the spike in inflation.

A.3.3 Simulations of a version of the FRB/US Model with nonlinearities in the slope of the Phillips curve

In this section, we present simulations of a version of the FRB/US model that incorporates nonlinear inflation dynamics. More specifically, the linearized price Phillips curve of the model is replaced by the system of nonlinear equations that describes price setting in the model of Blanco and others (2024). The model is simulated under the assumption of model-consistent expectations on the part of asset-market participants and wage- and price-setters.

We perform the same experiments as in appendix A.3.2. For this purpose, the supply shock under consideration is a 1 percent negative shock to the level of total factor productivity that unwinds gradually over five years. The optimal policy chooses the path of the federal funds rate that minimizes a loss function that penalizes movements in the inflation gap, the unemployment gap, and changes in the policy rate. Different weights are associated with these three objectives. We consider a weighting structure that assigns a lower penalty to movements in the unemployment gap and changes in the policy rate (low inertia).

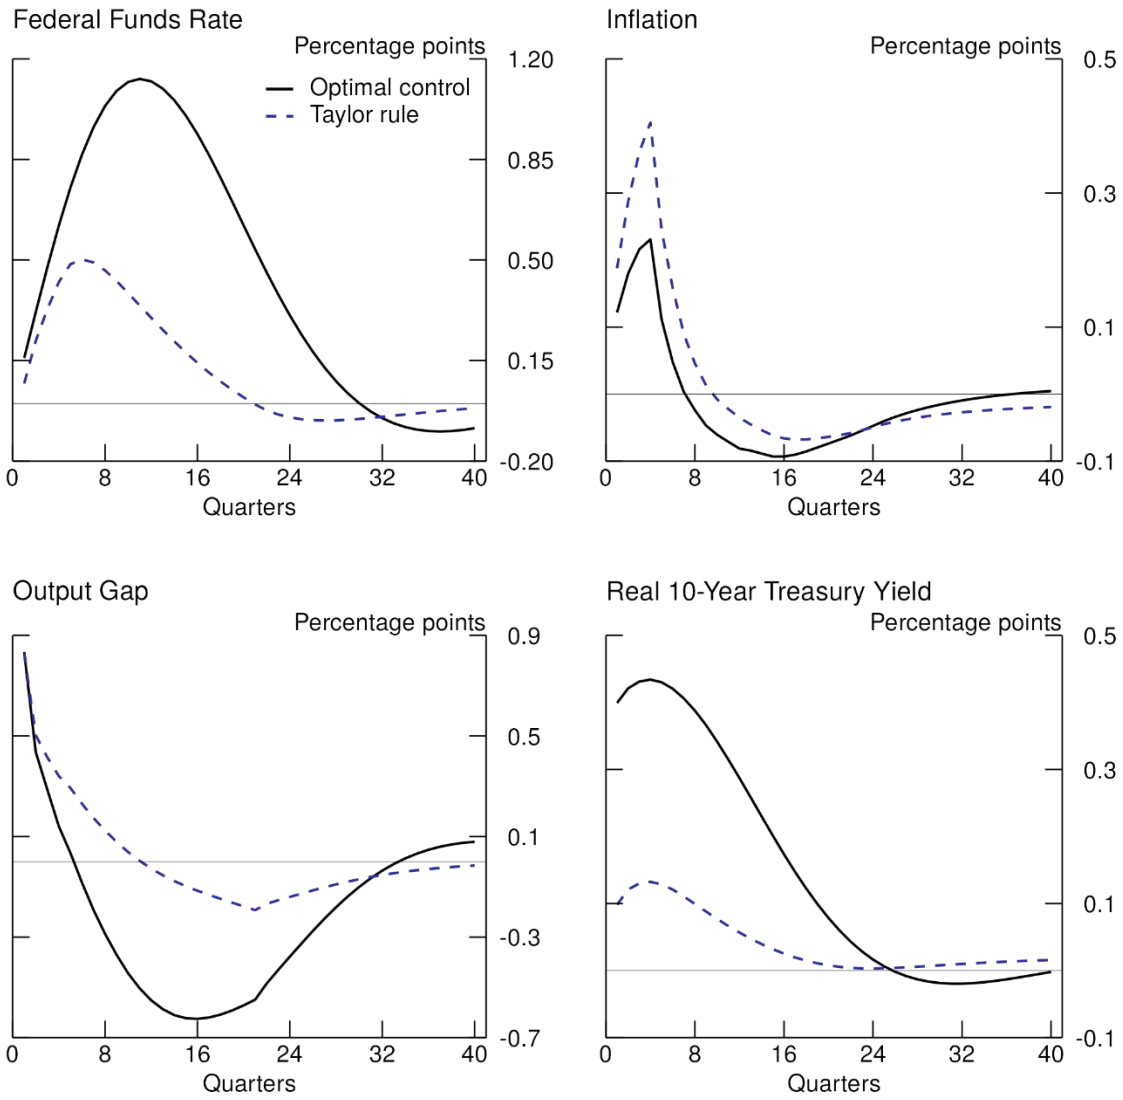
Panel A of figure A.3.3.1 compares outcomes under the optimal policy (the solid black lines) with those arising when monetary policy is conducted according to a standard Taylor rule (the dashed blue lines).[53] As can be seen from the responses of the federal funds rate and the 10-year real rate and as was the case in the model of the previous subsection, the optimal policy reacts more expeditiously and forcefully than the Taylor rule, which helps contain the rise in inflation induced by the shock.

Panel B of figure A.3.3.1 contrasts outcomes under optimal policy in the model with nonlinearities (the same black lines than as in panel A) with those in a counterfactual model where nonlinearities are shut down (the dashed blue lines). Once again, we see that policymakers react more forcefully to increases in inflation in the presence of nonlinearities. Unlike in the example of the previous subsection, this more forceful reaction of monetary policy is initially not sufficient to fully offset the additional inflation brought by nonlinearities. However, from period 6 onward, inflation is lower in the case with nonlinearities on account of both a lower output gap and a steeper Phillips curve (the sacrifice ratio is lower with nonlinearities).

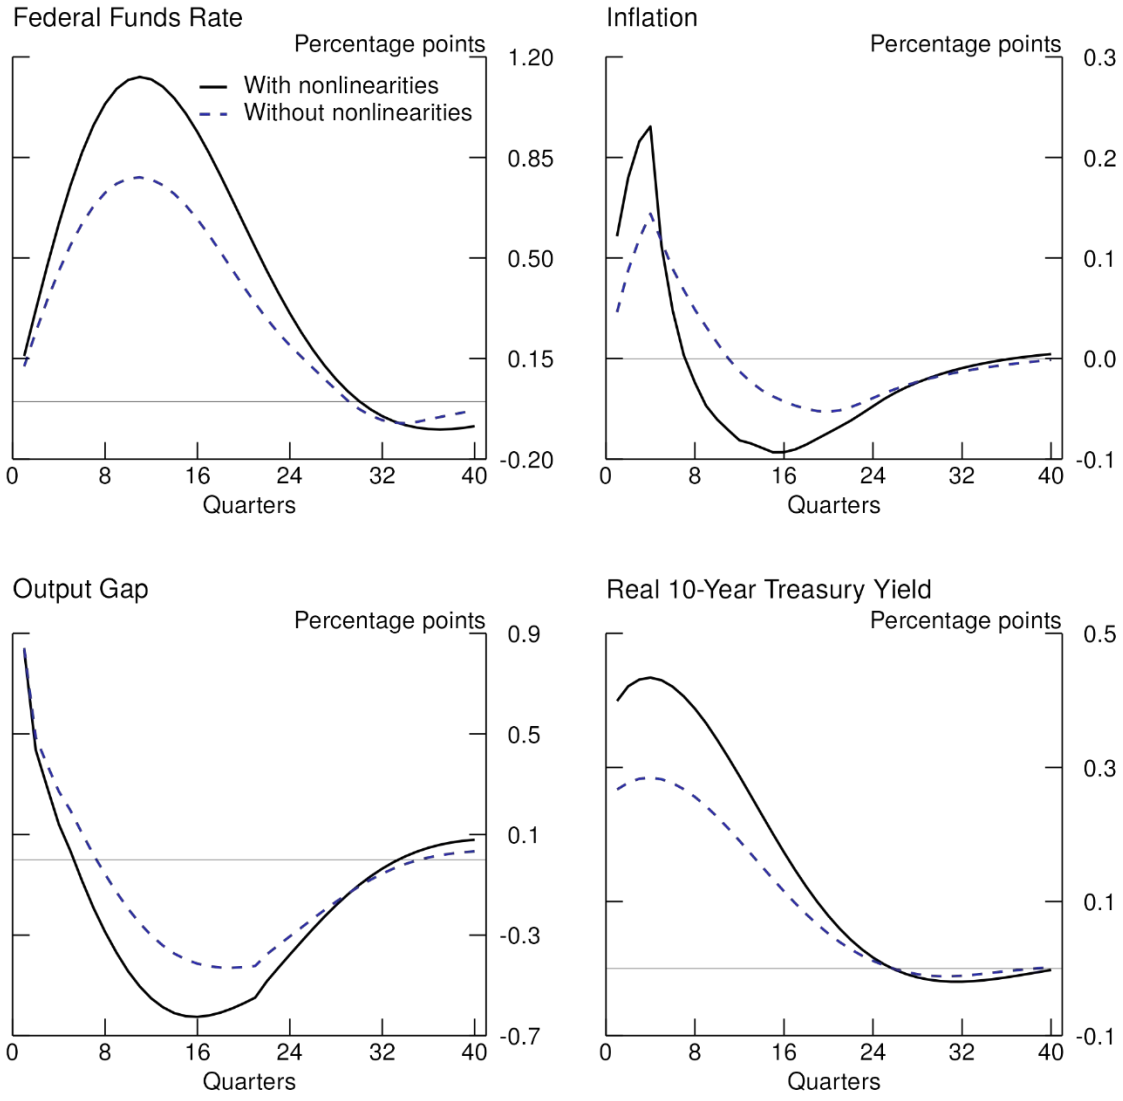
[53] The equation for the Taylor rule is the same as that under FIT in appendix A.1 above with $\phi_i = 0.81$, $\phi_\pi = 1.5$, and $\phi_y = 1$.

Figure A.3.3.1: Impulse responses to a supply shock in the FRB/US model

Panel A: Optimal policy and Taylor rule, model with nonlinearities



Panel B: Optimal policy, model with nonlinearities and without nonlinearities



Notes: Authors' estimates.

Journal of Economic Perspectives—Volume 11, Number 2—Spring 1997—Pages 97–116

Inflation Targeting: A New Framework for Monetary Policy?

Ben S. Bernanke and Frederic S. Mishkin

The world's central bankers and their staffs meet regularly, in venues from Basle to Washington, to share ideas and discuss common problems. Perhaps these frequent meetings help explain why changes in the tactics and strategy of monetary policymaking—such as the adoption of money growth targets in the 1970s, the intensification of efforts to reduce inflation in the 1980s, and the recent push for increased institutional independence for central banks—tend to occur in many countries more or less simultaneously. Whatever their source, major changes in the theory and practice of central banking are of great importance, for both individual countries and the international economy. In this article, we discuss a new strategy for monetary policy known as "inflation targeting," which has sparked much interest and debate among central bankers and monetary economists in recent years. This approach is characterized, as the name suggests, by the announcement of official target ranges for the inflation rate at one or more horizons, and by explicit acknowledgment that low and stable inflation is the overriding goal of monetary policy. Other important features of inflation targeting include increased communication with the public about the plans and objectives of the monetary policymakers, and, in many cases, increased accountability of the central bank for attaining those objectives.

■ *Ben S. Bernanke is Class of 1926 Professor of Economics and Public Affairs, Princeton University, Princeton, New Jersey, and Research Associate of the National Bureau of Economic Research, Cambridge, Massachusetts. Frederic S. Mishkin is Executive Vice President and Director of Research at the Federal Reserve Bank of New York, New York, New York, and Research Associate of the National Bureau of Economic Research, Cambridge, Massachusetts. He is on leave from his position as the A. Barton Hepburn Professor of Economics and Finance, Graduate School of Business, Columbia University, New York, New York.*

Inflation targeting in various forms has been adopted in recent years by a number of industrialized countries, including Canada, the United Kingdom, New Zealand, Sweden, Australia, Finland, Spain and Israel.[1] Table 1 offers some details about the specific plans in each country. There are also important elements of inflation targeting, as we discuss below, in the long-standing and well-regarded monetary policy approaches of Germany and Switzerland. In the United States, inflation targeting has been advocated by some influential policymakers, and Senator Connie Mack (R-Fla.) has introduced a bill that, if passed, would establish price stability as the primary goal of monetary policy [S.R. 1266, 104th Cong. 1st sess.]. Finally, the Maastricht treaty mandates price stability as the primary objective of the European Central Bank, and it seems likely—if European monetary union in fact occurs—that the ECB would incorporate major elements of the inflation targeting approach in its procedures (Issing, 1996).

We begin our discussion of inflation targeting with some details of how this approach has been implemented in practice. We focus on the practice of inflation targeting, rather than the theory, because we believe that the rhetoric associated with inflation targeting is often misleading. In particular, we will argue that actual experience with this approach shows that inflation targeting does not represent an ironclad policy *rule*, as some writers on the subject and even some advocates of this approach seem to assume. Instead, inflation targeting is better understood as a policy *framework*, whose major advantage is increased transparency and coherence of policy, and in which fairly flexible, even "discretionary" monetary policy actions can be accommodated.[2] We next discuss in more detail why viewing inflation targeting as a framework, rather than a rule, blunts some of the arguments that have been made against it and in general enhances the appeal of this approach. This is not to say that valid questions do not remain about this strategy for monetary policy; in the final portion of the paper we discuss some important additional issues and draw conclusions about the usefulness of the inflation targeting framework.

Inflation Targeting in Practice

Although every country that has adopted inflation targeting has customized the approach in various ways, certain empirical generalizations about this strategy can be made.

The hallmark of inflation targeting is the announcement by the government, the central bank, or some combination of the two that in the future the central bank will strive to hold inflation at or near some numerically specified level. As can be seen in Table 1, inflation targets are more often than not specified as ranges—

[1] Detailed analyses of experiences with inflation targeting can be found in Goodhart and Vinals (1994), Leiderman and Svensson (1995), Haldane (1995) and McCallum (1996), among others.

[2] King (1996) adopts a similar view.

Table 1
Operational Aspects of Inflation Targets

Country (date of adoption)	*Target Series Definition*	*Target Level (percentage annual inflation)*	*Time Horizon*
Australia (1993)	Underlying CPI (excluding fruit and vegetables, petrol, interest costs, public sector prices and other volatile prices)	2–3	Ongoing
Canada (February 1991)	Core CPI (excluding food, energy and first-round effects of indirect taxes)	1–3	18 months
Finland (February 1993)	Underlying CPI (excluding government subsidies, indirect taxes, housing prices and mortgage interest payments)	about 2	Ongoing
Israel (December 1991)	CPI	8–11	1 year
New Zealand (March 1990)	Underlying CPI (excluding changes in indirect taxes or government changes, significant changes in import or export prices, interest costs and natural disasters)	0–2 (until November 1996; 0–3 thereafter)	1 year
Spain (January 1995)	CPI (excluding first-round effects of indirect tax changes)	below 3	Through 1997
Sweden (January 1993)	CPI	2 ± 1	Ongoing
United Kingdom (October 1992)	RPIX (RPI excluding mortgage interest payments)	lower half of 1–4 until spring 1997; 2.5 or less thereafter	Until the end of this Parliament

for example, 1–3 percent—rather than single numbers, and they are typically established for multiple horizons ranging from one to four years. However, there are exceptions to both observations; indeed, Germany, with the longest experience with inflation-focused monetary policy, specifies its implicit inflation target as a point and only for a one-year horizon. Initial announcements of inflation targeting generally allow for a gradual transition from the current level of inflation to a desired steady-state level, usually the level deemed consistent with price stability. "Price stability" never in practice means literally zero inflation, however, but usually something closer to a 2 percent annual rate of price change, for reasons we discuss later.

There is a lively debate over whether targeting should be of the inflation rate per se or of the price level. Of course, a targeted price level need not remain constant indefinitely, but could be allowed to drift upward in a predetermined way over time (Goodhart and Vinals, 1994; Svensson, 1996). The relative disadvantage of targeting the inflation rate is that unanticipated shocks to the price level may be

treated as bygones and never offset; as a result, forecasts of the price level at long horizons might have a large variance under inflation targeting, which presumably impedes private-sector planning.[3] On the other hand, strict price-level targeting requires that overshoots or undershoots of the target be fully made up, which reduces the variance of long-run forecasts of prices but could impart significantly more volatility into monetary policy in the short run.[4] In practice, central banks tend to compensate partially for target misses, particularly at shorter horizons.

Associated with the announcement of inflation targets there is usually some statement to the effect that control of inflation is the "primary" or "overriding" goal of monetary policy and that the central bank will be held accountable for meeting the inflation targets. For example, Section 8 of the Reserve Bank of New Zealand Act of 1989 assigns the central bank the statutory responsibility "to formulate and implement monetary policy directed to the economic objective of achieving and maintaining stability in the general level of prices," with no mention of competing goals. Section 9 of the act requires the Minister of Finance and the Governor of the Reserve Bank to negotiate and make public a Policy Targets Agreement (PTA), setting out specific inflation targets. In other countries, such as Switzerland, Canada and the United Kingdom, the inflation goal is embodied in public statements by the central bank rather than mandated by law.

The rationale for treating inflation as the primary goal of monetary policy is clearly strongest when medium- to long-term horizons are considered, as most economists agree that monetary policy can affect real quantities, such as output and employment, only in the short run. Of course, some economists of new classical or monetarist persuasions might claim that inflation should be the sole concern of monetary policy in the short run as well, arguing that using monetary policy for short-run stabilization of the real economy is undesirable, infeasible, or both. However, in practice no central bank has of yet completely forsworn the use of monetary policy for short-run stabilization, and so the phraseology "primary" or "overriding" must be taken to refer to the longer term.

The degree to which the central bank is held formally accountable for inflation outcomes varies considerably. The New Zealand law links the tenure of the governor of the Reserve Bank to the achieving of the inflation targets, and thus comes closest to providing an explicit "incentive contract," as proposed by Persson and Tabellini (1993) and Walsh (1995).[5] In other countries, no explicit sanctions on the central bank for missing the target are given; presumably, however, missing the target badly would impose implicit institutional or personal costs in terms of lost reputation or prestige. It is rather early in many of the inflation-targeting experiments to judge

[3] Technically, ensuring only that the inflation rate is stationary may leave a unit root in the price level, so that the forecast variance of the price level grows without bound. This problem is analogous to the issue of "base drift" in the literature on money-growth targeting.

[4] However, Svensson (1996) gives examples in which price-level targeting actually reduces the volatility of output.

[5] Svensson (1997b) relates inflation targeting to the contracting approach.

the extent to which the prospective penalties for missing announced targets will constrain central bank behavior.

Despite the language referring to inflation control as the primary objective of monetary policy, as we have said, inflation-targeting central banks always make room for short-run stabilization objectives, particularly with respect to output and exchange rates.[6] This accommodation of short-run stabilization goals is accomplished through several means. First, the price index on which the official inflation targets are based is often defined to exclude or down-weight the effects of "supply shocks;" for example, the officially targeted price index may exclude some combination of food and energy prices, indirect tax changes, terms-of-trade shocks, and the direct effects of interest rate changes on the index (for example, through imputed rental costs). Second, as already noted, inflation targets are typically specified as a range; the use of ranges generally reflect not only uncertainty about the link between policy levers and inflation outcomes but is also intended to allow the central bank some flexibility in the short run. Third, short-term inflation targets can and have been adjusted to accommodate supply shocks or other exogenous changes in the inflation rate outside the central bank's control. A model here is the Deutsche Bundesbank's practice of stating its short-term (one-year) inflation projection as the level of "unavoidable inflation." In the aftermath of the 1979 oil shock, for example, the Bundesbank announced the "unavoidable" inflation rate to be 4 percent, then moved its target gradually down to 2 percent over a six-year period. In other cases, the central bank or government makes explicit an "escape clause," which permits the inflation target to be suspended or modified in the face of certain adverse economic developments.

In making inflation, a goal variable, the focus of monetary policy, the inflation-targeting strategy in most cases significantly reduces the role of formal intermediate targets, such as the exchange rate or money growth. To the extent that intermediate targets are used, it is emphasized that the inflation goal takes precedence in case of conflict. Unconditional commitment to an intermediate target is of course inconsistent with inflation targeting (except in the unusual case that the intermediate target effectively summarizes all current information about inflation at the forecast horizon). The fact that in most countries the relation between intermediate targets, such as money growth, and the central bank's goal variables has proven to be relatively unreliable—the so-called "velocity instability" problem—is a major motivation for dropping formal intermediate targets and instead attempting to target the goal variable directly.

On the other hand, since targeting inflation directly requires that the central bank form forecasts of the likely path of prices, close attention is typically paid to a variety of indicators that have shown predictive power for inflation in the past. For example, as an aid to inflation forecasting, monetary policymakers in Canada

[6] Another short-run objective that is almost always retained by inflation-targeting central banks is the maintenance of financial stability. For example, see Mishkin (forthcoming).

and Sweden make use of a "monetary conditions index," a weighted combination of the exchange rate and the short-term interest rate, in conjunction with other standard indicators such as money and credit aggregates, commodity prices, capacity utilization and wage developments.[7]

In most inflation-targeting regimes, the central bank publishes regular, detailed assessments of the inflation situation, including current forecasts of inflation and discussions of the policy response that is needed to keep inflation on track. A good example is the Bank of England's *Inflation Report*, published quarterly, which contains detailed analyses of factors likely to affect the inflation rate as well as probabilistic forecasts of inflation, assuming no change in interest rates. The central banks of Canada and Sweden release similar documents, and the Reserve Bank of New Zealand is required to issue a policy statement at least every six months. As we discuss further below, the use of such reports reflects a key objective of inflation targeting, which is improved communication with the public about monetary policy, its goals and, in particular, the long-run implications of current policy actions.

The adoption of inflation targeting is often linked with changes in the laws or administrative arrangements associated with the central bank. Typically, reforms are in the direction of increased independence for the central bank, particularly in respect to its choice of instrument settings.[8] This seems to be a logical consequence of making price stability the overriding goal of policy, since the central bank is the best place to make the technical decisions necessary to achieve price stability and to make judgments about whether the pursuit of other objectives is consistent with this goal. Exceptions to this observation are the United Kingdom and, to a lesser extent, Canada, where despite the commitment to inflation targeting, the government, rather than the central bank, retains the final control over monetary policy. However, even in the British case the adoption of inflation targeting seems to have increased the relative influence of the central bank, as the *Inflation Report* and the timely publication of the minutes of the monthly meeting between the Governor and the Chancellor of the Exchequer provide an independent forum for the bank to express its views; in effect, the government must rationalize for the public any deviations of its policies from those recommended by the bank.

Most or all of the characteristics of inflation targeting described in this section

[7] Users of the monetary conditions index would probably argue that treating the MCI simply as a forecasting variable is oversimple; they tend to view the MCI more specifically as a measure of how overall monetary conditions are affecting aggregate demand and thus as a potential guide to policy actions. See Freedman (1994) for further discussion.

[8] Debelle and Fischer (1994) make the useful distinction between goal independence and instrument independence for the central bank. Goal independence implies the unilateral ability of the central bank to set its inflation targets and other goals, while instrument independence means that, although goals may be set by the government or by the government in consultation with the central bank, the central bank is solely responsible for choosing the instrument settings (for example, the level of short-term interest rates) necessary to achieve those goals. Instrument independence would seem to be the form of independence that maximizes central bank accountability and minimizes opportunistic political interference, while still leaving the ultimate goals of policy to be determined by democratic processes.

apply to countries adopting this approach within the last eight years or so; as noted in the introduction, these include Canada, the United Kingdom, New Zealand, Sweden, Australia, Finland, Spain and Israel. Germany and Switzerland, which have conducted inflation-focused monetary policies since the mid-1970s, are better viewed as "hybrid" cases, which meet some but not all of the above criteria. These two countries differ from the "pure" inflation targeters primarily in their greater focus on money growth as an intermediate target, and indeed, the Bundesbank has emphasized the superiority (in their view) of money targeting as a means of insuring monetary discipline and transparency (for example, Deutsche Bundesbank, 1995, pp. 67–8). In fact, many observers (including ourselves) would argue that the distinction between inflation and money targeting is overstated and that monetary policies in both countries are driven in the medium and long term primarily by inflation goals. For example, the Bundesbank's money growth targets are derived, using the quantity equation, to be consistent with an annual inflation target, given projections of the growth of potential output and of possible changes in the velocity of money. This inflation target, in turn, has been brought down steadily over time and has remained at 2 percent—the level deemed consistent by the Bundesbank with price stability—since 1986. Further, the Bundesbank has shown itself quite willing to miss its money targets when pursuing these targets threatens to conflict with the control of inflation (von Hagen, 1995; Bernanke and Mihov, 1997).

All in all, the philosophy guiding German and Swiss monetary policies seems relatively consistent with the one motivating the self-declared inflation targeters. The main practical difference between the two sets of countries is that the Germans and Swiss believe that the velocity of money has been relatively more stable in their countries, and so they view money-growth targeting as a useful tool for implementing their inflation objectives. It is also true that Germany and Switzerland have been less explicit in stating their inflation targets; neither central bank publishes a regular inflation report per se. But this distinction seems relatively unimportant; inflation developments receive prominent attention in the regular publications of both banks. Moreover, there may be less need for public declarations given the long-standing commitment of the Bundesbank and Swiss National Bank—and the popular support for that commitment—to price stability. The examples of Germany and Switzerland are important because, unlike the other countries mentioned, these two countries have been following their monetary policy strategies fairly consistently for more than two decades, rather than for only a few years; thus, their experiences may provide researchers attempting to assess the value of inflation-focused monetary policy with useful information.

A Framework, Not a Rule

The motivations for an inflation-targeting approach have been varied. In a number of cases, such as those of the United Kingdom and Sweden, the collapse of an exchange rate peg led the monetary authorities to search for an alternative

"nominal anchor" for monetary policy, a way of reassuring the public that monetary policy would remain disciplined. The demise of a fixed-exchange-rate regime similarly motivated the adoption of a money-focused approach by Germany in the mid-1970s. Some countries, such as Canada, came to inflation targeting after unsuccessful attempts to use a money-targeting approach. For example, in the case of Canada, by 1980 inflation was as high as it was in 1975 (10 percent per year) despite adherence to monetary targets that led to lower money growth rates (Howitt, 1993). In other cases, countries that by tight monetary policies had succeeded in reducing their core rate of inflation adopted inflation targeting as an institutional means of locking in their inflation gains.

Developments in macroeconomic theory also played some role in the growing popularity of the inflation targeting approach. These familiar developments included reduced confidence in activist, countercyclical monetary policy; the widespread acceptance of the view that there is no long-run tradeoff between output (or unemployment) and inflation, so that monetary policy affects only prices in the long run; theoretical arguments for the value of precommitment and credibility in monetary policy (Kydland and Prescott, 1977; Calvo, 1978; Barro and Gordon, 1983); and an increasing acceptance of the proposition that low inflation promotes long-run economic growth and efficiency.

Unfortunately, the interpretation of inflation targeting in terms of some longstanding debates in monetary economics has also been the source of confusion. For many years the principal debate about the best approach for monetary policy was framed as an opposition between two polar strategies, termed "rules" and "discretion." Advocates of rules—such as the fixed rule for money growth proposed by Milton Friedman, or a gold standard—argued that "tying the hands" of policymakers will prevent the monetary authorities from implementing counterproductive attempts at short-run stabilization and will thus eliminate the inflationary bias inherent in discretionary monetary policy. Supporters of discretionary policymaking—under which the central bank is left free to "do the right thing" as economic conditions evolve—stress the inability of ironclad rules to deal with unforeseen shocks or changes in the structure of the economy.

For various reasons, including the rhetoric of some of its proponents, inflation targeting is sometimes interpreted as falling on the "rule" side of this traditional dichotomy (for example, Friedman and Kuttner, 1996). We view this characterization of inflation targeting as a mistake; indeed, we would go farther and say that the traditional dichotomy of monetary policy strategies into rules and discretion is itself misleading. In particular, some useful policy strategies are "rule-like," in that by their forward-looking nature they constrain central banks from systematically engaging in policies with undesirable long-run consequences; but which also allow some discretion for dealing with unforeseen or unusual circumstances. These hybrid or intermediate approaches may be said to subject the central bank to "constrained discretion." We argue below that inflation targeting should be viewed in this way, rather than as a rigid policy rule.

If inflation targeting is interpreted as a rule in the classic Friedman sense, then

it would have to be conceded that this approach is vulnerable to some important criticisms. First, the idea that monetary policy has (essentially) no legitimate goals besides inflation would find little support among central bankers, the public and most monetary economists. Second, given that central banks do care about output, employment, exchange rates and other variables besides inflation, treating inflation targeting as a literal rule could lead to very poor economic outcomes. As Friedman and Kuttner (1996) emphasize, much in the same way that money-growth targeting in the United States was done in by unpredicted shocks to the velocity of money, so an exclusive emphasis on inflation goals could lead to a highly unstable real economy should there be significant supply shocks, such as large changes in the price of oil.

Finally, critics of inflation targeting *as a rule* might well ask what is gained by the loss of flexibility entailed by precommitting monetary policy in this way. The academic literature on rules argues that tying the hands of policymakers will reduce the inflation bias of discretionary policy and perhaps allow for less costly disinflations, as increased credibility leads inflation expectations to moderate more quickly. However, critics of inflation targeting could point out that, although inflation-targeting countries have generally achieved and maintained low rates of inflation, little evidence supports the view that these reduced rates of inflation have been obtained at a lower sacrifice of output and employment than disinflations pursued under alternative regimes (at least so far). Even the Deutsche Bundesbank and the Swiss National Bank, whose pursuit of low inflation over the last two decades has presumably given the maximum credibility, have been able to achieve inflation reductions only at high costs in lost output and employment (Debelle and Fischer, 1994; Posen, 1995). Nor is there evidence that the introduction of inflation targets materially affects private-sector expectations of inflation, as revealed either by surveys or by the level of long-term nominal interest rates. Inflation expectations have come down, in most cases, only as inflation-targeting central banks have demonstrated that they can deliver low inflation (Posen and Laubach, 1996).

These objections are certainly important, as far as they go. However, again, they derive much of their force from the assumption that inflation targeting is to be viewed as an ironclad rule. As we have said, we believe that interpreting inflation targeting as a type of monetary policy rule is a fundamental mischaracterization of this approach *as it is actually practiced by contemporary central banks.* First, at a technical level, inflation targeting does not qualify as a policy rule in that it does not provide simple and mechanical operational instructions to the central bank. Rather, the inflation targeting approach enjoins the central bank to use its structural and judgmental models of the economy, in conjunction with all relevant information, to determine the policy action most likely to achieve the inflation target, and then to take that action. Unlike simple policy rules, inflation targeting never requires that the central bank ignore information that bears on its achieving its objectives. Second, and more importantly, inflation targeting as it is actually practiced contains a considerable degree of what most economists would define as policy discretion. Within the general constraints imposed by their medium- to long-term inflation

targets, central bankers have in practice left themselves considerable scope to respond to current unemployment conditions, exchange rates and other short-run developments.

The 1989 reform of the Reserve Bank of New Zealand, for example, is often held up as an example of the rule-making impulse. It is important to note that New Zealand is the most extreme of all the inflation-targeting countries in its use of formal institutional constraints on policy. Even so, the New Zealand law does provide the central bank some discretion and flexibility; for example, the target inflation series excludes movements in commodity prices; the target may be readjusted if necessary in the judgment of the bank in response to supply or terms-of-trade shocks; the inflation target is specified as a 3 percentage point range rather than as a single number; and there is an explicit escape clause that permits amending the target in the face of unexpected developments. In practice, inflation targeting in New Zealand has been implemented even more flexibly. Inflation was brought down to its current low level only gradually; and when inflation moved briefly above the target range in 1996, the Parliament did not seriously consider its option of replacing the governor of the central bank.

If inflation targeting is not a rule in the way this term is usually understood, then what is it, and what good is it? We believe that it is most fruitful to think of inflation targeting not as a rule, but as a framework for monetary policy within which "constrained discretion" can be exercised. This framework has the potential to serve two important functions: improving communication between policymakers and the public, and providing increased discipline and accountability for monetary policy.

In terms of communication, the announcement of inflation targets clarifies the central bank's intentions for the markets and for the general public, reducing uncertainty about the future course of inflation. (Of course, this assumes that the announcements are believable and believed; more on this later.) Arguably, many of the costs of inflation arise from its uncertainty or variability more than from its level. Uncertain inflation complicates long-term saving and investment decisions, exacerbates relative price volatility, and increases the riskiness of nominal financial and wage contracts. Uncertainty about central bank intentions may also induce volatility in financial markets—a common phenomenon in the United States, where stock market analysts parse every sentence uttered by the Fed chairman in search of hidden meanings. Inflation targets offer transparency of policy; they make explicit the central bank's policy intentions in a way that should improve private-sector planning, enhance the possibility of public debate about the direction of monetary policy, and increase central bank accountability. Transparency has been claimed as a positive feature of other policy strategies, such as money-growth targeting, but we doubt that concepts like the growth rates of particular money aggregates are nearly so understandable to the general public as is the predicted rate of change of consumer prices.

To see the practical advantage of policy transparency, consider the familiar scenario in which an upcoming election or a slow economic recovery induces the

government to pressure the central bank to apply some short-run stimulus. In an inflation-targeting regime, the central bank would be able—indeed, would be required—to make explicit that the short-run benefits of this policy (faster real growth) may well be purchased at the price of medium- and long-term inflation. These projections could then be debated by politicians, press and public, but at least the issue of long-run inflation effects would be on the table, serving as an explicit counterweight to the short-run benefits of monetary expansion. Making the linkage of short-term policies and long-term consequences explicit would clarify for the public what monetary policy can and cannot do.

Aggregate supply shocks, such as oil price shocks, present a thornier policy problem. If a severe supply shock hits the economy, keeping medium-term inflation close to the long-run target could well be very costly in terms of lost output. However, in practice, a well-implemented inflation-targeting regime need not strongly constrain the ability of the monetary authorities to respond to a supply shock. Remember, the inflation target itself can be and typically is defined to exclude at least the first-round effects of some important supply shocks, such as changes in the prices of food and energy or in value-added taxes; the use of target ranges for inflation gives additional flexibility. Escape clauses, which permit the central bank to change its medium-term targets in response to major developments, are another possibility. We have seen, for example, that the Bundesbank's one-year inflation targets were often defined by its view of how much inflation was "unavoidable," rather than by its long-run objective of price stability. Thus, intermediate-run inflation targets can be used to define a transition path by which the temporary inflation induced by a supply shock is eliminated gradually over time. Relative to a purely discretionary approach, the inflation-targeting framework should give the central bank a better chance of convincing the public that the consequences of the supply shock are only a one-time rise in the price level, rather than a permanent increase in inflation. A relevant example occurred in Canada in 1991, shortly after their implementation of inflation targeting, when a sharp increase in indirect taxes caused a blip in the price level but had no apparent effect on the underlying inflation rate.

The idea that inflation targeting requires an accounting of the long-run implications of short-run "discretionary" actions is also central to the argument that inflation targeting helps to discipline monetary policy. In practice, exactly who needs disciplining may differ from country to country, depending on politics, institutional arrangements and personalities. In the macroeconomic literature on central bank credibility, it is the central bank that needs discipline, because it is assumed to desire an unemployment rate lower than the natural rate. This desire leads the monetary authority to try to "fool" the public with surprise inflation, inducing producers (who confuse nominal and real price increases) to increase output and employment above the natural rate. If the public has rational expectations, however, it will anticipate the central bank's actions, and producers will not be fooled, so that in equilibrium the economy

will suffer higher-than-optimal inflation with no benefits in terms of lower unemployment.[9]

If a story along these lines describes the actual situation in a given economy, then an inflation-targeting framework will not *directly* prevent the counterproductive attempts of the central bank to engage in excessive short-run stimulus. In this respect, inflation targeting is inferior to an ironclad rule, if such could be implemented. However, in contrast to the purely discretionary situation with no explicit targets, under inflation targeting the central bank would be forced to calculate and to publicize the implications of its short-run actions for expected inflation in the long run (and again, these projections would be subject to scrutiny and debate). To the extent that the central bank governors dislike admitting publicly that they are off track with respect to their long-run inflation targets, the existence of this framework would provide an additional incentive for the central bank to limit its short-run opportunism.

Although the theoretical literature typically posits the central bank as the entity who chooses to inflate opportunistically, we suspect that in most cases the executive and legislative branches of the government have the greater incentive to engage in such behavior, often because of approaching elections. Central bankers, in contrast, tend to view themselves as defenders of the currency. This view may be the result of intentional appointments of "tough" central bankers (for reasons described by Rogoff, 1985), or it may just be that self-selection and socialization act to make central bankers relatively hawkish on inflation. But in either case, the existence of longer-term inflation targets can prove a useful device by which the central bank can protect itself politically from overexpansionist pressures. In particular, by making explicit the long-run, as well as the short-run, implications of overexpansionist policies, the central bank may be better able to get the support it needs to resist such policies. Our impression is that the Bank of England, for example, has on occasion used numerical inflation targets in precisely this way.

Further Issues with Inflation Targeting

If viewed as a framework rather than as a rule, inflation targeting can confer some important advantages. It provides a nominal anchor for policy and the economy. By communicating the central bank's objectives and views, it increases the transparency of monetary policy. It has the potential to provide increased discipline and accountability for policymakers. Importantly, it may be able to achieve all this without entirely giving up the benefits of discretionary policies in the short run. These optimistic conclusions notwithstanding, important questions and controver-

[9] McCallum (1997) argues that the central bank can simply choose not to behave myopically, and the public's expectations will come to reflect this more farsighted behavior. He also points out, however, that to the extent time inconsistency is a problem, it will affect the government as well as the central bank; we agree, as we discuss below.

sies remain around inflation targeting, even when interpreted in the way that we prefer. Let us consider a few of these.

Which Inflation Measure? What Target Value?

A critical aspect of the design of an inflation-targeting regime is the definition of the price series to be used in the inflation target. The series needs to be considered accurate, timely and readily understood by the public, but may also need to allow for individual price shocks or one-time shifts that do not affect trend inflation, which is what monetary policy should influence. As Table 1 indicates, all inflation-targeting countries have chosen some variant of the consumer price index (CPI) as their target series. However, this choice is not typically the "headline" CPI figure, but an index that excludes some components or focuses on "core" inflation; clearly, it is incumbent on the central bank to explain its choice of index and to help the public understand its relation to the headline index.

In all inflation-targeting regimes, the inflation objective has been set at a low number, 4 percent or less. Is this the ideal range for the inflation target? Or would a somewhat higher range for inflation, which might involve lower initial output cost to attain, be acceptable?

Obtaining direct empirical confirmation of a link between inflation and economic performance is very difficult. Inflation is, after all, an endogenous variable; and so we rarely if ever see variation in inflation that is not associated with some third factor, such as supply shocks or political instability, which would plausibly affect other elements of economic performance as well.[10] As a result, economists' views on the subject have been based largely on prior arguments, intuition and indirect evidence. That conceded, it is nevertheless clear that the professional consensus, which at one time did not ascribe substantial costs to moderate inflation, has over the past few decades begun to take the costs of inflation more seriously. For example, Feldstein (1996) has emphasized the importance of inflation-induced inefficiencies, via the tax code, on capital formation. Fischer (1993) and others have provided some evidence that macroeconomic stability, including control of inflation, is an important precondition for economic growth. Shiller's (1996) opinion surveys of public attitudes about inflation, while confirming economists' suspicions that the public is confused about even the definition of inflation, also show that people believe inflation to be highly uneven in its distributional impacts and hence corrosive of the social compact. A strengthening preference for low inflation is quite visible in policy circles, perhaps most strikingly in the tough limits on inflation imposed by the Maastricht treaty on countries that want to join the European currency union.

Given the growing consensus that the long-term goal of monetary policy should be a low inflation rate, there remains the question of how low it should

[10] Studies that attempt to overcome these problems include Lebow, Roberts and Stockton (1992) and Barro (1995).

be. It seems clear that an inflation target of zero or near zero is not desirable, for several reasons. First, much recent research suggests that official CPI inflation rates tend to overstate the true rate of inflation, due to various problems such as substitution bias in the fixed-weight index and failure to account adequately for quality change. Studies for the United States have estimated this overstatement of inflation to be in the range of 0.5 to 2.0 percentage points per year.[11] Thus, as a practical matter, even if the central bank chooses to pursue a zero rate of true inflation, the target for the measured inflation rate should be greater than zero.

Putting aside measurement issues, there are other risks of setting the inflation target too low. In a much discussed recent article, Akerlof, Dickens and Perry (1996) point out that if nominal wages are rigid downward (a possibility that they argue is consistent with the evidence), then reductions in real wages can occur only through inflation in the general price level. Very low inflation therefore effectively reduces real-wage flexibility and hence may worsen the allocative efficiency of the labor market; indeed, the authors perform simulations suggesting that inflation rates near zero would permanently increase the natural rate of unemployment.[12] Another danger of setting the inflation target too low is that there is a greater chance that the economy will be tipped into deflation, with the true price level actually falling—as may have happened during the recent recession in Japan. As pointed out in the literature on financial crises, persistent deflation—particularly if unanticipated—can create serious problems for the financial system, interfering with its normal functioning and precipitating an economic contraction (Bernanke and James, 1991; Mishkin, 1991).

These risks suggest that the inflation target, even when corrected for measurement error, should be set above zero, as has been the practice of all inflation-targeting countries to date. Indeed, a potentially important advantage of inflation targeting is that it provides not only a ceiling for the inflation rate, but also a floor. Inflation targeting thus acts to attenuate the effects of negative, as well as positive, shocks to aggregate demand. An interesting historical example is that of Sweden in the 1930s, which adopted a "norm of price stabilization" after leaving the gold standard in 1931. As a result, Sweden did not undergo the devastating deflation experienced by other countries during the Great Depression (Jonung, 1979).

[11] This bias was the subject of an official report to the Senate Finance Committee, the so-called Boskin report (Boskin et al., 1996). See also Moulton (1996) and Shapiro and Wilcox (1997).

[12] The force of this argument should not be overstated. First, the inflation rates which Akerlof, Dickens and Perry (1996) argue would significantly affect the natural rate of unemployment are really quite low, for example, measured rates (as opposed to "true" rates) of inflation of 2 percent per annum or less. Second, their simulation studies do not take into account forces that may work in the opposite direction: for example, Groshen and Schweitzer (1996) point out that high and variable inflation rates may increase the "noise" in relative wages, reducing the efficiency of the process by which workers are allocated across industries and occupations; thus higher inflation can represent sand as well as grease in the wheels of the labor market.

Is Inflation Sufficiently Predictable and Controllable to be "Targeted"?

It has been noted by several authors that inflation is very difficult to predict accurately, particularly at both very short and very long horizons (Cecchetti, 1995). This lack of predictability poses two important problems for the inflation targeting strategy. The first is strictly operational: given the long lags between monetary policy actions and the inflation response, low predictability suggests that accurate targeting of inflation could be extremely difficult. The second issue has to do with the central bank's credibility: if inflation is largely unpredictable, and hence not finely controllable, then it will be difficult to judge whether the central bank has made its best effort to hit the inflation targets. For example, the central bank could always argue that wide misses were the result of bad luck, not bad faith; since central bank forecasts of inflation contain substantial judgmental components, such claims would be difficult to disprove. This possible escape hatch for the central bank weakens the argument that inflation targeting increases accountability of monetary policy and suggests that building up credibility for its inflation-targeting framework could be a long and arduous process.

While we agree that inflation targeting is less effective, the less predictable or controllable is the inflation rate, several observations should be made. First, statistical measures of predictability are themselves likely to be sensitive to the monetary policy regime in place. Inflation was no doubt difficult to predict during the 1970s, when monetary policymakers tried to deal with oil price shocks and other stagflationary pressures without a coherent, clearly articulated framework. In contrast, the stability of the inflation rate in the United States and other industrialized countries since the mid-1980s, a period during which the maintenance of low and steady inflation has received much greater weight in central bank decision making, suggests that inflation will be easier to predict in the future.

Second, the relative unpredictability of goal variables is not in itself an argument for the use of intermediate targets in the conduct of monetary policy. As Svensson (1997a) points out, from an optimal control perspective, the best possible intermediate target is the current forecast of the goal variable itself—in this context, inflation. Using an intermediate target such as money growth is acceptable in an optimal control framework only if the intermediate target contains all information relevant to forecasting the goal variable; in this extreme case, using the intermediate target is equivalent to targeting the forecast of the goal variable. However, if any variable other than the intermediate target contains marginal information about the future values of the goal variable, then targeting the inflation forecast strictly dominates using any single intermediate target. Thus, from a strictly operational point of view, while it is unfortunate if the goal variable is hard to predict or to control, no improvement is available by using an intermediate target.[13]

[13] In characterizing the forecast of inflation as the intermediate target, Svensson (1997a) is careful to define "forecast" to mean the forecast derived internally by the central bank using its structural model of the economy. An intriguing alternative would be to try to "target" private-sector forecasts of inflation, that is, set short-run policy instruments so that private-sector forecasts of inflation equal the announced

When the credibility of the central bank is at issue, the problem of whether to target inflation directly or to rely on an intermediate target becomes more complex. By Svensson's argument, use of the intermediate target must increase the variance of the goal variable, which is a cost of the intermediate targeting approach; the benefit, however, is that by hitting its announced target for the intermediate variable, the central bank can demonstrate the seriousness of its intentions to the public more quickly and reliably (Cukierman, 1995; Laubach, 1996). If credibility building is an important objective of the central bank, and if there exists an intermediate target variable—such as a monetary aggregate—that is well controlled by the central bank, observed and understood by the public and the financial markets, and strongly and reliably related to the ultimate goal variable, then targeting the intermediate variable may be the preferred strategy. All of these are big "ifs," particularly the last one. However, this analysis may help to explain the continued use of money-growth targets by Germany and Switzerland, where financial institutions and hence velocity have evolved rather slowly, while countries such as the United Kingdom, with a history of unstable velocity, have opted for targeting inflation directly.

Is Inflation the Right Goal Variable for Monetary Policy?

The consensus that monetary policy is neutral in the long run restricts the set of feasible long-run goal variables for monetary policy, but inflation is not the only possibility. Notably, a number of economists have proposed that central banks should target the growth rate of nominal GDP rather than inflation (Taylor, 1985; Hall and Mankiw, 1994). Nominal GDP growth, which can be thought of as "velocity-corrected" money growth (that is, if velocity were constant, nominal GDP growth and money growth would be equal, by definition), has the advantage that it does put some weight on output as well as prices. Under a nominal GDP target, a decline in projected real output growth would automatically imply an increase in the central bank's inflation target, which would tend to be stabilizing.[14] Also, Cecchetti (1995) has presented simulations that suggest that policies directed to stabilizing nominal GDP growth may be more likely to produce good economic outcomes, given the difficulty of predicting and controlling inflation.

Nominal GDP targeting is a reasonable alternative to inflation targeting, and one that is generally consistent with the overall strategy for monetary policy discussed in this article. However, we have three reasons for mildly preferring

target. Unfortunately, as shown by Woodford (1994) and Bernanke and Woodford (1996), such a policy is usually not consistent with the existence of a unique rational expectations equilibrium. However, Bernanke and Woodford also show that, while targeting private-sector forecasts is not a good idea, private-sector forecasts can typically be combined with the central bank's own information to improve the efficiency of its operating procedure. Further, private-sector forecasts that the public observes to be close to the central bank's official targets may help to provide some validation of the bank's internal procedures for forecasting and controlling inflation.

[14] Hall and Mankiw (1994) point out, however, that the equal weighting of real output growth and inflation implied by a nominal GDP targeting is not necessarily the optimal one; in general, the relative weight put on the two goal variables should reflect social preferences.

inflation targets to nominal GDP targets. First, information on prices is more timely and frequently received than data on nominal GDP (and could be made even more so), a practical consideration that offsets some of the theoretical appeal of the nominal GDP target. Although collection of data on nominal GDP could also be improved, measurement of nominal GDP involves data on current quantities as well as current prices and thus is probably intrinsically more difficult to accomplish in a timely fashion. Second, given the various escape clauses and provisions for short-run flexibility built into the inflation-targeting approach, we doubt that there is much practical difference in the degree to which inflation targeting and nominal GDP targeting would allow accommodation of short-run stabilization objectives. Finally, and perhaps most important, it seems likely that the concept of inflation is better understood by the public than is the concept of nominal GDP, which could easily be confused with real GDP. If this is so, the objectives of communication and transparency would be better served by the use of an inflation target. As a matter of revealed preference, all central banks that have thus far adopted this general framework have chosen to target inflation rather than nominal GDP.

If It's Not Broke, Why Fix It?

Friedman and Kuttner (1996) decry the tendency of economists to want to impose restrictions and rules on central bank policymaking. They survey the problems with policy rules in the past, notably the failure of money-growth targeting to become a reliable policy framework in the United States, and they correctly point out that U.S. monetary policy has performed quite well in the recent past without the benefit of a formal rule or framework. Why, they ask, should we change something that is working well, especially given our inability to know what types of challenges will confront monetary policy in the future?

We would respond that a major reason for the success of the Volcker-Greenspan Fed is that it has employed a policymaking philosophy, or framework, which is de facto very similar to inflation targeting. In particular, the Fed has expressed a strong policy preference for low, steady inflation, and debates about short-run stabilization policies have prominently featured consideration of the long-term inflation implications of current Fed actions.

To take the next step and to formalize this framework would have several advantages. It would increase the transparency of the Fed's decision-making process, allowing more public debate and discussion of the Fed's strategy and tactics and, perhaps, reducing the financial and economic uncertainty associated with the Fed's current procedures. It would create an institutional commitment to the current approach that would be less dependent on a single individual's philosophy and might thus be expected to survive when, inevitably, new leadership takes over at the Fed. Finally, inflation targeting will be easiest to implement in a situation, like the current one, in which inflation is already low and the basic approach has been made familiar to the public and the markets. By adopting this approach now when it is relatively easy politically, we could ensure that the new procedures will

be in place to provide guidance when the next difficult decisions about monetary policy have to be made.

Conclusion

It is too early to offer a final judgment on whether inflation targeting will prove to be a fad or a trend. However, our preliminary assessment is that this approach—when construed as a framework for making monetary policy, rather than as a rigid rule—has a number of advantages, including more transparent and coherent policymaking, increased accountability, and greater attention to long-run considerations in day-to-day policy debates and decisions.

■ *We thank Alan Blinder, Brad De Long, Mervyn King, Don Kohn, Alan Krueger, Bennett McCallum, Michael Peytrignet, Adam Posen, Georg Rich, Julio Rotemberg, Lars Svensson and Timothy Taylor for their helpful comments. Any opinions expressed are those of the authors and not those of Princeton University, Columbia University, the National Bureau of Economic Research, the Federal Reserve Bank of New York, or the Federal Reserve System.*

References

Akerlof, George, William Dickens, and George Perry, "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity,* 1996, *1,* 1–59.

Barro, Robert, "Inflation and Economic Growth," *Bank of England Quarterly Bulletin,* May 1995, *35,* 166–76.

Barro, Robert, and David Gordon, "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics,* July 1983, *12,* 101–21.

Bernanke, Ben, and Harold James, "The Gold Standard, Deflation, and Financial Crisis in the Great Depression: An International Comparison." In Hubbard, R. G., ed., *Financial Markets and Financial Crises.* Chicago: University of Chicago Press for NBER, 1991, pp. 33–68.

Bernanke, Ben, and Ilian Mihov, "What Does the Bundesbank Target?," *European Economic Review,* forthcoming 1997.

Bernanke, Ben, and Michael Woodford, "Inflation Forecasts and Monetary Policy," unpublished paper, Princeton University, September 1996.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale Jorgenson, "Toward a More Accurate Measure of the Cost of Living: The Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index," December 4, 1996.

Calvo, Guillermo, "On the Time Consistency of Optimal Policy in a Monetary Economy," *Econometrica,* November 1978, *46,* 1411–28.

Cecchetti, Stephen, "Inflation Indicators and Inflation Policy," *NBER Macroeconomics Annual,* 1995, 189–219.

Cukierman, Alex, "Towards a Systematic Comparison Between Inflation Targets and Money Targets." In Leiderman, L., and L. Svensson, eds., *Inflation Targets.* London: Centre for Economic Policy Research, 1995, pp. 192–209.

Debelle, Guy, and Stanley Fischer, "How Independent Should a Central Bank Be?" In Fuh-

rer, Jeffrey, ed., *Goals, Guidelines, and Constraints Facing Monetary Policymakers.* Boston: Federal Reserve Bank of Boston, 1994, pp. 195–221.

Deutsche Bundesbank, *The Monetary Policy of the Bundesbank.* Frankfurt am Main: Deutsche Bundesbank, 1995.

Feldstein, Martin, "The Costs and Benefits of Going from Low Inflation to Price Stability." NBER Working Paper No. 5469, February 1996.

Fischer, Stanley, "The Role of Macroeconomic Factors in Growth." NBER Working Paper No. 4565, December 1993.

Freedman, Charles, "The Use of Indicators and of the Monetary Conditions Index in Canada." In Balino, T., and C. Cottarelli, eds., *Frameworks for Monetary Stability: Policy Issues and Country Experiences.* Washington, D.C.: International Monetary Fund, 1994, pp. 458–76.

Friedman, Ben, and Kenneth Kuttner, "A Price Target for U.S. Monetary Policy? Lessons from the Experience with Money Growth Targets," *Brookings Papers on Economic Activity,* 1996, *1,* 77–125.

Goodhart, Charles, and José Vinals, "Strategy and Tactics of Monetary Policy: Examples from Europe and the Antipodes." In Fuhrer, Jeffrey, ed., *Goals, Guidelines, and Constraints Facing Monetary Policymakers.* Boston: Federal Reserve Bank of Boston, 1994, pp. 139–87.

Groshen, Erica, and Mark Schweitzer, "The Effects of Inflation on Wage Adjustments in Firm-Level Data: Grease or Sand?" Staff Report No. 9, Federal Reserve Bank of New York, January 1996.

Haldane, Andrew G., ed., *Targeting Inflation.* London: Bank of England, 1995.

Hall, Robert, and N. Gregory Mankiw, "Nominal Income Targeting." In Mankiw, N. G., ed., *Monetary Policy.* Chicago: University of Chicago Press for NBER, 1994, pp. 71–94.

Howitt, Peter W., "Canada." In Fratianni, Michelle U., and Dominick Salvatore, eds., *Monetary Policy in Developed Economies, Handbook of Comparative Economic Policies.* Vol. 3, Westport: Greenwood Press, 1993, pp. 459–508.

Issing, Otmar, "Monetary Policy Strategies: Theoretical Basis, Empirical Findings, Practical Implementation." In Deutsche Bundesbank, ed., *Monetary Policy Strategies in Europe.* München: Verlag Franz Vahlen, 1996, pp 197–202.

Jonung, Lars, "Knut Wicksell's Norm of Price Stabilisation and Swedish Monetary Policy in the 1930s," *Journal of Monetary Economics,* October 1979, *5,* 459–96.

King, Mervyn, "Direct Inflation Targets." In Deutsche Bundesbank, ed., *Monetary Policy Strategies in Europe.* München: Verlag Franz Vahlen, 1996, pp. 45–75.

Kydland, Finn, and Edward Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy,* June 1977, *88,* 473–92.

Laubach, Thomas, "Signalling with Monetary and Inflation Targets," unpublished paper, Princeton University, August 1996.

Lebow, David, John Roberts, and David Stockton, "Economic Performance Under Price Stability." Working Paper No. 125, Division of Research and Statistics, Federal Reserve Board, 1992.

Leiderman, Leonardo, and Lars E. O. Svensson, eds., *Inflation Targets.* London: Centre for Economic Policy Research, 1995.

McCallum, Bennett, "Inflation Targeting in Canada, New Zealand, Sweden, the United Kingdom, and in General." NBER Working Paper No. 5579, May 1996.

McCallum, Bennett, "Crucial Issues Concerning Central Bank Independence." NBER Working Paper No. 5597; *Journal of Monetary Economics,* forthcoming 1997.

Mishkin, Frederic S., "Asymmetric Information and Financial Crises: A Historical Perspective." In Hubbard, R. Glenn, ed., *Financial Markets and Financial Crises.* Chicago: University of Chicago Press, 1991, pp. 69–108.

Mishkin, Frederic S., "What Monetary Policy Can and Cannot Do." In *Monetary Policy in Transition: Strategies, Instruments and Transmission Mechanisms.* Vienna: Oesterreichische Nationalbank, forthcoming.

Moulton, Brent, "Bias in the Consumer Price Index: What is the Evidence?," *Journal of Economic Perspectives,* Fall 1996, *10,* 159–77.

Persson, Torsten, and Guido Tabellini, "Designing Institutions for Monetary Stability," *Carnegie-Rochester Conference Series on Public Policy,* 1993, *39,* 53–84.

Posen, Adam, "Declarations are Not Enough: Financial Sector Sources of Central Bank Independence," *NBER Macroeconomics Annual,* 1995, 253–74.

Posen, Adam, and Thomas Laubach, "Some Comparative Evidence on the Effectiveness of Inflation Targets," unpublished paper, Federal Reserve Bank of New York, 1996.

Rogoff, Kenneth, "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics,* November 1985, *100,* 1169–89.

Shapiro, Matthew, and David Wilcox, "Mismeasurement in the Consumer Price Index: An

Evaluation," *NBER Macroeconomics Annual,* forthcoming 1997.

Shiller, Robert, "Why Do People Dislike Inflation?" Cowles Foundation Discussion Paper No. 1115, March 1996.

Svensson, Lars E. O., "Price Level Targeting vs. Inflation Targeting: A Free Lunch?" NBER Working Paper No. 5719, 1996.

Svensson, Lars E. O., "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." NBER Working Paper No. 5797; *European Economic Review,* forthcoming 1997a.

Svensson, Lars E. O., "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts." NBER Working Paper No. 5251; *American Economic Review,* forthcoming 1997b.

Taylor, John, "What Would Nominal GDP Targeting do to the Business Cycle?" In *Carnegie-Rochester Conference Series on Public Policy.* Vol. 22, Amsterdam: North-Holland, 1985, pp. 61–84.

von Hagen, Jurgen, "Inflation and Monetary Targeting in Germany." In Leiderman, L., and L. Svensson, eds., *Inflation Targets.* London: Centre for Economic Policy Research, 1995, pp. 107–21.

Walsh, Carl, "Optimal Contracts for Central Bankers," *American Economic Review,* March 1995, *85,* 150–67.

Woodford, Michael, "Nonstandard Indicators for Monetary Policy: Can Their Usefulness be Judged from Forecasting Regressions?" In Mankiw, N. G., ed., *Monetary Policy.* Chicago: University of Chicago Press for NBER, 1994, pp. 95–116.

NBER WORKING PAPER SERIES

MAKING DISCRETION IN MONETARY POLICY MORE RULE-LIKE

Frederic S. Mishkin

Working Paper 24135
http://www.nber.org/papers/w24135

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2017

This paper was being presented at the Federal Reserve Bank of Boston, Annual Conference, "Are Rules Made to Be Broken? Discretion and Monetary Policy," October 13-14, 2017. It's original title was "Improving the Use of Discretion in Monetary Policy." I thank participants in the macro lunch seminar at Columbia University for their helpful comments. The views expressed here are my own and are not necessarily those of the Federal Reserve Bank of Boston, Columbia University or the National Bureau of Economic Research. Disclosure of my outside compensated activities can be found on my website at http://www0.gsb.columbia.edu/faculty/fmishkin/. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Frederic S. Mishkin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Making Discretion in Monetary Policy More Rule-Like
Frederic S. Mishkin
NBER Working Paper No. 24135
December 2017
JEL No. E5,E52,E58

ABSTRACT

This paper argues that the rules versus discretion debate has been miscast because a central bank does not have to choose only between adopting a policy rule versus pure discretion, both of which have serious shortcomings. Rather it can choose a constrained discretionary regime that has rule-like attributes. Monetary policy discretion can be made more rule-like, by 1) adopting a nominal anchor such as an inflation target, and 2) communication of a monetary policy reaction process, especially through data-based forward guidance, in which the monetary policy authorities describe how the future policy path will change as economic circumstances change.

Frederic S. Mishkin
Columbia University
Graduate School of Business
Uris Hall 817
3022 Broadway
New York, NY 10027
and NBER
fsm3@columbia.edu

1 INTRODUCTION

The debate on whether a central bank should have a rule to set its policy instruments or alternatively to conduct monetary policy with discretion has heated up in recent years with proposed legislation in Congress, which has been passed by the House, to require the Federal Reserve to report on a "directive" rule similar to a Taylor (1993) rule for its policy instruments. In this paper, I discuss where I think the rules versus discretion debate is currently and argue that in a sense, this debate has been miscast, because a central bank does not have to choose only between adopting a policy rule versus pure discretion. Instead it can choose to adopt a discretionary regime that has rule-like attributes, a regime that Ben Bernanke and I referred to as "constrained discretion" in Bernanke and Mishkin (1997). However, how can be discretion be constrained so it avoids the disadvantages of pure discretion? The answer I provide here is that monetary policy discretion can be made more rule-like by pursuing monetary policy communication that not only constrains discretion, but also has additional benefits in enabling the markets to respond to shocks to the economy so both inflation and output are stabilized.

2 THE RULES VERSUS DISCRETION DEBATE

First, let's define our terms. A *rule* requires that monetary policy is essentially automatic: it involves a precise prescription for how monetary policy should react to a set of economic circumstances. One example of a monetary policy rule is the constant-money-growth rule advocated by Milton Friedman, in which the money supply is set by the central bank to grow at a constant rate. A more recent alternative is the classic Taylor (1993) rule in which the policy interest rate, the federal funds rate, is set to be a weighted average of an output gap (actual output minus potential output) and an inflation gap (actual inflation minus the target inflation rate.) The polar opposite of a monetary policy rule, according to the traditional classification of policy regimes is based on *discretion*. Discretion, in its purist form, involves monetary policymakers setting their policy instruments on a day-to-day basis as economic events unfold, with no public commitments about its objectives or actions.

2.1 The Case for Rules

There are two basic arguments for monetary policy rules. First is that monetary policy makers making discretionary decisions cannot be trusted. In their classic study of monetary policy history, Friedman and Schwartz (1963) document many instances where Federal Reserve policies led to sharp contractions in economic activity, as does Meltzer (2004, 2014) . Overly expansionary monetary policy that followed the period discussed by Friedman and Schwartz and Meltzer, led to the Great Inflation that lasted from the late 1960s to 1979, until actions by the Federal Reserve under the chairmanship of Paul Volcker led to a low inflation period since the early 1980s. Monetarists such as Friedman, Schwartz and Meltzer have argued that if the Federal Reserve had pursued a constant-money-growth rule these episodes of severe economic contractions and high inflation would have been avoided.

The second, and I would argue more powerful, argument for monetary policy rules results from the literature on the time-inconsistency problem described by Kydland and Prescott (1977), Calvo (1978) and Barro and Gordon (1983) in the context of monetary policy. The time-inconsistency problem occurs because economic agents or policy makers always have a temptation to deviate from an optimal long-run plan, which is therefore time-inconsistent, when they operate with discretion. In other words, even when an optimal long-run plan is formulated,

when tomorrow comes and the economic agent or policymaker reoptimizes, they renege on the optimal plan. In the case of monetary policy, monetary policy cannot achieve higher economic growth and lower unemployment in the long run by pursuing expansionary policy to produce higher inflation. Thus in the Barro-Gordon (1983) framework, the optimal plan pursues price stability, that is, a low and stable inflation rate. However, monetary policy makers are tempted to pursue a discretionary monetary policy that tries to exploit the short-run tradeoff between unemployment and inflation and so pursue more expansionary monetary policy than firms and people expect because such a policy would boost economic output and raise employment, thereby lowering unemployment in the short run. However, because there is no long-run tradeoff between unemployment and inflation, this discretionary policy only leads to the higher inflation in the long run, while it is unsuccessful in producing lower unemployment.

Elsewhere I have criticized this argument for the time-inconsistency problem in monetary policy because I do not believe that central bankers want to renege on an optimal plan to keep inflation low and stable. After all, by recognizing that there is no long-run tradeoff between unemployment and inflation, monetary policy makers can avoid the temptation to pursue overly expansionary monetary policy. Indeed, I have argued (Mishkin, 2016) that central bankers tend to be too conservative along the lines described by Rogoff (1985) and so have a tendency to not worry sufficiently about inflation being too low. Nonetheless, I believe that the time-inconsistency problem is a serious problem for monetary policy, not because it stems from central bankers, but because politicians are short-sighted and put pressure on central bankers to pursue overly expansionary monetary policy to lower unemployment in the short run. This pressure can lead to central bankers reneging on the optimal, low inflation plan, even if they would not renege otherwise. In other words, the time-inconsistency problem arises because central banks cannot be completely insulated from political pressure (Mishkin and Westelius, 2008).

A commitment to an instrument, monetary policy rule that embeds a nominal anchor is one way of avoiding the time-inconsistency problem. Once a monetary policy rule such as a constant-money-growth rule or a Taylor rule is adopted, monetary policy no longer can try to exploit the short-run tradeoff between unemployment and inflation because as the nominal anchor of either the money supply or the target level of inflation is exceeded, monetary policy

automatically tightens, so that inflation is stabilized. Furthermore, the presence of a nominal anchor in an instrument rule results in an anchoring of expected inflation, which results in a a smaller tradeoff between output gaps and inflation gaps.

There can also be a time-inconsistency problem in the other direction, when the inflation rate is too low, particularly when the policy interest rate has hit the zero lower bound and this nominal interest rate cannot go below zero. This tme-inconsistency problem can be avoided bu what Woodford (2003) refers to as history-dependent policy: that is, if the central bank could commit to keep the policy rate "lower for longer," so that when a the inflation objective is reached, the policy rate would be lower than it otherwise would be to keep the inflation rate at the desired level. However, once the desired inflation rate occurs, the central bank would have a temptation to renege on this commitment and raise the policy rate to a level that is consistent with keeping the inflation rate at the desired level. An instrument rule that sets the policy rate at a lower rate when the price level is below a target path for the price level would overcome this time-inconsistency problem.

2.2 The Case for Discretion

There are five main arguments against adoption of a monetary policy, instrument rule, and so argue in favor of the conduct of monetary policy with some discretion.

1. A rule requires a reliable model of the macroeconomy.
2. A rule requires that the structure of the economy is stable.
3. A rule cannot foresee every contingency.
4. A rule does not allow judgement.
5. Monetary policymakers are not less trustworthy than rules.

We look at each of these in turn.

2.2.1 A Rule Requires a Reliable Model of the Macroeconomy

For an instrument rule to produce stable inflation and output, policymakers must have a reliable model of the macroeconomy so that they can have confidence that the instrument rule they choose is close to the policy rule that minimizes output and inflation gaps.. For example, deriving a reliable Taylor rule requires that the central bank has confidence in its estimate of 1) the natural (equilibrium) rate of interest, 2) the natural rate of unemployment and 3) that there is a stable Phillips-curve relationship.

As we have seen recently, there have been major reassessments of the value of the natural (equilibrium) rate of interest. The Summary of Economic Projections, which provides projections from all the FOMC participants of long-term inflation and the long-term policy rate under appropriate policy , shows that the FOMC participants have substantially reduced their estimates of the natural rate of interest in recent years. Academic studies (Barsky, et. al., 2014, Curdia, et. al., 2014 and Hamilton, et. al., Harris, Hatzius and West, 2015) also suggest that the natural rate of interest has fallen in recent years, and they also emphasize how uncertain estimates of the natural rate of interest are.

Research also indicates that estimates of the natural rate of unemployment are highly uncertain (Staiger, Stock and Watson (1997) and Federal Reserve officials have cast doubts on whether the Phillips curve is sufficiently stable to provide a reliable guide to monetary policy (e.g., recent remarks by Brainard (2015) and Tarullo (2015). Indeed, Orphanides (2002) has

argued that the very high inflation outcomes in the United States in the 1970s were due to an underestimate of the natural rate of unemployment on the part of Federal Reserve policymakers.

2.2.2 A Rule Requires that the Structure of the Economy is Stable

A successful instrument rule requires that the structure of the economy does not undergo substantial changes so the instrument rule remains valid. The failure of monetary targeting in many countries in the 1980s indicates the dangers of adopting instrument rules. A particularly striking example occurred in Switzerland in the late 1980s, as documented in Bernanke, Laubach, Mishkin and Posen (1999). In 1980, the Swiss National Bank adopted a growth rate target for the monetary base. In 1988, the Swiss introduced a new interbank payment system, the Swiss Interbank Clearing (SIC), and a wide ranging revision of commercial banks liquidity requirements. These structural changes caused a severe drop in banks' desired holdings of deposits at the Swiss National Bank, and so a smaller amount of the monetary base was now needed relative to aggregate spending. The resulting upward shift in velocity meant that adherence to the monetary base target led to very high inflation, with Swiss inflation rising to above 5%, well above that of the rest of Europe, which of course horrified the anti-inflation Swiss. The problem with monetary targeting instrument rules is exemplified by the colorful quote from Gerald Bouey, the governor of the Bank of Canada in the 1980s, who said, "We didn't abandon monetary aggregates, they abandoned us."

2.2.3. A rule cannot foresee every contingency.

An instrument rule can be too rigid because it cannot foresee every contingency. For example almost no one could have predicted that problems in one small part of the financial system, subprime mortgage lending, would lead to the worst financial crisis since the Great Depression. The unprecedented steps that the Federal Reserve took during the crisis to prevent it from escalating into an even deeper crisis (Mishkin and White, 2016) could not have been written into a policy rule ahead of time.

2.2.4. A rule does not allow judgement.

An instrument rule does not easily incorporate the use of judgement. Monetary policy is as much an art as a science. Monetary policymakers look at a wide range of information in order to decide on the best course for monetary policy, and some of this information is not easily quantifiable. As an illustration, in the first go-round of every FOMC meeting, the Federal Reserve Bank presidents provide anecdotal information about what their contacts their Federal Reserve district tell them about the state of the economy, and this information is used by FOMC participants in their decisions about setting the monetary policy instrument. Another illustration is Alan Greenspan's use of judgement in the mid to late 1990s to argue against monetary policy tightening despite the rapidly growing economy and falling unemployment rate. Greenspan was able to convince the FOMC to refrain from raising rates despite recommendations from the Board's models that the FOMC do so. Greenspan was proved to be right, earning him the moniker of "maestro" (Woodward, 2000).

Judgement, which in its nature is discretionary, is thus an essential element of monetary policy to stabilize inflation and output, as has been emphasized by Svensson (2003, 2005). But

how do you put judgement into an instrument rule, when judgement necessarily is based on non-quantifiable information? There is no way that I can see how this might be done and so the use of judgement in monetary policy provides another strong argument against adoption of an instrument rule.

2.2.5 Monetary Policymakers are Not Less Trustworthy than Rules

One argument for adoption of rules is that they are more trustworthy than policymakers. There are certainly cases where this has been true in the past, as Friedman and Schwartz (1963) and Meltzer (2004, 2014) have documented. However, I would argue that beginning in the 1980s, monetary policy in the United States has produced more stable inflation and output, although not everyone would agree (e.g., Taylor, 2007). Inflation has been low and until the global financial crisis in 2007 business cycle fluctuations have been quite muted. Although I agree with John Taylor (2007) that monetary policy was too easy in the run up to the Great Recession because inflation did overshoot the 2% level, I do not agree with Taylor (2007) that monetary policy was a key reason why the housing bubble occurred for the reasons outlined by Bernanke (2010). However, the Federal Reserve's actions during the global financial crisis and in its aftermath resulted in smaller output and inflation gaps than would have occurred if monetary policy had been dictated by instrument rules, as the next subsection discusses.[1] Indeed, given the difficult circumstances, I think future historians will give the Fed very high grades on its performance during this period. Indeed, as argued below, the Federal Reserve was able to stabilize inflation and output during this period than would have a Taylor rule. Clearly the Federal Reserve learned a lot from the policy outcomes during the Great Depression and the

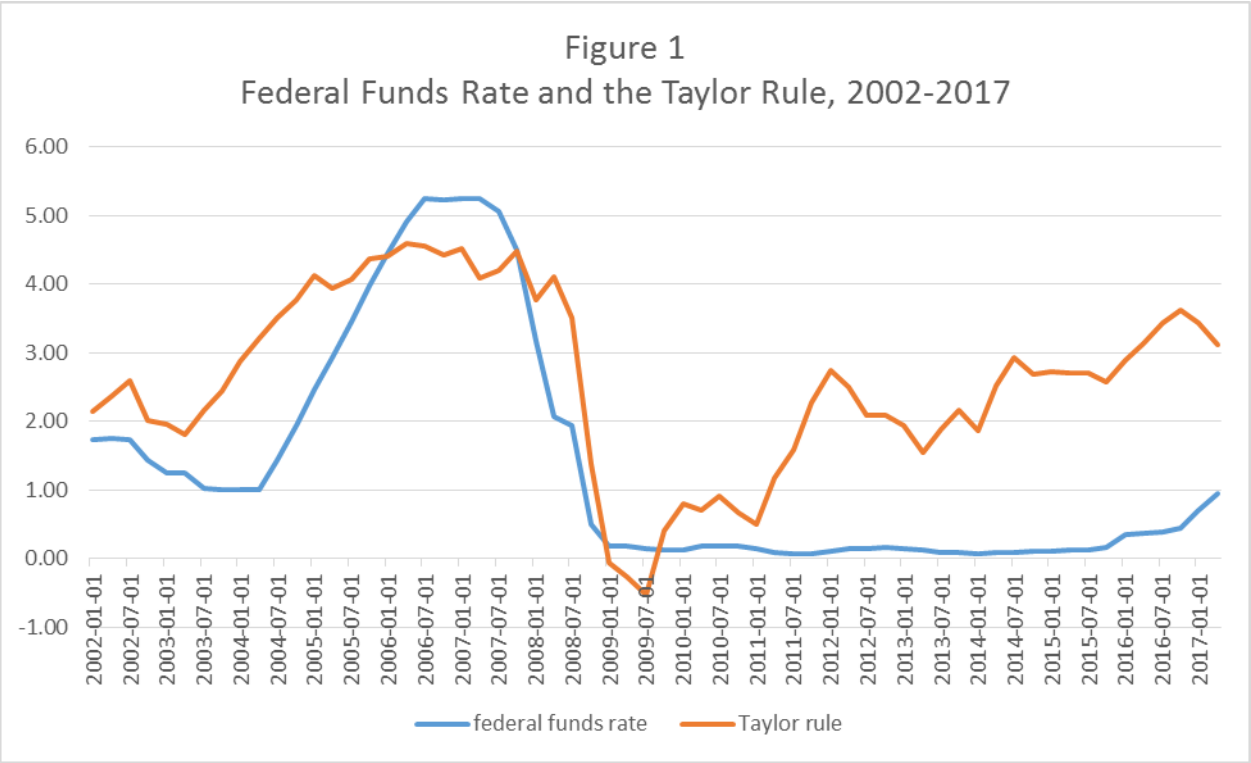
[1] Note that I have a conflict of interest in making this statement because I was an active participant in Federal Reserve decisions during the early stages of the financial crisis.

Great Inflation of the 1970s, which led to monetary policy that produced more stable inflation and output. Arguing that the Fed would be less successful at stabilizing output and inflation than would a policy instrument rule is now harder to make.

2.2.6 An Illustration Using a Traditional Taylor Rule

To illustrate the arguments above, let's see what a traditional Taylor (1993) rule would have recommended for monetary policy during the global financial crisis and its aftermath. This is only an illustration because more sophisticated instrument rules might do result in more stable inflation and output than the simple, traditional Taylor rule discussed here. However, I think that even more sophisticated Taylor or other instrument rules that would have reasonably been developed *before* the crisis would likely have run into similar problems.

Figure 1 provides a comparison of the actual federal funds rate with a Taylor (1993) rule using CBO estimates of potential output provided in the FRED database.



The global financial crisis started on August 2007 with the seizing up of interbank markets as a result of an announcement by the French bank BNP Paribas that is was suspending redemption of shares held in some of its money market funds. At the time, the U.S. economy was growing rapidly and inflation was rising. The Taylor rule which raises the federal funds rate when either the output gap or the inflation gap rises, would recommend a rise in the federal funds rate. Indeed, the 1993 Taylor rule shown in the figure would have recommended a rise in the federal funds rate of 38 basis points from the second quarter of 2007 to the fourth quarter of 2007. Instead, the FOMC decided to lower the federal funds rate by 25 basis points (0.25 percentage points) at the September FOMC meeting, and then by another 50 basis points in the fourth quarter, exactly the opposite of what the Taylor rule would have suggested.

As you can see in Figure 1, from the fourth quarter of 2007 to the second quarter of 2008, the Taylor rule kept the federal funds rate at fairly high level, around 4%, while the Federal Reserve lowered the federal funds 300 basis from the 5% level to the 2% level. The Taylor rule only recommended a substantial decline in the federal funds after the second quarter of 2008, but even then the FOMC's setting of the federal funds rate was well below the recommendation of the Taylor rule through the fourth quarter of 2008.

I think that very few economists would argue that the economy would have been grown faster and had inflation closer to the Fed's 2% objective if the Fed had followed the Taylor rule during this episode and kept the federal funds rate so high during the financial crisis. Indeed, given the course of the economy once the financial crisis began, the Federal Reserve can be criticized for not lowering the federal funds rate fast enough. Indeed, as is revealed by the FOMC transcripts, at the December 2007 FOMC meeting, in that meeting I argued strenuously for more pre-emptive easing of monetary policy and indicated that I would have preferred to dissent from the FOMC decision to lower the federal funds rate by only by 25 basis points, but I did not do so because it might have suggested to the markets that I was less supportive of the Chair, Ben Bernanke, then I was. Then in a speech that I gave on January 11, 2008 (Mishkin, 2008a), I provided the arguments why a more preemptive monetary policy was required to cope with the financial crisis. Subsequently, the Federal Reserve pursued a far more expansionary monetary policy, although it was insufficient to keep the economy going into the most severe recession of the post World War II period, with inflation falling well below the 2% inflation objective, which is also embedded in the Taylor rule. If the Federal Reserve had not used discretion and departed from the Taylor rule, inflation would have fallen even lower, so that there would have been a real possibility of a deflation setting in, as occurred in Japan after 1998,.

Furthermore, the decline in GDP would have been even more severe. Indeed, the possibility of the Great Recession turning into a full-scaled depression cannot be ruled out.

Why did the Federal Reserve depart from the Taylor rule? First, the global financial crisis revealed that the standard general equilibrium models macroeconomic models used by both central banks and academic economists were unreliable because they had ignored financial disruptions as an important factor in the evolution of the economy (see Mishkin, 2011). Second, is that the global financial crisis led to a change in the structure of the economy that made the economy more nonlinear. (However, because financial factors had not been built into general equilibrium macroeconomic models before the financial crisis, the change in structure could just be thought of as the pre-crisis macro models being misspecified.) Third, is that the financial crisis was not anticipated by macroeconomists, even those like myself who had studied financial crises in emerging market economies, but couldn't imagine that such a crisis would occur in a developed economy like the United States. The contingency of this financial crisis was therefore not even comtemplated in the literature on monetary policy rules. Fourth, is that the relatively rapid reactions of the Federal Reserve to the financial crisis and the departure from the Taylor rule were not based on hard data. Instead it was the judgement, very often based on anecdotal evidence and conversations with financial market participants, that led the Fed to depart from the Taylor rule. This was certainly true in my case, where my advocacy of more expansionary monetary policy to contain the crisis was not based primarily on hard data, but rather on my judgement that the financial disruption was going to have very adverse effects on the economy as a result of my reading and research on what had happened in previous financial crises.

All of the objections to adoption of an instrument rule that were articulated at the beginning of this section therefore came to play in the departure from the Taylor rule during the global financial crisis.

What happened after the crisis was over in 2009 also provide support for discretion and departures from the Taylor rule. As you can see in Figure 1, starting in the fourth quarter of 2009, the Taylor rule would have suggested that the federal funds rate should rise steadily to over a 3% rate currently. In contrast, the Federal Reserve has been raising rates very gradually, so by the second quarter of 2017 they are more than 200 basis points below the Taylor rule recommendation. Again, I think there are very few economists who would advocate that the Fed would have produced higher economic growth and inflation closer to the Fed's 2% target if it had followed the Taylor rule and raised the federal funds rate by 300 basis points, instead of the around 100 basis points that they actually did.

What is the rationale for this departure from the Taylor rule? First is that the equilibrium, natural rate of interest has been falling over time, although we don't know exactly by how much, while the Taylor rule has the equilibrium rate fixed at 2%. We could either see this as a misspecification of the model or alternatively that the structure of the economy has changed, leading to a decline in the equilibrium interest rate. In addition, inflation has remained below the 2% desired level, despite very low unemployment rates. This suggests that either the model of the Phillips curve is wrong or that the estimate of the natural rate of unemployment is too high. Stabilizing inflation and output in this environment where we are not sure what is going on requires the use of judgement.

We thus see that the objections to an instrument rule come into play again in arguing for the use of discretion in setting monetary policy in the period from late 2009 to the present.

2.3 Constrained Discretion: Making Discretion More Rule-Like

The arguments above argue against adoption of an instrument rule for monetary policy. However, we have also seen that pure discretion can lead to high inflation. But is there something in between?

Bernanke and Mishkin (1997) argue that the rules-versus-discretion debate has been miscast because the dichotomy between rules and discretion is too simple. Advocates of rules argue against *pure* discretion which is subject to the time-inconsistency problem, while advocates of discretion argue against *rigid* rules. Bernanke and I argued that by imposing a structure that imposes discipline on monetary policy, but does not eliminate flexibility, what we called *constrained discretion*, monetary policy could avoid some of the disadvantages of either rigid rules or pure discretion. Another way of thinking about constrained discretion is that it is an attempt to achieve the best of both worlds of rules and discretion by making discretion have rule-like properties, so that it avoids the time-inconsistency problem.

3 HOW CAN MONETARY POLICY DISCRETION BE MADE MORE RULE-LIKE?

Constrained discretion as an approach to the conduct of monetary policy can help stabilize output and inflation, but there is still the question of how a framework of constrained discretion can be designed so that it is rule-like, and so avoids the time-inconsistency problem. A first step in making monetary policy discretion more rule-like is for the monetary authorities to constrain discretion to avoid the time inconsistency problem by adopting a nominal anchor and being accountable to not deviate very far from it. However, higher accountability of monetary policy to achieve the nominal anchor can be enhanced by communication of the monetary policy reaction process.

Let's look at each in turn.

3.1 Adoption of a Nominal Anchor

An increasingly popular approach to adopting a nominal anchor and being accountable to not deviate from it is inflation targeting. As emphasized in Bernanke and Mishkin (1997), Bernanke, et. al. (1999) and Mishkin (1999), inflation targeting is a form of constrained discretion that can avoid the time-inconsistency problem by not only announcing an inflation target, but also by being accountable to achieve the target through communication about how the target is to be achieved and how past policy actions were consistent with achieving the inflation target. Other similar approaches adopt other target criteria with a nominal anchor such as price-level targeting or nominal GDP targeting, or alternatively, as in Woodford (2003), a target criterion that involves a tradeoff between output gaps and inflation gaps. These approaches are sometimes referred to as *target rules*. This is because they have rule-like properties that allow them to mitigate the time-inconsistency problem. However, these approaches are not rules in the sense that they provide an automatic prescription for how monetary policy is conducted. Instead they allow a lot of discretion on the part of monetary policymakers, including a lot of judgement as to how monetary policy instruments are set to achieve the target criterion.

It should also be pointed out that even if monetary policy makers do not explicitly adopt a nominal anchor by announcing a target for a nominal variable, they may do nearly as well by implicitly adopting a nominal anchor. For example, until January 2012, when the Federal Reserve adopted a 2% inflation target, the Federal Reserve did not have an explicit nominal anchor. Nonetheless, the Federal Reserve did emphasize that price stability was its most important long-term objective and expected inflation did become reasonably well-anchored both before and immediately after the global financial crisis. Indeed, one of the great successes of the Bernanke-Fed before 2012 was that it was able to anchor inflation expectations during the financial crisis through highly active, discretionary actions that prevented inflation expectations from plummeting, as they did in Japan. When I was a member of the Board of Governors of the Federal Reserve System, I nonetheless argued that adopting an explicit inflation target would help stabilize output and inflation relative to the Fed's monetary policy strategy at the time because it would make it more likely that inflation expectations would be anchored in the future (Mishkin, 2008b).

Advocates of instrument rules criticize inflation targeting for being too discretionary. I do not want to go into great detail as to how inflation targeting or its variants can constrain discretion to ensure that monetary policymakers are less tempted to renege on optimal plans. This is well-covered territory. The evidence, however, does support that countries that have adopted inflation targeting have been able to anchor inflation expectations well (Gürkaynak, et. al, 2010), which only occurs if this strategy is rule-like and overcomes the time-inconsistency problem. Furthermore, countries that have adopted inflation targeting have had inflation that is low and stable without bearing the cost of larger fluctuations of output (Mishkin and Schmidt-Hebbel, 2002).

Indeed, the excellent performance of inflation targeting, which can be viewed as a new technology for monetary policy, provides a counter to advocates of instrument rules who claim that central banks cannot be trusted to conduct discretionary monetary policy and more stable inflation and output would be achieved with an instrument rule. It is certainly true that central banks, such as the Federal Reserve, have found that some of their past policies have led to less stable inflation and output, but it is much harder to claim that central banks, particularly in developed countries with solid institutional frameworks, cannot be trusted to conduct monetary

policy when they are using the new technologies such as inflation targeting. Inflation targeting central banks have been able to avoid the time-inconsistency problem and keep inflation low and stable. I would argue that advocates of instrument rules are fighting the last war. They argue for an instrument rule to prevent a problem that is no longer occurring, high inflation.

3.2 Communicating a Monetary Policy Reaction Process

Although adopting a nominal anchor such as an inflation target can go a long way to constraining discretion and making it rule-like, it might not create enough accountability to constrain discretion sufficiently; the long lags from monetary policy to inflation mean that it may be several years before the monetary policy authorities can be monitored to see whether they took appropriate steps to achieve the target. Indeed, I suspect that this is why advocates of instrument rules do not feel that inflation targeting is sufficiently rule-like and that adoption of a rule would provide more stable inflation and output outcomes.

This concern that just announcing a nominal anchor is not sufficient to stabilize inflation can be addressed by improving the use of discretion with more communication about the monetary policy reaction process, that is, how monetary policy instruments would change as economic circumstances change. (Note that I use the term monetary reaction *process* rather than monetary policy reaction *function*. A function is a mathematical construct that generates an output from quantifiable inputs. Since, as I argue later, judgement requires that not all inputs into monetary policy decisions have to be quantifiable. Thus a description of monetary policy reactions to evolving economic circumstances is better described as a reaction process rather than a reaction function.) If the markets and the public have a better understanding of the policy reaction process, they can evaluate whether the setting of the policy instruments is consistent with achieving the nominal anchor target, thus increasing the accountability of the monetary policy authorities for achieving this target.

It is true that the markets can glean some information about the policy reaction process by seeing how actual monetary policy actions react to the incoming data, as the evidence presented below indicates. However, the monetary authorities can provide even more information about their policy reaction process and increase accountability by communicating

how the policy instruments would change as economic circumstances change and then explaining how the current setting of their policy instruments is consistent with this policy reaction process. Then as the economy evolves, the public and the markets can assess whether monetary policy is trying to achieve the nominal anchor, even before the outcomes on this anchor are revealed. However, communication about the policy reaction process can be increased by the use of what Feroli et. al. (2017) refer to as *data-based forward guidance*: that is, providing information on the future path of the policy rate conditional on the data that could occur over the policy horizon.

3.2.1 Data-Based Forward Guidance

Data-based forward guidance requires not only that the central bank provides information on the policy path given the central bank's forecasts, but also communicates how that path would change if the central bank's forecast changes. One central bank that does this is the central bank of Norway, the Norges Bank. The Norges Bank does this by first providing a baseline projected policy path along with a fan chart showing the confidence intervals around the policy path. More importantly, the Norges Bank explains why the policy path may deviate from the baseline path by providing several scenarios as to how the policy path would change when economic outcomes change. It is true that because the governance structure of the FOMC includes up to nineteen particpants (seven governors and twelve Federal Reserve Bank presidents) that make policy decisions, it would be extremely difficult to derive one probability distribution for the policy path outcomes or an agreement on how the policy path would change when economic outcomes change. This difficulty occurs because the Federal Reserve does not speak with one voice, as occurs with the Norges Bank. However, FOMC participants could increase communicating how their projections of the policy path would react to economic events. Alternatively, the FOMC might delegate to the Chair to provide information on how the committee's view of the future policy path might change under different scenarios for how the economy evolves.

This form of data-based forward guidance would not only make it easier for the public and the markets to evaluate whether the actual setting of policy instruments is consistent with the

objective of hitting the nominal anchor target, but also has expectations dynamics that can help stabilize output and inflation.

To see why data-based forward guidance leads to expectations dynamics that stabilize inflation and output, consider a negative shock to aggregate demand when both the inflation gap and output gap are at zero. The result would be that both the inflation and output gaps would turn negative in the future and an monetary policy reaction process that minimizes output and inflation gaps would indicate that the federal funds rate path would be lowered. If the central bank's reaction process is well understood by the markets and the public, then without -the Fed taking any actions, expectations of the future federal funds rate would decline, which would result in lower longer-term interest rates and stimulate the economy. The result would then be an immediate offset to the negative aggregate demand shock which would help stabilize the economy. Another way of stating this result is that successful central bank communication about the monetary policy reaction process would enable the markets to do a lot of the work (heavy lifting) for the central bank. If the monetary policy reaction to shocks is predictable, expectation dynamics work to tighten or loosen financial conditions that stabilize output and inflation appropriately when there are shocks to the economy.

However, forward guidance is often not done in a data-based way. As described in Feroli et. al. (2017) the Federal Reserve has often engaged in a second type of forward guidance, *time-based forward guidance,* in which a central bank commits to set the policy rate at specific levels at specific calendar dates. An extreme version of time-based forward guidance would be a central bank committing not to raise interest rates from their current level for several years. Such a commitment would ignore incoming information, which is why the forward guidance is *time-based*.

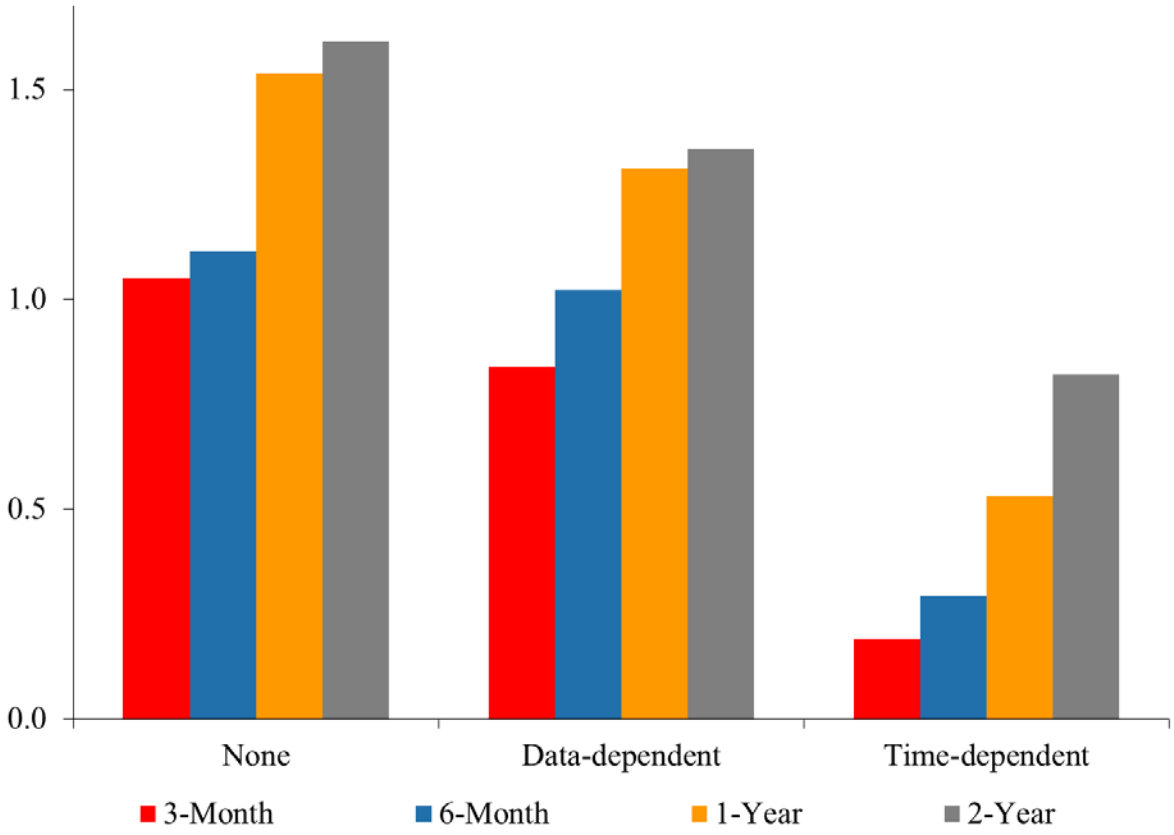
Not only does this time-based forward guidance provide less information about the policy reaction process so that there is less accountability of the central bank to meet its nominal anchor objectives, but it results in expectations dynamics that destabilize output and inflation. Again consider the situation in which the positive employment report leads to expectations that inflation will be higher than previously expected. With time-dependent forward guidance, the projected policy path does not change, but expected inflation rises. This means that the expected path of future *real* interest rates, policy interest rates minus expected inflation, now declines. The

effect of the positive employment report shock is then an effective easing of monetary policy, the opposite to the monetary policy response that would stabilize output and inflation.

This feature of time-dependent forward guidance is exactly the same problem created by the zero lower bound for the policy rate, as discussed in Eggertson and Woodford (2003). They point out that when there is a negative aggregate demand shock and the policy rate is at the zero lower bound, then a negative aggregate demand shock leads to a decline in expected inflation and therefore a rise in real interest rates, which further weakens aggregate demand. Negative aggregate demand shocks when the zero lower bound is binding therefore can lead to prolonged economic downturns. Time-dependent forward guidance creates a similar problem because, just as occurs when the policy rate is at the zero lower bound, a negative aggregate demand shock leaves the projected future path of the policy rate unchanged, so that real interest rates rise, thereby propagating the negative aggregate demand shock further.

Does empirical evidence support the theory that time-based forward guidance leads to bad expectation dynamics because it leads to interest rates becoming insensitive to macroeconomic news? Feroli et. al. (2017) find that the answer is yes. Using the methodology developed by Swanson and Williams (2014), they evaluate how responsive interest rates were to economics news during periods when the Federal Reserve used time-based forward guidance, data-based forward guidance or no forward guidance at all. Chart 3.3 from Feroli et. al. (2017) reproduced as Figure 2 below shows the sensitivity of Treasury bond rates to macroeconomic news when there is no forward guidance, data-based forward guidance, or time-based forward guidance. As the figure shows, time-based forward guidance is associated with lower sensitivity of interest rates to macroeconomic news at all of the maturities they examine. Time-based forward guidance not only leads to less sensitivity to macroeconomic news than does data-based forward guidance, but also less than when there is no forward guidance at all. Indeed, the results in Feroli et. al. (2017) indicate that, even without forward guidance, markets are able to glean some information about the monetary policy reaction process, which does not occur with time-based forward guidance.

Figure 2: Sensitivity of Interest Rates to Economic News and Forward Guidance, 2001-2015



*Source: Feroli et. al. (2016).

Note that the results in Figure 2 are not driven by the zero-lower bound constraint during the post Great-Recession period. Even excluding the zero lower bound period, the sensitivity of interest rates to macroeconomic news is lower during periods in which FOMC communication on forward guidance is more strongly time-dependent.

3.2.2 Why Communicating a Monetary Policy Reaction Process is Not an Instrument Rule

There is an important subtle issue about the benefits of a central bank communicating a monetary policy reaction process. At first glance, a monetary policy reaction process appears to be very similar to instrument rule like the Taylor rule. After all, a Taylor rule is a very simple way of specifying a monetary policy reaction process. So why is data-based forward guidance

very different from adoption of an instrument rule, such as a Taylor rule. The answer is that monetary policy that minimizes output and inflation gaps results in a policy reaction process to change over time, either as monetary policymakers learn more about how the economy works or when the structure of the economy changes. Furthermore, monetary policy that minimizes output and inflation gaps leads to a modification of the policy reaction process when there are unforeseen contingencies that were previously not part of the reaction process. Judgement is also a feature of such as policy as demonstrated by Svensson (2005) and would then be part of a monetary policy reaction process. A Taylor rule, which does not change over time, can therefore be far from a policy that minimizes output and inflation gaps.

Unlike a Taylor rule, data-dependent forward guidance can be consistent with monetary policy that minimizes output and inflation gaps, but this requires that it changes the monetary policy reaction process. Monetary policy that minimizes output and inflation gaps requires that projections of the future policy path not only must be altered when forecasts of the economy change, but also when the central bank has reasons to expect that the model of the economy is changing. Data-dependent forward guidance thus requires substantial communication to explain not only the past policy reaction process, but also any reasons for changes in the reaction process.

To understand this point, consider what data-dependent forward guidance might have looked like when the global financial crisis started in August of 2007. At the time, inflation was rising and the economy was still growing rapidly in the third quarter. The Federal Reserve dramatically deviated from its previous reaction process by aggressively cutting the federal funds rate even before the economy and inflation had turned down. If the Fed had been providing data-based forward guidance, it would have explained the disruption to financial markets required a change in the policy reaction process, with much easier monetary policy in the future in response to financial shocks than had been anticipated earlier. Indeed, such a shift in the policy reaction process is exactly what would have helped stabilize both aggregate output and inflation. If this communication led to the markets understanding that there had been a shift in the policy reaction process, longer-term interest rates would have fallen more rapidly in response to news that the financial disruption was getting worse. This would have helped monetary policy to be even more

expansionary than it otherwise would have been, by helping offset some of the negative shocks to the economy from the ongoing financial crisis.

4 CONCLUDING REMARKS

In this paper, I have argued that the rules versus discretion debate is miscast because monetary policy does not have to choose only between adopting a policy instrument rule or pure discretion, both of which have serious shortcomings. Instrument rules can lead to poor economic outcomes if the model of the economy is not reliable or the structure of the economy is unstable, or because rules cannot foresee every contingency or allow judgement. Pure discretion can lead to policies that destabilize inflation and output and is subject to the time-inconsistency problem where there monetary policy reneges on the long-run plan that minimizes output and inflation gaps. Instead of making the stark choice between an instrument rule and pure discretion, another choice is to constrain discretion and make it more rule-like. But how can discretion be constrained and be accountable to be rule-like? I argue that monetary policy discretion can be made more rule-like by 1) adopting a nominal anchor such as an inflation target, a monetary policy strategy that has proved to be very successful at stabilizing inflation and output in recent years, and 2) communication of a monetary policy reaction process, especially through data-based forward guidance, in which the monetary policy authorities describe how the future policy path will change as economic circumstances change.

REFERENCES

Barsky, R., A. Justiniano, and L. Melosi, (2014), "The Natural Rate of Interest and Its Usefulness for Monetary Policy," American Economic Review: Papers & Proceedings 104(5): 37-43.

Barro, R.J. and D.B. Gordon (1983), "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12, 1: 101-121.

Bernanke, B.S. (2010), "Monetary Policy and the Housing Bubble", speech given at the annual meeting of the American Economic Association, Atlanta, Georgia, 3 January 2010, available at http://www.federalreserve.gov

Bernanke, B.S., and F.S. Mishkin, (1997), "Inflation Targeting: A New Framework for Monetary Policy," *Journal of Economic Perspectives,* Vol 11, #2 (Spring): 97-116.

Bernnake, B.S., Laubach, T., F.S. Mishkin, and A. Posen (1999), *Inflation Targeting: Lessons from the International Experience* (Princeton, N.J., Princeton University Press).

Brainard, L. (2015), "Normalizing Monetary Policy When the Neutral Interest Rate is Low," http://www.federalreserve.gov/newsevents/speech/brainard20151201a.htm.

Calvo, G.A., (1978), "On the Time-Inconsistency of Optimal Policy in a Monetary Economy," *Econometrica,* 46, 6: 1411-1428.

Curdia, V., A. Ferrero, G.C. Ng, and A. Tambalotti, 2014, "Has U.S. Monetary Policy Tracked the Efficient Interest Rate?", FRB San Francisco Working Paper 2014-12.

Eggertsson, G. B. and Woodford, M. (2003), "The Zero Bound on Interest Rates and Optimal Monetary Policy", *Brookings Papers on Economic Activity*, Vol. 2003, No 1, pp. 139-211.

Eggertsson, G.B. and Woodford, M. (2004), "Policy Options in a Liquidity Trap", *American Economic Review*, 94(2), pp. 76-79.

Feroli, M., Greenlaw, D., Hooper, P., Mishkin, F.S. and A. Sufi (2017), "Language After Liftoff: Fed Communication Away from the Zero Lower Bound," *Research in Economics*, volume 71, issue 3, September, pp. 452-490.

Friedman, M. and A. Schwartz (1963), *A Monetary History of the United States, 1867–1960* (Princeton, NJ: Princeton University Press).

Gürkaynak, Refet S., Andrew Levin, and Eric T. Swanson (2010). "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-Term Bond Yields in the U.S., U.K., and Sweden." *Journal of the European Economic Association*, 8:1208-1242.

Hamilton, J.D., Harris, E.S., Hatzius, J. and K.D. West, (2015) "The Equilibrium Real Funds Rate: Past, Present and Future," U.S. Monetary Policy Forum, February 2015, https://research.chicagobooth.edu/-/media/research/igm/docs/2015-usmpf.pdf?la=en&hash=E5C27E7E0FA17AA169B268EBC157039DC3D662C8

Kydland, F.E and E.C. Prescott, (1977), "Rules Rather Than Discretion: The Inconsistency of Optimal Plans, *Journal of Political Economy,* 85, 3: 473-492.

Mishkin, F.S., (1999) "International Experiences with Different Monetary Policy Regimes," *Journal of Monetary Economics,* Vol 43, #3: 579-606.

Mishkin, F.S. and K. Schmidt-Hebbel, "One Decade of Inflation Targeting in the World: What Do We Know and What Do We Need to Know?" in Norman Loayza and Raimundo Soto, eds., *Inflation Targeting: Design, Performance, Challenges* (Central Bank of Chile: Santiago 2002): 171-219.

Mishkin, F.S. (2008a), "Monetary Policy Flexibility, Risk Management, and Financial Disruptions," speech delivered at the Federal Reserve Bank of New York, January 11, 2008. https://www.federalreserve.gov/newsevents/speech/mishkin20080111a.htm

Mishkin, F.S. (2008b), "Whither Federal Reserve Communication", speech delivered at the Peterson Institute for International Economics, Washington, DC, 28 July, available at https://www.federalreserve.gov/newsevents/speech/mishkin20080728a.htm.

Mishkin, F.S. (2011), "Monetary Policy Strategy: Lessons From the Crisis," in Jarocinski, M, Smets, F., and C. Thimann, eds. *Moneary Policy Revisited: Lessons from the Crisis,* Sixth ECB Central Banking Conference (Frankfort, European Central Bank): 67-118.

Mishkin, F.S.(2016), "2% Forever? Rethinking the Inflation Target" in Oesterreichische Nationalbank, *Central Banking in Times of Change:* 92-101.

Mishkin, F.S. and E. White, (2016), "Unprecedented Action: The Federal Reserve's Response to the Global Financial Crisis in Historical Perspective," in M. Wynne and M. Bordo, eds., *The Federal Reserve System's Role in the Global Economy: An Historical Perspective* (Cambridge, U.K., Cambridge University Press): 220-258.

Mishkin, F.S. and N. Westelius, "Inflation Band Targeting and Optimal Inflation Contracts," *Journal of Money, Credit and Banking* Volume 40, No. 4 (June 2008), pp. 557-582.

Meltzer, A.H. (2004), *A History of the Federal Reserve, Volume 1: 1913-1951* (Chicago, University of Chicago Press).

Meltzer, A.H. (2014), *A History of the Federal Reserve, Volume 2: 1951-1969* (Chicago, University of Chicago Press)

Rogoff, K. (1985) "The Optimal Degree of Commitment to an Intermediate Monetary Target," Quarterly Journal of Economics 100 1169-1189.

Staiger, D, J.H. Stock and M.W. Watson, (1997), "The NAIRU, Unemployment and Monetary Policy" *Journal of Economic Perspectives,* 11, 1: 33-49.

Svensson, L. (2003) "What is Wrong with Taylor Rules? Using Judgment in Monetary Policy Through Targeting Rules," *Journal of Economic Literature* 41: 426-477.

Svensson, L. (2005), "Monetary Policy with Judgment: Forecast Targeting," *International Journal of Central Banking* 1(1) (2005), 1-54.

Swanson, E.T and J. C. Williams (2014), "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, 104, 10: 3154-3185.

Tarullo, D. (2015) Remarks on CNBC's "Power Lunch Today," October 13, http://www.cnbc.com/2015/10/13/.

Taylor, J.B., (1993), "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39: 195-214.

Taylor, J. B. (2007), "Housing and Monetary Policy", *Housing, Housing Finance and Monetary Policy*, Federal Reserve Bank of Kansas City, Kansas City, pp. 463-476.

Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, Princeton.

Woodward, R. (2000), *Greenspan's Fed and the American Boom*, (New York, Simon and Schuster, 2000)

Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Taylor Rules

Athanasios Orphanides

2007-18

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Taylor Rules

Athanasios Orphanides*

Board of Governors of the Federal Reserve System

January 2007

Abstract

Taylor rules are simple monetary policy rules that prescribe how a central bank should adjust its interest rate policy instrument in a systematic manner in response to developments in inflation and macroeconomic activity. This paper reviews the development and characteristics of Taylor rules in relation to alternative monetary policy guides and discusses their role for positive and normative monetary policy analysis.

KEYWORDS: Monetary policy, simple rules.

JEL Classification: E52.

Correspondence: Division of Monetary Affairs, Federal Reserve Board, Washington, D.C. 20551, USA. Tel.: (202) 452-2654. E-mail: Athanasios.Orphanides@frb.gov.

*Prepared for the second edition of *The New Palgrave: A Dictionary of Economics*

The opinions expressed are those of the author and do not necessarily reflect views of the Board of Governors of the Federal Reserve System.

1 Introduction

Taylor rules are simple monetary policy rules that prescribe how a central bank should adjust its interest rate policy instrument in a systematic manner in response to developments in inflation and macroeconomic activity. They provide a useful framework for the analysis of historical policy and for the econometric evaluation of specific alternative strategies that a central bank can use as the basis for its interest rate decisions.

A perennial question in monetary economics has been how the monetary authority should formulate and implement its policy decisions so as to best foster ultimate policy objectives such as price stability and full employment over time. It is widely accepted that well designed monetary policy can counteract macroeconomic disturbances and dampen cyclical fluctuations in prices and employment, thereby improving overall economic stability and welfare. In principle, when economic growth unexpectedly weakens below the economy's potential, accommodative monetary policy can stimulate aggregate demand and restore full employment. Likewise, when inflationary pressures develop, monetary restriction can restore the central bank's price stability objective. In practice, however, given the limited knowledge that economists have about the macroeconomy—for example, about macroeconomic dynamics, about the monetary transmission mechanism, and even about the measurement of fundamental concepts such as the natural rates of output, employment and interest—there is substantial disagreement about the scope of stabilization policy and about policy design.

One approach is to decide upon what seems to be the best policy on a period-by-period basis, without appeal to any specific policy guide. A seeming advantage of this approach is that it gives policymakers the discretion to use their judgment period by period. However, a basic tenet of modern research is that systematic policy—that is,

policy based on a contingency plan or policy rule—has important advantages over a purely discretionary policy approach. By committing to follow a rule, policymakers can avoid the inefficiency associated with the time-inconsistency problem that arises when policy is formulated in a discretionary manner. Following a rule allows policymakers to communicate and explain their policy actions more effectively. Policy based on a well-understood rule enhances the accountability of the central bank and improves the credibility of future policy actions. Also, by making future policy decisions more predictable, rule-based policy facilitates forecasting by financial market participants, businesses, and households, thereby reducing uncertainty.

Various proposals for monetary policy rules have been made over time, and a vast literature continues to examine the relative advantages and drawbacks of alternatives in abstract theoretical terms, in the context of empirical macroeconometric models, and in terms of the practical experience accumulated from past policy practice. To appreciate the appeal and limitations of Taylor rules, it is useful to relate their development to other proposals for systematic monetary policy.

2 Development of Monetary Policy Rules

Some proposals suggest postulating a rule in terms of the main objectives of monetary policy, for example "maintain economic stability" or "maintain a constant aggregate price level." (See Simons, 1936, for early arguments favoring price-level targeting over discretionary policy.) One important practical difficulty with these proposals, however, is that the concepts involved are not under the control of the central bank and thus the proposals are not operational. In essence, these proposals fail to draw a clear distinction between the objectives of monetary policy and the policy instruments that are at least under the approximate control of the central bank. As a result, the suggested rules are only implicit in nature and are difficult to monitor and to distinguish from discretionary policy in a meaningful manner.

To be useful in practice, policy rules must be simple and transparent to communicate, implement, and verify. This requires a clear choice of what should serve as the policy instrument—for example the money supply, m, or the short-term interest rate, i—and clear guidance as to how any other information necessary to implement the rule—for instance, recent readings or forecasts of inflation and economy activity—should be used to adjust the policy instrument.

Perhaps the simplest example of a policy rule is the proposal that the central bank maintain a constant rate of growth of the money supply—Milton Friedman's k-percent rule (Friedman, 1960). The rule draws on the equation of exchange expressed in growth rates:

$$\Delta m + \Delta v = \pi + \Delta q \tag{1}$$

where $\pi \equiv \Delta p$ is the rate of inflation and p, m, v, and q are (the logarithms of), respectively, the price level, money stock, money velocity and real output. Selecting the constant growth of money, k, to correspond to the sum of a desired inflation target, π^*, and the economy's potential growth rate, Δq^*, and adjusting for any secular trend in the velocity of money, Δv^*, suggests a simple rule that can achieve, on average, the desired inflation target, π^*:

$$\Delta m = \pi^* + \Delta q^* - \Delta v^*. \tag{2}$$

Further, if the velocity of money were fairly stable this simple rule would also yield a high degree of economic stability. An early illustration of this rule appeared in 1935 in the work of Carl Snyder, a statistician at the Federal Reserve Bank of New York. After estimating that the normal rate of growth of trade in the United States was about 4 percent per year at the time and observing that the velocity of money was stable, Snyder argued that "the highest attainable degree of general industrial and economic stability will be gained by an expansion of currency and credit ... at this rate [4 percent]" (Snyder, 1935, p. 198). During the 1960s and early 1970s, Milton Friedman's recommendation that the Federal Reserve control the rate of money

growth to equal 4 percent per year was similarly based on the assumption that potential output growth in the Unites States roughly equaled 4 percent—the prevailing estimate at that time.

Another way to interpret this policy rule is in terms of the growth of nominal income, $\Delta x = \pi + \Delta q$. With the economy's natural growth of nominal income defined as the sum of the natural growth rate of output and the central bank's inflation objective, $\Delta x^* = \pi^* + \Delta q^*$, a rule for constant money growth can be seen as targeting this natural growth rate. An advantage of a constant money growth rule is that very little information is required to implement it. If velocity does not exhibit a secular trend, the only required element for calibrating the rule is the economy's natural growth of output. In addition, while the calibration of this rule does not rest on the specification of any particular model, the rule is remarkably stable across alternative models of the economy. In this sense, the policy of maintaining a constant growth rate of money is arguably the ultimate example of a rule that is robust to possible model misspecification.

Simple modifications allowing for some automatic response of money growth to economic developments have also been proposed as simple rules that could deliver improved macroeconomic performance. (See, e.g. Cooper and Fischer, 1972). Among the simplest such alternatives is the rule associated with Bennett McCallum(1988, 1993):

$$\Delta m = \Delta x^* - \Delta v^* - \phi_{\Delta x}(\Delta x - \Delta x^*). \tag{3}$$

McCallum showed that if a rule such as this (for example with $\phi_{\Delta x} = 0.5$) had been followed, the performance of the U.S. economy likely would have been considerably better than actual performance, especially during the 1930s and 1970s—the two periods of the worst monetary policy mistakes in the history of the Federal Reserve.

A factor that complicates the use of the money stock as a policy instrument is the potential for instability in the demand for money either due to temporary

disturbances or due to persistent changes resulting from financial innovation. In part for this reason, central banks generally prefer to adjust monetary policy using an interest rate instrument.

A policy rule quite as simple as Friedman's k-percent rule cannot be formulated with an interest rate instrument. As early as Wicksell's (1898) monumental treatise on *Interest and Prices*, it was recognized that attempting to peg the short-term nominal interest rate at a fixed value does not constitute a stable policy rule. (Indeed, this was one reason why Friedman, 1968, and others expressed a preference for rules with money as the policy instrument.) Wicksell argued that the central bank should aim to maintain price stability, which in theory could be achieved if the interest rate were always equal to the economy's natural rate of interest, r^*. Recognizing that the natural rate of interest is merely an abstract, unobservable concept, however, he noted: "This does not mean that the bank ought actually to *ascertain* the natural rate before fixing their own rates of interest. That would, of course, be impracticable, and would also be quite unnecessary." Rather, Wicksell pointed out that a simple policy rule that responded systematically to prices would be sufficient to achieve satisfactory, though imperfect, stability: "If prices rise, the rate of interest is to be raised; and if prices fall, the rate of interest is to be lowered; and the rate of interest is henceforth to be maintained at its new level until a further movement in prices calls for a further change in one direction or the other." (Wicksell, 1898 [1936], p. 189, emphasis in the original). In algebraic terms, Wicksell proposed what is arguably the simplest reactive monetary rule with an interest rate instrument:

$$\Delta i = \theta \pi. \tag{4}$$

Wicksell's simple interest rate rule did not attract much attention in policy discussions, perhaps because of its exclusive focus on price stability and lack of explicit reference to developments in real economic activity.

3 The Classic Taylor Rule and its Generalizations

The policy rules that are commonly referred to as Taylor rules are simple reactive rules that adjust the interest rate policy instrument in response to developments in both inflation and economic activity. An important advance in the development of these rules can be identified with the policy regime evaluation project reported in a volume published by the Brookings Institution (Bryant, Hooper and Mann, 1993). The objective of the project was to identify simple reactive interest rate rules that would deliver satisfactory economic performance for price stability and economic stability across a range of competing estimated models. The Brookings project examined rules that set deviations of the short-term nominal interest rate, i, from some baseline path, i^*, in proportion to deviations of target variables z, from their targets, z^*:

$$i - i^* = \theta(z - z^*). \tag{5}$$

The collective findings pointed to two alternatives as the most promising in delivering satisfactory economic performance across models. One targeted nominal income, while the other targeted inflation and real output:

$$i - i^* = \theta_\pi(\pi - \pi^*) + \theta_q(q - q^*) \tag{6}$$

The potential usefulness of this particular rule as a benchmark for setting monetary policy was further highlighted in the celebrated contribution by John B. Taylor (1993) at the Fall 1992 Carnegie-Rochester Conference on Public Policy. Taylor developed a "hypothetical but representative policy rule" (p. 214) by using the sum of the equilibrium or natural rate of interest, r^*, and inflation, π, for i^* and setting the inflation target and equilibrium real interest equal to two and the response parameters to one half. The result was what became known as the classic Taylor rule:

$$i = 2 + \pi + \frac{1}{2}(\pi - 2) + \frac{1}{2}(q - q^*). \tag{7}$$

Taylor noted that if one used the deviation of real quarterly output from a linear trend to measure the output gap, $(q - q^*)$, and the year-over-year rate of change of the output deflator to measure inflation, π, this parameterization appeared to describe Federal Reserve behavior well in the late 1980s and early 1990s.

The confluence of the econometric evaluation evidence supporting the stabilization properties of this rule and its usefulness for understanding historical monetary policy in a period generally accepted as having good policy performance generated tremendous interest, and numerous central banks began to monitor this policy rule or related variants to provide guidance in policy decisions. These developments also greatly influenced monetary policy research and teaching. By linking interest rate decisions directly to inflation and economic activity, Taylor rules offered a convenient tool for studying monetary policy while abstracting from a detailed analysis of the demand and supply of money. This allowed the development of simpler models (see the survey in Clarida, Gali, and Gertler, 1999 and papers in Taylor, 1999) and the replacement of the "LM curve" with a Taylor rule in treatments of the Hicksian IS-LM apparatus. (It should be noted, however, that this abstraction is overly simplistic when the short-term interest rate approaches zero. At the zero bound, the stance of monetary policy can no longer be measured or communicated with a short-term interest rate instrument; see, for example, Orphanides and Wieland 2000). Subsequent research (see Orphanides, 2003b, for a survey) suggested that a generalized form of Taylor's classic rule could provide a useful common basis both for econometric policy evaluation across diverse families of models and for historical monetary policy analysis over a broad range of experience:

$$i = (1 - \theta_i)(r^* + \pi^*) + \theta_i i_{-1} + \theta_\pi(\pi - \pi^*) + \theta_q(q - q^*) + \theta_{\Delta q}(\Delta q - \Delta q^*) \qquad (8)$$

The generalized Taylor rule (8) nests rule (6) as a special case but introduces two additional elements. First, it allows for inertial behavior in setting interest rates, $\theta_i > 0$, which proves particularly important for policy analysis in models with strong

expectational channels (Woodford, 2003). Second, it allows the policy response to developments in economic activity to take two forms: a response to the level of the output gap, $(q - q^*)$, or its difference, which can also be restated as a response to the difference between output growth and its potential, $(\Delta q - \Delta q^*)$. The generalized Taylor rule also nests another simplification of special interest, $\theta_i = 1$ and $\theta_q = 0$, which yields a family of difference rules similar to Wicksell's original proposal:

$$\Delta i = \theta_\pi(\pi - \pi^*) + \theta_{\Delta q}(\Delta q - \Delta q^*). \tag{9}$$

These difference rules are also of interest because, like money-growth rules, their implementation does not require estimates of the natural rate of interest or the level of potential output (and the output gap) but only of the growth rate of potential output. Indeed, these rules may be viewed as a reformulation of money-growth rules in terms of an interest rate instrument. To see the relationship of (9) to money growth targeting note that by substituting the money growth in rule (3) into the equation of exchange, that rule can be stated in terms of the velocity of money:

$$\Delta v - \Delta v^* = (1 + \phi_{\Delta x})(\Delta x - \Delta x^*). \tag{10}$$

To reformulate this strategy in terms of an interest rate rule, consider the simplest formulation of money demand as a (log-) linear relationship between velocity deviations from its equilibrium and the rate of interest. In difference form this is

$$\Delta v - \Delta v^* = a\Delta i + e, \tag{11}$$

where $a > 0$ and e summarizes short-run money demand dynamics and temporary velocity disturbances. An interest-rate-based strategy that avoids the short-run velocity fluctuations, e, may be obtained by substituting the remaining part of (11) into (10). This yields

$$\Delta i = \theta((\pi - \pi^*) + (\Delta q - \Delta q^*)) \tag{12}$$

for some $\theta > 0$, which, as can be readily seen, has the same form as rule (9).

In light of this flexibility in nesting a wide range of alternative monetary policy strategies and the relative simplicity of the form (8), Taylor rules have been used to discuss a variety of policy regimes, from money growth targeting (see, e.g., Clarida and Gertler, 1997) to inflation targeting (see, e.g. Orphanides and Williams, 2007).

4 Operational Implementation

A crucial element for the design and operational implementation of a Taylor rule is the detailed description of its inputs. This requires specificity regarding the measures of inflation and economic activity that the policy rule should respond to, whether forecasts or recent outcomes of these variables are to be employed, and the source of these data or forecasts. In addition, the source of information and updating procedures regarding the unobservable concepts required for implementing the rule must be stipulated. Specificity in these dimensions is essential for practical analysis because there is often a multitude of competing alternatives and a lack of consensus about the appropriate concepts and sources of information that ought to be used for policy analysis. This situation is particularly vexing in regard to the treatment of unobservable concepts, such as the output gap. Unfortunately, econometric policy evaluation exercises suggest that inferences regarding the performance of a particular Taylor rule often depend sensitively on assumptions regarding the availability and reliability of these inputs. Differences in underlying assumptions complicate comparisons across studies and often explain differences in reported findings.

An illustrative example of this sensitivity relates to improper treatment of information regarding the current state of the economy. A common pitfall in theoretical policy evaluation exercises is to assume that the current state of the economy, e.g. the current output gap, can be perfectly observed. Under this assumption, a Taylor rule with a vigorous response to the output gap is often recommended as "optimal" in model-based policy evaluations. However, naive adoption of such recommenda-

tions would be counterproductive. Available real-time estimates of the output gap are imperfect, and historical experience suggests that the mismeasurement is often substantial. Under these circumstances, better stabilization outcomes would result if policy did not respond to the output gap at all or if it responded to output growth instead (Orphanides, 2003a). If the natural rate of interest is also unknown and its real-time estimates are subject to significant mismeasurement, the difference variant of the Taylor rule, (9), proves considerably more robust than the Brookings variant, (6), reversing the ranking of the two alternatives that is implied under perfect knowledge (Orphanides and Williams, 2002).

Another example of such sensitivity relates to the use of forecasts in the Taylor rule. Because of lags in the monetary policy transmission mechanism, preemptive policy reaction is generally recommended, especially with respect to inflation. But inferences regarding the performance of forecast-based policy are sensitive to the quality of the forecasts. In some models, Taylor rules responding to several-quarters-ahead forecasts of inflation appear more promising for stabilization than rules focusing only on near-term conditions. However, this conclusion is not robust and is overturned once the potential unreliability of longer-term forecasts due to model misspecification is factored into the analysis (Levin, Wieland, and Williams, 2003).

As already noted, Taylor rules have proven valuable for historical policy analysis. Following Taylor (1993), numerous authors have examined historical monetary policy in the United States using either calibrated or estimated versions of Taylor rules (8). Studying the characteristics of policy in periods associated with good or bad economic performance helps identify aspects of policy that may be associated with such differences in performance. A complicating factor is the need for real-time data and forecasts for proper inference (Orphanides, 2001). The pitfall of using ex post revised data and retrospective estimates of unobserved concepts in estimating Taylor rules is not uncommon. However, interpretations of historical policy based on infor-

mation that was unavailable to policymakers when policy decisions were made is of questionable value. Policy prescriptions from a fixed rule are distorted as the inputs to the rule are revised from those originally available to policymakers, and therefore counterfactual comparisons of alternative policy rules can be misleading when they are based on revised data.

5 Concluding Remarks

Despite these challenges, some useful elements of policy design emerge from historical analysis of Taylor rules. First, and arguably most important, good stabilization performance is associated with a strong reaction to inflation. Second, good performance is associated with policy rules that exhibit considerable inertia. Third, a strong reaction to mismeasured output gaps has historically proven counterproductive. Fourth, successful policy could still usefully incorporate information from real economic activity by focusing on the growth rate of the economy. To be sure, such broad principles provide insufficient guidance for identifying the precise policy rule that might be ideal in a specific context. But this is not the objective of policy design with Taylor rules. Rather, the goal is the identification of simple guides that are robust to misspecification and other sources of error experienced over history.

In summary, Taylor rules offer a simple and transparent framework with which to organize the discussion of systematic monetary policy. Their adoption as a tool for policy discussions has facilitated a welcome convergence between monetary policy practice and monetary policy research and proved an important advance for both positive and normative analysis.

References

Bryant, R.C., Hooper, P., Mann, C. (Eds), 1993. *Evaluating Policy Regimes: New Research in Empirical Macroeconomics.* Washington DC: Brookings.

Clarida, R., Gali, J., Gertler, M., 1999. The Science of Monetary Policy. *Journal of Economic Literature* 37(4), 1661-1707, December.

Clarida, R., and Gertler, M., 1997. How the Bundesbank Conducts Monetary Policy. In: Romer, C. and Romer, D., (Eds), *Reducing Inflation: Motivation and Strategy.* Chicago: University of Chicago Press.

Cooper, J.P., and Fischer, S., 1972. Simulations of Monetary Rules in the FRB-MIT-Penn Model. *Journal of Money, Credit and Banking* 4(2), 384-396. May.

Friedman, Milton, 1960. *A Program for Monetary Stability.* New York: Fordham University Press.

Friedman, Milton, 1968. The Role of Monetary Policy. *American Economic Review* 58(1), March, 1-17.

Levin, A., Wieland, V., Williams, J.C., 2003. The Performance of Forecast-based Monetary Policy Rules under Model Uncertainty. *American Economic Review* 93(3), 622-645, June.

McCallum, B.T., 1988. Robustness Properties of a Rule for Monetary Policy. *Carnegie-Rochester Conference Series on Public Policy* 29, Autumn, 173-203.

McCallum, B.T., 1993. Specification and Analysis of a Monetary Policy Rule for Japan. *Bank of Japan Monetary and Economic Studies*, November, 1-45.

Orphanides, A., 2001. Monetary Policy Rules Based on Real-Time Data. *American Economic Review* 91(4), 964-85, September.

Orphanides, A., 2003a. Monetary Policy Evaluation with Noisy Information. *Journal of Monetary Economics* 50(3), 605-631, April.

Orphanides, A., 2003b. Historical Monetary Policy Analysis and the Taylor Rule. *Journal of Monetary Economics* 50(5), 983-1022, July.

Orphanides, A., Wieland, V., 2000. Efficient Monetary Policy Design Near Price Stability. *Journal of the Japanese and International Economies* 14, 327-365.

Orphanides, A., Williams, J.C., 2002. Robust Monetary Policy Rules with Unknown Natural Rates *Brookings Papers on Economic Activity* 2:2002, 63-145.

Orphanides, A., Williams, J.C., 2007. Inflation Targeting under Imperfect Knowledge. in Mishkin, F. and Schmidt-Hebbel, K., (Eds) *Monetary Policy under Inflation Targeting,* Santiago: Central Bank of Chile (forthcoming).

Simons, H.C. 1936. Rules vs Authorities in Monetary Policy. *Journal of Political Economy* 44(1), 1-30.

Snyder, Carl, 1935. The Problem of Monetary and Economic Stability. *Quarterly Journal of Economics* 49(2) 173-205, February.

Taylor, J.B., 1993. Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy* 39, December, 195-214.

Taylor, J.B. (Ed), 1999. *Monetary Policy Rules.* Chicago: University of Chicago.

Wicksell, Knut, 1898. *Interest and Prices.* (1936 translation from the German by R. F. Kahn.) London: Macmillan.

Woodford, Michael, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

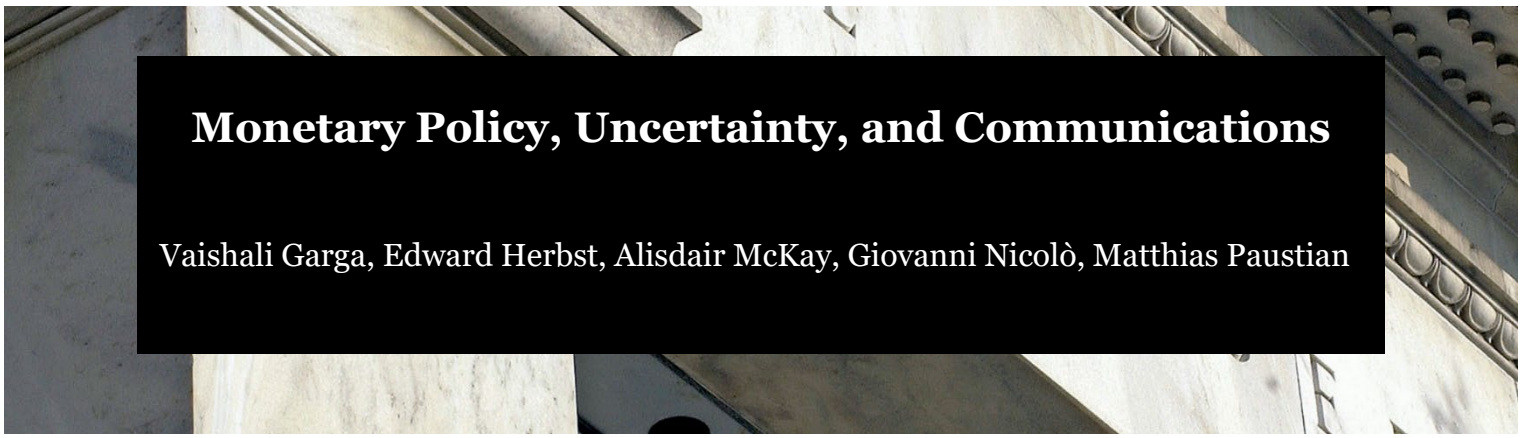
Monetary Policy, Uncertainty, and Communications

Vaishali Garga, Edward Herbst, Alisdair McKay, Giovanni Nicolò, Matthias Paustian

2025-074

Please cite this paper as:
Garga, Vaishali, Edward Herbst, Alisdair McKay, Giovanni Nicolò, and Matthias Paustian (2025). "Monetary Policy, Uncertainty, and Communications," Finance and Economics Discussion Series 2025-074. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.074.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.



Monetary Policy, Uncertainty, and Communications

Vaishali Garga, Edward Herbst, Alisdair McKay, Giovanni Nicolò, Matthias Paustian

The analysis in this paper was presented to the Federal Open Market Committee as background for its discussion of the Federal Reserve's 2025 review of its monetary policy strategy, tools, and communications.

Abstract: We review the design and communication of monetary policy strategies that take into account risks and uncertainty. A key element in a robust monetary strategy is the concept of risk management, which is the weighing of key risks when setting policy. When risks to the outlook are balanced, the baseline outlook may be sufficient to guide policy decisions. However, risk-management considerations become important when risks are asymmetric. We discuss how robust simple interest rate rules and optimal control policy can incorporate risk-management considerations into the design of a monetary policy strategy. Alternative scenarios can illustrate salient risks and how monetary policy might respond if those risks were to materialize. However, using alternative scenarios in policy deliberations and communications requires important implementation choices.

JEL Classification: E31, E32, E52, E58.

Keywords: Uncertainty, risk management, robust monetary policy strategies, scenario analysis, monetary policy communication.

Note: Authors' affiliations are Board of Governors of the Federal Reserve System (Herbst, Nicolò, and Paustian), Federal Reserve Bank of Boston (Garga), and Federal Reserve Bank of Minneapolis (McKay), respectively. The views expressed are those of the authors and do not necessarily reflect the views of anyone else affiliated with the Federal Reserve System. The authors benefited from the comments and suggestions of Jonas Fisher, Christopher Gust, Benjamin Johannsen, Michael Kiley, Sylvain Leduc, Kurt Lewis, David López-Salido, Edward Nelson, Anna Paulson, Trevor Reeve, Robert Tetlow, Stacey Tevlin, Paula Tkac, and Rebecca Zarutskie. The authors would like to thank James Hebden for designing and implementing the simulations, Karen Page for excellent research assistance, and David Jenkins for editorial assistance.

1. Introduction and overview

This paper analyzes the design and communication of monetary policy strategies, taking into account risks and uncertainty.[1] As outlined in the companion paper "Accounting for Uncertainty and Risks in Monetary Policy" by Bauer and others (2025), policymakers face many types of risks and uncertainty, including regarding the state and structure of the economy and the formation of expectations. This paper identifies key elements of a monetary policy strategy that is robust in the face of these risks and uncertainty. A robust monetary policy strategy shares key elements of any well-designed monetary policy strategy: It is systematic yet flexible, firmly anchors longer-term inflation expectations, and clearly communicates the rationale for the policy decision. In addition, it incorporates risk-management considerations through the weighing of key risks in setting monetary policy.

In the context of model-based analysis, two benchmark approaches, robust simple interest rate rules and optimal control policy, can guide how risk-management considerations are incorporated into the design of a monetary policy strategy. Robust simple rules incorporate risk-management considerations indirectly because they perform well in a broad array of economic environments instead of being narrowly tailored to particular settings. However, robust simple rules often lack the flexibility to appropriately respond to a specific evolution of the economic outlook. By contrast, optimal control policy can explicitly incorporate risk-management considerations: As economic developments reshape the economic outlook and the surrounding risks, the prescription of optimal control adjusts accordingly. However, the prescription is closely tied to modeling assumptions and can be complex, and contemplating the full spectrum of risks is challenging.

Although optimally managing all risks is challenging, alternative scenarios can be a useful tool with which to illustrate salient risks to the baseline outlook and how monetary policy might respond if those risks were to materialize. However, using alternative scenarios in policy deliberations and communications requires important implementation choices. These choices include what forecast to show as the baseline scenario, what risks to highlight, and what monetary policy assumptions to use in constructing the alternative scenarios.

2. Key elements of a robust monetary policy strategy

A robust monetary policy strategy aims to stabilize economic activity and inflation across a wide range of highly uncertain developments. Over the past several decades, the experiences of major central banks and extensive economic research have informed the development of key principles for any well-designed monetary policy strategy.[2] These principles can also help guide

[1] In this paper, we use these terms as defined in the companion paper by Bauer and others (2025).

[2] A series of short notes on "Monetary Principles and Practice" are available on the Board's website at https://www.federalreserve.gov/monetarypolicy/monetary-policy-principles-and-practice.htm. See also Williams (2025) for a related discussion.

policymaking in the face of inherent uncertainty about the economic outlook. They can foster public understanding of monetary policy and help the public to anticipate policy actions. A well-designed policy that takes into account uncertainty should include the following elements.

- ***Systematic yet flexible.*** A robust policy strategy needs to trade off being systematic with being flexible. The commitment to achieving clear objectives via a strategy that is predictable and well understood by the public fosters credibility, which in turn imparts greater influence on expectations of future policy and outcomes, improving policy transmission. There are consequently major benefits to monetary policy that is systematic. However, flexibility is also important. Uncertain economic developments call for policy to be flexible enough to respond to unusual movements in economic conditions and a broad range of scenarios.
- ***Forward-looking and data dependent.*** As changes in monetary policy affect the economy with a delay, policy is necessarily forward-looking and adjusted with an eye to the expected evolution of the economy. At the same time, policy decisions need to depend on the incoming data and their implications for the outlook and the balance of risks.[3]
- ***Firm nominal anchor.*** A commitment to a well-understood inflation target fosters credibility and anchors longer-term inflation expectations, which helps stabilize inflation without requiring excessive movements in economic activity.
- ***Follows the Taylor principle.*** The nominal policy rate should adjust, over time, more than one-for-one in response to a persistent movement in inflation. Broadly speaking, real interest rates lean against persistent fluctuations in inflation and help stabilize economic activity and anchor longer-term inflation expectations.
- ***Clear communication.*** Clear and transparent communication of policy intentions and of the rationale for policy decisions helps enhance the public's understanding and promotes confidence in the central bank. Transparent communication about risks and uncertainty can also make clear that future policy will be adjusted as needed to achieve longer-run goals. However, care should be exercised so that communication itself does not exacerbate uncertainty or create confusion.
- ***Incorporates risk management.*** A cornerstone of monetary policy under uncertainty is "risk management." Under this approach, policymakers weigh possibilities—for example, that different shocks might hit the economy, or that the structure of the economy could differ from their baseline assessment. The modal outlook is typically a good guide for setting policy if the distribution of risks is symmetric around that mode. However, the mode is not sufficient when risks are asymmetric. The asymmetry of risks can include both the distribution of possible outcomes and the severity of losses associated with those outcomes.

[3] See the companion paper by Bauer and others (2025) for a definition of the concept of balance of risks.

3. Incorporating risk-management considerations into monetary policy analysis

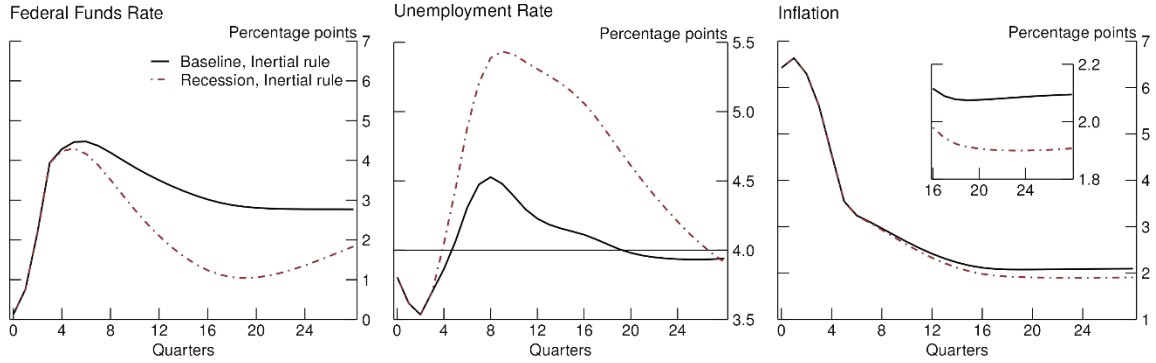
Incorporating risk-management considerations into policymaking can be formalized using the concept of a loss function, which attempts to quantify the welfare costs that arise when the objectives of a monetary policy strategy are not achieved. Typically, loss functions include deviations of inflation from the policymaker's inflation target and the percentage difference in employment from its natural level.[4] Importantly, the loss function includes both current and discounted future costs, because the lags in the transmission of monetary policy imply that the current policy affects the future values of inflation and employment and consequently future losses. In the presence of uncertainty, the dominant framework used to evaluate a policy strategy is *expected loss*, in which costs are weighed according to the assessment of the probability of their realization.[5] A robust policy strategy—one that manages risk well—will result in small expected losses across a range of risks and uncertainty.

[4] Recent work has studied the implications of responding to shortfalls, rather than deviations, of activity from its natural level using simple policy rules (Bundick, Cairó, and Petrosky-Nadeau, 2025) or using a loss function and solving for optimal policy (Kiley, 2024a; Gust, López-Salido, and Meyer, 2017).

In addition, the loss function often also penalizes changes in the policy rate. The inclusion of a desire for interest rate smoothing in an assumed loss function could be motivated in several ways. Such smoothing can improve the predictability of future policy actions, helping to keep the public's inflation expectations well anchored and to influence longer-term interest rates. The desire to smooth variation in the policy rate could also proxy for mitigating financial stability risk stemming from highly volatile interest rates. Moreover, such smoothing implies policy prescriptions that more closely resemble the observed persistence in the actual federal funds rate in the past. Policymakers may put considerable weight on interest rate smoothing as a hedge against model uncertainty or out of concerns for not-perfectly-rational expectations formation (such as learning behavior).

[5] Greenspan (2004) describes such a risk-management approach and how it influenced policy choices in practice. The framework we have described requires the assignment of (potentially subjective) probabilities to events. An alternative approach, known as robust control (or, relatedly, minimax regret), is appropriate for situations in which policymakers are unwilling or unable to assign probabilities to outcomes and seek to achieve the best outcome in the worst-case scenario among those that are considered. The policy prescriptions that emerge from the robust control approach can be sensitive to modeling assumptions about the worst case. In this paper, we focus on the Bayesian approach that evaluates policy strategies based on expected losses, taking as given subjective probabilities of designated events. A related possibility is that policymakers may not even contemplate some possible outcomes (see Kay and King, 2020). Formal analysis of decisionmaking in this setting is challenging. This paper does not specifically engage with "unknown unknowns," but we note that it is important to avoid a false sense of precision that can arise when making decisions assuming only "known unknowns."

Figure 1: Baseline and recession scenario



Source: Authors' calculations.

In order to calculate the expected loss associated with a particular policy strategy, one needs to specify the set of potential outcomes. Many approaches can be used to describe risks, but we focus on the use of alternative scenarios as commonly practiced by central banks and regularly shown in the Tealbook. Figure 1 illustrates a deliberately stylized example of this approach, beginning in a situation of elevated inflation and low unemployment. The figure shows two sets of outcomes: a baseline projection and an alternative scenario, both constructed using the "Inertial rule"—that is, a version of the Taylor (1999) rule that includes a response to the lagged policy rate.[6] The baseline projection (solid black line) represents the modal outlook—that is, the most likely evolution of the economy. Under this projection, the policy rate responds to elevated inflation and the unemployment rate rises modestly above its longer-run level. The recession projection (dash-dotted red line) represents the alternative scenario, in which the economy enters a recession and the policy rate falls more rapidly in response to the higher unemployment rate and somewhat lower inflation. In this example, the recession represents the only risk to the baseline projection, and the risk to unemployment is skewed upwards—that is, the mean unemployment rate is necessarily higher than in the baseline (mode) if there is any chance of the recession occurring.

The concept of expected loss provides a unified framework for evaluating the performance of monetary policy in two benchmark approaches to designing monetary strategies: robust simple interest rate rules and optimal control policy. These two approaches offer complementary risk-management perspectives when designing and communicating monetary policy. We will later use the stylized example in figure 1 to explore the policy prescriptions of

[6] All simulations shown in the paper are constructed using a linearized version of the FRB/US model. In the model, private-sector decisionmakers—such as wage and price setters and financial market participants—form expectations using small-scale statistical (vector autoregressive) models. In the simulations, the "Inertial rule" sets the policy rate according to an inertial Taylor (1999) rule with a coefficient on the lagged policy rate of 0.85 and in which the unemployment rate gap is measured by the deviation of the unemployment rate from an assumed longer-run unemployment rate of 4 percent. For further details about the expectations formation process or the Taylor-type rule, see Hebden and others (2020).

different simple rules and optimal control policies as well as the implications for risk management.

Robust simple interest rate rules

Simple policy rules—which set the short-term policy rate in response to a few key economic variables—provide useful benchmarks for monetary policy deliberations.[7] Several benefits follow from the simplicity of these rules. By design, policy rules are systematic and thus predictable. These rules are relatively easy to communicate and understand, and the relationship of the policy strategy to policy objectives can be made explicit by, for example, directly including the policy objective in the rule. A simple rule also has the advantage that setting policy in accordance with the Taylor principle is typically straightforward.[8] A broad benefit obtains from the principle of parsimony: By focusing on key principles, the rule avoids being narrowly tailored to particular settings or fine-tuned for specific models.

Risk-management considerations can be incorporated into the choice or design of the rule by taking into account its performance in a broad range of situations and environments.[9] A robust policy rule is one that performs well across a broad range of environments. Although economic research has identified some broad features that a robust simple policy rule should possess, such as satisfying the Taylor principle, there is no single simple rule that delivers low expected losses in all settings. Even the design and evaluation of robust simple rules is premised on the set of models and shocks used in the simulations and thus may not be robust to situations or circumstances outside of those explicitly considered.

To illustrate how the performance of simple policy rules can depend on economic circumstances, we build on the alternative simulations used in figure 1. Specifically, we demonstrate how inertial rules—which the literature has found to perform well when private-sector decisionmakers are forward-looking—may lack the flexibility required to respond rapidly to economic downturns. In figure 1, both the baseline scenario and the recession scenario use the Inertial rule. In figure 2, we compare the outcomes (under both projections) when monetary policy instead follows a "Non-inertial rule"—that is, the non-inertial version of the Taylor (1999) rule.

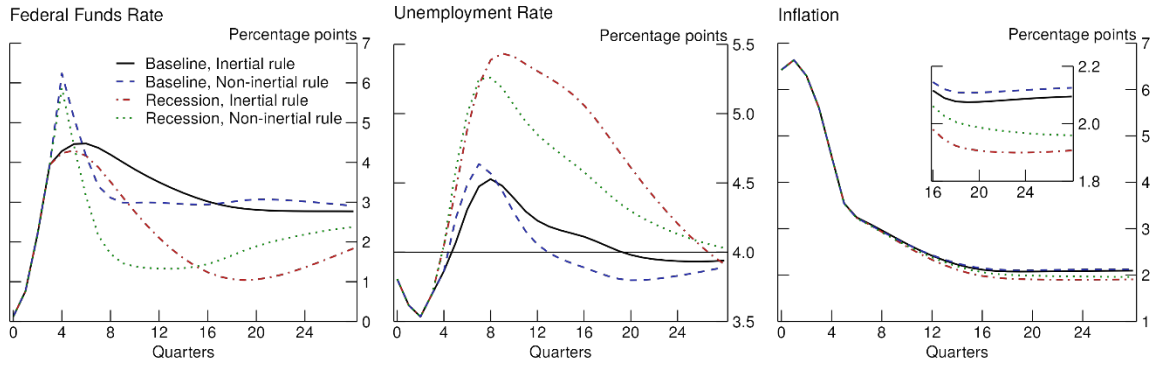
---

[7] Rules that relate the policy rate to a set of macroeconomic indicators are known as instrument rules. These rules are the focus of this subsection. An alternative approach, targeting rules, instead specify a condition to be fulfilled between the target variables for the central bank. Optimal targeting rules are optimality conditions for the setting of monetary policy that are derived from an objective function and an economic model. Targeting rules are implicit in our discussion of optimal control in the next subsection in this paper. See Svensson (2002) and Giannoni and Woodford (2002).

[8] For background on simple policy rules, see McCallum (2002) and Taylor and Williams (2011). A discussion on "Policy Rules and How Policymakers Use Them" is available on the Board's website at https://www.federalreserve.gov/monetarypolicy/policy-rules-and-how-policymakers-use-them.htm.

[9] In principle, this approach can be used to construct optimal simple rules by choosing the coefficients of a given rule to minimize the expected losses in a given simulation study. See, for example, Levin, Wieland, and Williams (1999, 2003) and Orphanides and Williams (2002).

Figure 2: Outcomes under different simple policy rules



Source: Authors' calculations.

Under the baseline projection, the Inertial rule (solid black line) calls for a more gradual increase of the policy rate than the Non-inertial rule (dashed blue line) in response to the elevated level of inflation at the beginning of the simulation and a more gradual return to its longer-run level thereafter. Although these policy rate prescriptions lead to similar inflation paths, the Inertial rule prescribes more moderate policy rate changes than the Non-inertial rule, resulting in smaller fluctuations of the unemployment rate around its assumed longer-run level of 4 percent. For a hypothetical loss function that places equal weights on deviations of inflation from 2 percent and of unemployment from its longer-run level, the somewhat larger cumulated unemployment deviations associated with the Non-inertial rule amount to an underperformance of that rule relative to its inertial counterpart. In contrast, when the recession materializes, the opposite conclusion applies: After calling for a fast increase in the policy rate at the beginning of the simulation in response to elevated inflation, the Non-inertial rule (dotted green line) prescribes rapid policy rate cuts to stimulate the economy in the face of negative demand shocks. The resulting path of unemployment rate deviations is lower than under the Inertial rule (dash-dotted red line), implying better economic performance under the Non-inertial rule.

Going beyond this illustrative example, the literature has evaluated the performance of a wide set of rules in response to general business cycle fluctuations and to specific risks. This literature has proposed ways to refine simple benchmark rules in a manner that makes them robust in the face of specific risks, but none of these robust rules are applicable in all situations. For example, a drawback of simple benchmark rules is that because they are typically specified to respond to contemporaneous indicators in a fixed manner, they do not explicitly incorporate information about the future course of the economy or about risks to the outlook that are not

captured by those indicators. Forward-looking rules that respond to expected future outcomes represent an attempt to overcome this challenge.[10]

Another strand of the literature has suggested that when the policy rate is near the effective lower bound (ELB), monetary policy behavior should differ—roughly speaking, be looser—from its behavior when the ELB constraint is not a prominent concern. For example, by incorporating a threshold criterion or a makeup variable cumulating past misses of inflation from target when the policy rate is at the ELB, policy rules can implement the well-known lower-for-longer strategy.[11] These kinds of robust rules serve as useful guides when central banks are facing specific risks, though, again, none of them are applicable in all situations.

Another challenge in using simple rules is that the key indicators they focus on may be measured with error. In many cases, simple rules depend on fundamentally unobservable variables like the natural rate of unemployment or the equilibrium level of interest rates. In the presence of measurement error, some studies suggest that robust simple rules should respond relatively less to the mismeasured variables so as to not impart additional volatility to the economy.[12] First-difference rules go a step further and completely omit these unobservable policy "guideposts" like the equilibrium interest rate and are thus robust to the (unavoidable) mismeasurement of these variables.[13] However, these rules are highly inertial and, as a result, tend to underperform when the economic outlook changes rapidly. Overall, although simple rules are unlikely to be robust to all possible circumstances, they can be designed to be robust to some specific risks and measurement challenges.

Optimal control policy

Under optimal control, monetary policy is set to minimize expected losses while accounting for the structure of the economy and the process by which the private sector forms expectations. Depending on the complexity of the model used to calculate the optimal control policy, the policy rate may be a function of many more macroeconomic variables than a simple policy rule and may depend on expected future as well as lagged values of these variables.[14] The

[10] Although well-designed policy should be forward-looking to account for the lagged effects of monetary policy changes, policy that is overly forward-looking may not be responsive enough to incoming information and can result in undesirable fluctuations in economic activity and inflation; see, for example, Levin, Wieland, and Williams (2003). In addition, the performance of forward-looking rules depends on the accuracy of the forecasts used in the rules.

[11] Arias and others (2020) and Hebden and others (2020) discuss makeup strategies—policies that aim to offset, at least in part, past misses of inflation from its objective—and their robustness. Kiley (2024b) argues that ELB risk can be mitigated by incorporating longer-run inflation expectations into a simple policy rule.

[12] See the discussion in Boehm and House (2019).

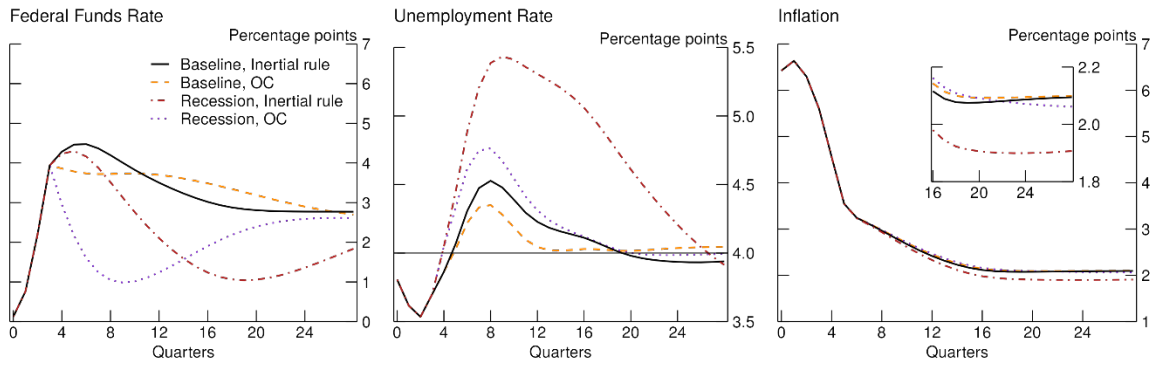
[13] See Orphanides and Williams (2002).

[14] In models in which private-sector decisionmakers are forward-looking, economic outcomes can be improved if current policymakers are able to make credible promises upon which future policymakers follow

typical implementation of optimal control in the Tealbook is one in which the evolution of the economy—according to the baseline outlook or alternative scenario—is taken as known with certainty, precluding formal risk-management considerations. Later in the paper, we discuss how the Tealbook optimal control simulations could be enhanced to explicitly consider the management of specific risks.

By their nature, optimal control policies lead to outcomes that are the best possible under the extensive assumptions—including the specific loss function—with which they are constructed. The forward-looking nature of these policies makes them potentially well suited to design a policy that can explicitly incorporate risk-management considerations. As economic developments reshape the economic outlook and the risks surrounding it, the prescriptions of optimal control adjust accordingly. In this way, the policy prescriptions are both forward-looking and data dependent.

Figure 3: Optimal control simulation under baseline and recession scenario



Source: Authors' calculations.

Figure 3 illustrates the flexibility of optimal control policies using the baseline and alternative scenarios of figure 1.[15] The optimal policies are constructed under the assumption that either the baseline outlook (dashed orange line) or the recession scenario (dotted purple line) materializes with certainty, consistent with the long-standing perfect foresight assumption in the Tealbook. Under the baseline, the optimal policy rate path is close to that of the Inertial rule, leading to similar unemployment and inflation outcomes. Under the recession scenario, the optimal policy path (dotted purple line) falls much more quickly than prescribed by the Inertial

through, even if they believe it is not optimal to do so—that is, the promises are time inconsistent along the lines described by Kydland and Prescott (1977); see also Duarte and others (2020). By contrast, under discretion, policymakers do not keep past promises and therefore lack the commitment required to directly influence private-sector expectations. As a result, policymakers take those expectations as given when setting policy optimally.

[15] The optimal control simulations posit that policymakers choose a path for the federal funds rate to minimize a discounted equally weighted sum of squared inflation gaps (measured as the difference between four-quarter headline personal consumption expenditures price inflation and the 2 percent target), squared unemployment gaps (measured as the difference between the unemployment rate and its longer-run level), and squared changes in the federal funds rate.

rule (dash-dotted red line). Consequently, the rise in the unemployment rate is more muted and inflation remains closer to the 2 percent target. The optimal policy delivers superior outcomes in both cases: It mimics the gradual response of the Inertial rule under the baseline as well as the rapid fall of the policy rate prescribed under the Non-inertial rule in response to the recession (shown in figure 2).

These simulations illustrate the optimal policy response to one specific adverse future outcome that is perfectly anticipated. However, they do not encompass risk management in the form of weighing multiple possible future outcomes, and later we extend the analysis to do so. However, optimal control analysis also has limitations. The prescriptions can hinge crucially on assumptions about the model, the nature of the uncertainty, and the loss function.[16] And though some guidelines emerge—discussed in the next subsection—optimal control policies under uncertainty can also often be complex and difficult to communicate because of the contingent nature of these strategies.

3.1 Optimal monetary policies under uncertainty in specific applications

Two important theoretical benchmarks in the discussion of monetary policy under uncertainty are certainty equivalence and the Brainard attenuation principle. We discuss these concepts in the context of optimal monetary policies.

Under certainty equivalence, which applies only under certain restrictive conditions, the policy that minimizes the expected loss is the same as the optimal policy for the case in which the uncertain outlook is replaced by its expected (mean) value.[17] In this context, it is in the process of computing the mean outlook—weighing outcomes by their probabilities—that policymaking incorporates risk management. In particular, certainty equivalence applies to mean outcomes, but monetary policy analysis often centers around a baseline modal path for the economy with a distribution of risks around this baseline. As discussed in Bauer and others (2025), the modal and mean outlooks will generally not coincide if risks are distributed asymmetrically around the modal projection, as in the case in the stylized example illustrated in figure 1.

A prominent example of deviation from certainty equivalence is the Brainard attenuation principle. Brainard (1967) offers a seminal contribution on how uncertainty about the structure of the economy—and therefore the effectiveness of policy—can affect the choice of an optimal policy. In the simplest case, Brainard (1967) considers a static problem in which a policymaker

[16] Although optimal control policies are more closely tied to a specific model, these limitations also apply to the evaluation of simple policy rules. Additionally, when monetary policy is set by a committee, agreeing on a single loss function (or a set of probabilities) is challenging. We abstract from this issue, though we discuss similar considerations related to alternative scenarios in section 4.

[17] In particular, in the case of quadratic policymaker objectives and a linear economic model, optimal policy is certainty equivalent. The essential feature of linear models is that the tradeoffs between policy objectives are constant. This precludes some salient sources of uncertainty such as those associated with policymaker misspecification of economic relationships.

seeks to optimally use a policy instrument—say, the federal funds rate—to bring an outcome variable—say, the inflation rate—to target. The policymaker is uncertain about the effect of the chosen policy on the target variable but knows the mean and variance of the distribution of possible effects. The policymaker's loss is measured in (squared) deviations of the outcome variable from target. The minimization of the expected loss establishes the Brainard attenuation principle: The more uncertainty there is about the effectiveness of a policy instrument, the less aggressive the policymaker should be in seeking to achieve the policy target. Intuitively, because the effects of policy are uncertain, the actions taken by the policymaker can affect the distribution—and, in this case, the mean and variance—of the target variable. A policymaker who more aggressively seeks to meet the target, on average, creates a higher volatility of the target variable. Therefore, in contrast to the case of certainty equivalence, the policymaker optimally chooses an attenuated policy response that accounts not only for the expected effect of the policy actions, but also for the uncertainty—and particularly the variance—of that effect. The attenuation principle is often referred to in policy discussions as the basis for gradualism in the face of uncertainty.[18]

Although certainty equivalence and the Brainard principle serve as useful reference points, there are limits to the guidance they can provide about the design of policy under uncertainty. Instead, more comprehensive analysis is typically required, and, in some cases, greater uncertainty calls for more aggressive action instead of attenuation. The rest of this section surveys some of this related literature and summarizes its findings on the optimal monetary policy response in the presence of some prominent sources of uncertainty.

Uncertainty about future shocks in the presence of nonlinearities

As discussed before, when certainty equivalence applies, the optimal control approach to risk management involves probability-weighting future outcomes into an expected value—that is, the mean—and basing the policy decision solely on the mean. However, basing policy solely on the assessment of the mean is suboptimal when there are nonlinearities or asymmetries affecting the economic outlook. Specifically, asymmetries can arise when future shocks push the economy, otherwise well approximated by a linear model, into regions of nonlinearities.[19] A linear model is often viewed as a reasonable approximation for normal times, but there are times when the economy displays nonlinear relationships. For example, several studies suggest that the Phillips curve may display nonlinearities at very low unemployment rates.[20] Another particularly salient nonlinearity is the possibility that the nominal interest rate will be constrained

[18] See Blinder (1998) for an example.

[19] A loss function that penalizes shortfalls of employment without penalizing high levels of employment would not be symmetric, and this asymmetry would push the optimal policy toward higher levels of employment than under certainty equivalence, with attendant consequences for the path of inflation; see the discussion in Bundick, Cairó, and Petrosky-Nadeau (2025).

[20] See Blanco and others (2024a, 2024b); Karadi and others (2024).

by the ELB. An implication of the ELB for risk management is that monetary policy should adopt a looser stance than typically implied by the mean outlook.[21]

Uncertainty about the state of the economy

Latent variables like potential output or the natural rate of unemployment, which are key to monetary policymaking, can be inferred from observable data only with considerable uncertainty, and important data on spending or the labor market are typically measured with error. Svensson and Woodford (2003) consider an environment in which there is imperfect information about the state of the economy.[22] They show that optimal policy in such a stylized environment involves two stages. In the first stage, the policymaker estimates the state of the economy from a set of indicators, placing less weight on noisier indicators. In the second stage, the policymaker sets optimal policy taking as given the best estimate of the state of the economy. Under the assumed structure of the economy, the optimal policy determined in the second stage is certainty equivalent and unaffected by the amount of noise. One implication of this framework is that the policymaker's understanding of the state of the economy may evolve only gradually, inducing inertia in policy.

A separate branch of the literature examines whether uncertainty about the natural rate of unemployment or potential output can motivate policies of experimentation or probing when a policymaker is also uncertain about the structure of the economy.[23] Conceptually, this entails setting policy with an eye toward actively learning about these latent variables for better future setting of policy. Broadly speaking, the literature has found that benefits from such an approach are often small. At the same time, active experimentation has costs and could have unintended consequences.

Uncertainty about expectations formation and the structure of the economy

As described in Bauer and others (2025), the transmission of monetary policy depends crucially on the public's expectations about the future course of monetary policy, the policy reaction function, and future inflation. A well-designed monetary policy strategy needs to be robust to uncertainty about the expectations formation process.

In many modern macroeconomic models, economic decisionmakers are assumed to be extremely perceptive: They consider outcomes far into the future and correctly assess the likelihood of these outcomes. In this setting, monetary policy can be very powerful (perhaps implausibly so) by influencing the public's expectations about policy far into the future.[24] In

[21] See Adam and Billi (2007); Nakov (2008); Evans and others (2015); Gust, Johannsen, and López-Salido (2015); and Yellen (2016).

[22] Partial information about the state is assumed to be symmetric—the private sector and the central bank both observe noisy signals of the true state.

[23] See the discussion in Wieland (2006).

[24] See Del Negro, Giannoni, and Patterson (2023) for a detailed description of the so-called forward guidance puzzle. See also McKay, Nakamura, and Steinsson (2016).

reality, the process by which economic decisionmakers form expectations, particularly about the longer run, is uncertain, and the stabilization of longer-term expectations is not necessarily straightforward.

The literature has generally found that monetary policy should give greater weight to the possible states in which longer-term inflation expectations may become unanchored, because although a state with unanchored inflation expectations might not be very likely, losses incurred in that state are high. This is an especially important consideration when there is uncertainty about the inflation expectations formation process. The literature has examined the robustness of different monetary policy strategies and compared their relative performance under different assumptions about expectations formation and found that a robust policy places a larger weight on inflation stabilization over the stabilization of real activity compared to policies that are not robust to this form of uncertainty.[25] Particularly in settings in which inflation may be "intrinsically" persistent through expectations formation—when expectations of future inflation depend on past inflation—and there is uncertainty about this relationship, the literature has found that policy should attempt to act forcefully to avoid inflation at very high or very low levels. This is because uncertainty about the persistence of inflation interacts with the level of inflation. When current inflation is further away from target, the possibility of high inflation persistence implies the risk of sustained future misses in the inflation objective. A similar result emerges when the source of the intrinsic persistence is not the expectations formation process but instead the structure of the economy—for example, if wage and price contracts are indexed to past inflation. In this case, too, the literature finds that optimal policy under uncertainty about inflation persistence does not follow the Brainard attenuation principle and instead reacts forcefully to inflation.

Another case in which the optimal policy under uncertainty may not follow the Brainard attenuation principle is when there is uncertainty about the effectiveness of policy in a dynamic setting. Note that Brainard's original analysis considered uncertainty about the effectiveness of policy within a static setting. However, in a dynamic setting with sufficiently forward-looking economic decisionmakers, an attenuated policy strategy could result in more volatile inflation expectations that destabilize the economy. This channel can work in the opposite direction of the Brainard principle and call for a more aggressive policy rate response.[26]

4. Alternative scenarios in monetary policy deliberations and communications

From a theoretical perspective, risk management would ideally involve accounting for the full spectrum of possible outcomes. In practice, this type of information is not available. Considering the implications of key risks to the outlook in the form of alternative scenarios is a practical route to incorporate risk management into policy deliberations. In addition, as

[25] Orphanides and Williams (2004), Gáti (2023), and Gust and López-Salido (2024) emphasize the importance of inflation stabilization in the context of the management of longer-term inflation expectations.

[26] See Dupraz, Guilloux-Nefussi, and Penalver (2023).

discussed in Bauer and others (2025), some central banks have made use of scenario analysis in public communications. We now discuss these two uses of alternative scenarios.

4.1 Alternative scenarios in monetary policy deliberations

In preparation for the meetings of the Federal Open Market Committee (FOMC), Board staff include alternative scenarios in the Risks and Uncertainty section of the Tealbook. Typically, the scenarios highlight the implications of alternatives to key assumptions in the baseline projection and other salient risks that are particularly important for the outlook. The design of the conditioning assumptions in the scenarios is judgmental, while the implications for the economic outlook are analyzed using a set of macroeconomic models maintained by the staff. In the alternative scenarios, monetary policy is typically governed by the interest rate rule used in the baseline projection.[27]

The use of alternative scenarios in deliberations has several benefits. Alternative scenarios illustrate salient risks to the outlook that can serve as useful points of reference in policy discussions. This may be particularly important under scenarios in which the dual-mandate goals are in conflict and therefore important tradeoffs emerge regarding the speed with which goal variables are returned to target. Policymakers can use the scenarios in conjunction with their own judgment of these risks to engage in risk management. Furthermore, as illustrated in section 3, alternative scenarios can be useful in evaluating macroeconomic outcomes under different monetary policy responses, as is done on occasion in the Monetary Policy Strategies section of the Tealbook.

Alternative scenarios as typically presented in the Tealbook also have some limitations. Scenarios are necessarily a subset of the full range of possible outcomes, and a particular concern is that the scenarios may understate tail risk, as tail events are typically not salient.[28] The emphasis on tail events highlights risks that could be highly consequential for macroeconomic outcomes, but there is a tradeoff: Frequent discussion of tail events could result in overrepresentation of such risks at the expense of other risks that are more likely to occur.

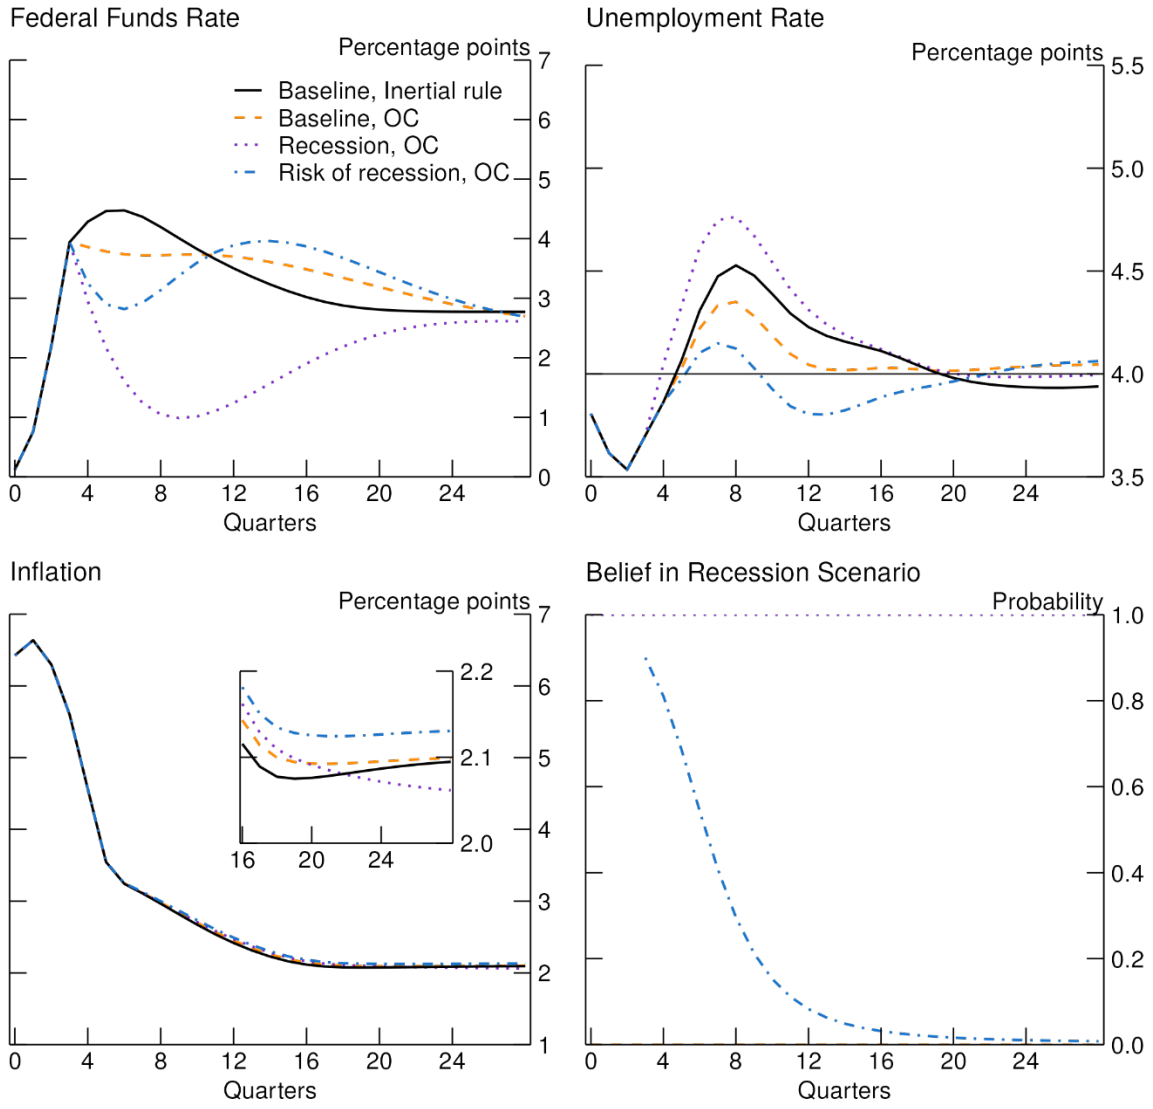
Furthermore, the risk-management approach we described in section 3 involves assigning probabilities to outcomes, but calculating scenario probabilities is challenging. Currently, probabilities are not attached to the alternative scenarios. To use the scenarios for the purpose of risk management, policymakers must then assign their own probabilities. Lastly, the typical assumption that the federal funds rate in the scenarios follows a simple interest rate rule may not give a realistic depiction of how monetary policy is likely to be set under the circumstances presented in the scenarios.

[27] In scenarios in which recessions occur, an adjustment to the baseline rule is usually made to capture the speed with which policymakers typically respond to the recession in the data.

[28] To better capture extreme events, the range of alternative scenarios considered might be informed by extreme outcomes from statistical models, such as those outlined in Bauer and others (2025).

One way in which scenario analysis in deliberations could be enhanced is by formally incorporating risk-management considerations. In that case, the likelihoods of alternative future outcomes play a central role, and uncertainty itself can shape the evolution of monetary policy and hence of the economy, even when a particular scenario materializes. When monetary policy incorporates risk-management considerations, the evolution of the economy in one scenario depends on beliefs about possible alternatives.

Figure 4: Optimal control adjusted for the risk of recession



Source: Authors' calculations.

Figure 4 illustrates how risk-management considerations may be incorporated into scenario analysis by accounting for the policymakers' beliefs and how this can lead to outcomes that are different from the typical Tealbook optimal control simulations (under which a given

scenario is assumed to materialize with certainty). For reference, we show the optimal control policies under the baseline (dashed orange line) and the alternative scenario (dotted purple line) from figure 3. We then superimpose on the figure the results obtained from a risk-adjusted optimal policy simulation, in which the policymaker minimizes expected losses in the face of uncertainty about which scenario will materialize, weighs the probability of each scenario, and updates these probabilities over time as shocks transpire and reveal which scenario is more likely.[29] In the simulation, we assume that the policymaker initially attaches a high likelihood to the recession scenario. However, the baseline scenario in fact materializes, and the policymaker learns this gradually over time. The policy rate path in the risk-adjusted simulation (dash-dotted blue line) is initially more accommodative than when the baseline materializes with certainty (dashed orange line), as the policymaker using a risk-adjusted strategy takes into account the possibility of a recession. However, as policymakers come to realize that the recession is less likely to materialize and that their risk-adjusted policy has been overly accommodative relative to the baseline (which becomes more likely over time), they eventually raise the policy rate above its baseline path.[30] A key aspect of this figure is that it demonstrates that risk-management considerations do not simply lead to a weighted average of the policy paths under the contemplated scenarios. Instead, the risk-adjusted policy path reflects the evolving beliefs of policymakers and the impact of previous policy decisions on current outcomes.[31]

4.2 Alternative scenarios in communications

The Federal Reserve does not currently use scenario analysis in its formal public communications. However, it has been suggested that it would be useful for the FOMC to do so.[32] Below, we describe potential roles that scenario analysis could play in central bank communications as well as some practical considerations in publishing scenarios.

One use of alternative scenarios in communications is in highlighting the message that a range of outcomes for the economy is possible, thereby putting less emphasis on the baseline outlook. This would underline the data-dependent nature of monetary policy, which may not be

[29] The policymaker updates beliefs about the scenarios as new information is acquired using Bayes rule: The posterior belief at the end of each period reflects the (prior) belief at the beginning of the period and the likelihood that a given scenario is true in light of the data received over the period. For more details, see Cairó and others (2025).

[30] Under a loss function that does not penalize interest rate changes, the initial accommodation and the subsequent reversal in the risk-adjusted policy are more accentuated, and the unemployment rate returns to its longer-run level earlier in the simulation. Conversely, the longer it takes for the policymaker to learn that the risk of recession does not materialize, the more inertial the risk-adjusted policy is, potentially leading to larger expected losses.

[31] Because households and firms use VAR-based expectations, they do not take uncertainty into account in their own spending and pricing decisions, isolating the role of uncertainty in risk management for policy.

[32] Bernanke (2025) offers proposals for changes to FOMC communications, including the use of alternative scenarios. See also Bordo, Levin, and Levy (2020) and Bernanke (2024). We note that scenarios are sometimes discussed by policymakers in the context of speeches and other communications.

easily conveyed through a single baseline forecast. De-emphasizing a baseline scenario could be particularly valuable at times of heightened uncertainty when there are one or more alternatives that are nearly as likely as the baseline.

Another use of alternative scenarios could be to explain the rationale for policy decisions that are influenced by risk-management considerations. More generally, alternative scenarios would provide the public with additional information about how the monetary policy strategy might respond to particular future events, thereby supplying more information about the reaction function. This might be particularly useful for communicating how policymakers would respond in situations in which the policy objectives are in tension.

Publishing alternative scenarios would require making a series of implementation choices. We list some of the key choices below.

- **What to show as the baseline?** Publishing alternative scenarios could require a decision about what to show as the baseline. A baseline forecast could represent the views of the FOMC or be a staff product.
- **What scenarios to show and when to show them?** The choice of a particular scenario is as much a part of the communication as the details of what the analysis shows. There may not always be salient risks that policymakers wish to highlight, raising the question of whether scenarios are used on a routine basis or only in special circumstances. Another practical consideration is the timing of the release relative to the FOMC meeting or other monetary policy communications and when to stop highlighting particular scenarios.
- **Is the policy rate path shown? What other variables are published?** If the intent is to clarify the monetary policy reaction function, then publishing the policy rate path would be essential.
- **How is monetary policy determined in the scenario analysis?** What is most useful to the public is to understand how monetary policy would respond if these scenarios were to materialize. However, it seems unlikely that policymakers, even when following a common strategy, would necessarily have a clear consensus on the appropriate policy rate path under all proposed hypothetical scenarios. Hence, monetary policy would likely have to be specified using an assumed monetary policy rule. Internationally, the Riksbank analyzes alternative scenarios using a macroeconomic model that includes a simple interest rate rule. In contrast, the Reserve Bank of Australia and the European Central Bank present scenarios under the assumption that the policy rate is not adjusted.
- **Whose views are represented?** The scenarios could be presented as the views of the staff or the policymakers. The Federal Reserve's internal scenario analysis is currently conducted using Board staff views of salient risks (partly informed by interactions with policymakers) and macroeconomic models, highlighting the important role that the staff might play in implementing scenario analysis. Alternatively, policymakers could be surveyed on how they

would likely respond to different scenarios or conditioning assumptions in a manner similar to the SEP process for the modal outlook.

- **How would the analysis be conducted?** If models are used in constructing the scenarios, how much information about those models would be made public? Other central banks have typically used their staff models for conducting alternative scenario analysis. The public communication often highlights the main model assumptions and mechanisms driving the results without revealing quantitative details of the model.

Implementation issues aside, the publication of alternative scenarios could have some unintended consequences. First, publishing a path of the policy rate under a scenario designed by Board staff may be misunderstood as communicating the reaction function of the FOMC and additionally misperceived as a commitment of the FOMC to follow that path if the risk materializes. Second, publishing a set of scenarios may not fully convey the uncertainty and all salient risks. At times of heightened focus on a specific risk, publishing or not publishing a relevant scenario may send an undesired signal about the FOMC's view on the scenario itself. Third, scenario analysis could result in the Federal Reserve implicitly commenting on the economic consequences of events and policies that are outside the purview of the central bank. Finally, if alternative scenarios are presented without reference to their likelihood, the public may misperceive the policymakers' judgement of risks and uncertainty. For example, an extreme scenario could be interpreted as likely, with unintended implications for private-sector expectations. As we have noted, there is some experience with publishing alternative scenarios by foreign central banks, but this is a relatively recent phenomenon, and it is too early to assess their utility and performance. Moreover, assessing the effectiveness of alternative scenarios as a communications tool is challenging, because there does not exist a well-established framework for evaluating the performance of most communication tools.

5. Conclusion

Setting policy under uncertainty requires flexibility to adapt to changing circumstances, but also a systematic approach so that policy is predictable and longer-term inflation expectations remain anchored. The implications for monetary policy of uncertainty depend crucially on the specific source of uncertainty and the modeling framework chosen for the analysis. Greater uncertainty does not generically tilt optimal monetary policy decisions to become more gradual, nor is a more aggressive approach universally preferable. At times of heightened risk, central bank communications play an important role by connecting current and future policy actions to risk-management considerations in the context of an evolving outlook. Alternative scenarios can be used to communicate how policy is responding or will respond to developments in a way that best achieves policy goals; however, using alternative scenarios in this way requires making several important implementation choices and weighing the associated benefits and costs.

References

Adam, Klaus, and Roberto M. Billi (2007). "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics,* vol. 54 (April), pp. 728–52.

Arias, Jonas, Martin Bodenstein, Hess Chung, Thorsten Drautzburg, and Andrea Raffo (2020). "Alternative Strategies: How Do They Work? How Might They Help?" Finance and Economics Discussion Series 2020-068. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2020.068.

Bauer, Michael, Travis Berge, Giuseppe Fiori, Francesca Loria, and Molin Zhong (2025). "Accounting for Uncertainty and Risks in Monetary Policy," Finance and Economics Discussion Series 2025-073. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.073.

Bernanke, Ben (2024). *Forecasting for Monetary Policy Making and Communication at the Bank of England: A Review.* London. Bank of England, April, https://www.bankofengland.co.uk/independent-evaluation-office/forecasting-for-monetary-policy-making-and-communication-at-the-bank-of-england-a-review/forecasting-for-monetary-policy-making-and-communication-at-the-bank-of-england-a-review.

——— (2025). "Improving Fed Communications: A Proposal," Hutchins Center Working Paper 102. Washington: Brookings Institution, May 16, https://www.brookings.edu/articles/improving-fed-communications-a-proposal.

Blanco, Andrés, Corina Boar, Callum J. Jones, and Virgiliu Midrigan (2024a). "Non-Linear Inflation Dynamics in Menu Cost Economies," NBER Working Paper Series 32094. Cambridge, Mass.: National Bureau of Economic Research, January, https://doi.org/10.3386/w32094.

——— (2024b). "The Inflation Accelerator," NBER Working Paper Series 32531. Cambridge, Mass.: National Bureau of Economic Research, May (revised March 2025), https://doi.org/10.3386/w32531.

Blinder, Alan S. (1998). *Central Banking in Theory and Practice.* Cambridge, Mass.: MIT Press.

Boehm, Christoph E., and Christopher L. House (2019). "Optimal Taylor Rules When Targets Are Uncertain," *European Economic Review,* vol. 119 (October), pp. 274–86.

Bordo, Michael D., Andrew T. Levin, and Mickey D. Levy (2020). "Incorporating Scenario Analysis into the Federal Reserve's Policy Strategy and Communications," NBER Working Paper Series 27369. Cambridge, Mass.: National Bureau of Economic Research, June, https://www.nber.org/papers/w27369.

Brainard, William C. (1967). "Uncertainty and the Effectiveness of Policy," *American Economic Review,* vol. 57 (May), pp. 411–25.

Bundick, Brent, Isabel Cairó, and Nicolas Petrosky-Nadeau (2025). "Labor Market Dynamics, Monetary Policy Tradeoffs, and a Shortfalls Approach to Pursuing Maximum Employment," Finance and Economics Discussion Series 2025-068. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.068.

Cairó, Isabel, Christopher Gust, James Hebden, Edward Herbst, Scott Konzem, and Giovanni Nicolò (2025). "Risk-Adjusted Optimal Policy for Scenario Analysis," unpublished paper, Board of Governors of the Federal Reserve System, Division of Monetary Affairs, June.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson (2023). "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics,* vol. 1 (March), pp. 43–79.

Duarte, Fernando, Benjamin K. Johannsen, Leonardo Melosi, and Taisuke Nakata (2020). "Strengthening the FOMC's Framework in View of the Effective Lower Bound and Some Considerations Related to Time-Inconsistent Strategies," Finance and Economics Discussion Series 2020-067. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2020.067.

Dupraz, Stéphane, Sophie Guilloux-Nefussi, and Adrian Penalver (2023). "A Pitfall of Cautiousness in Monetary Policy," *International Journal of Central Banking,* vol. 19 (August), pp. 269–323.

Evans, Charles, Jonas Fisher, François Gourio, and Spencer Krane (2015). "Risk Management for Monetary Policy Near the Zero Lower Bound," *Brookings Papers on Economic Activity,* Spring, pp. 141–96, https://www.brookings.edu/wp-content/uploads/2015/03/2015a_evans.pdf.

Gáti, Laura (2023). "Monetary Policy & Anchored Expectations—An Endogenous Gain Learning Model," *Journal of Monetary Economics,* vol. 140 supplement (November), pp. S37–S47.

Giannoni, Marc P., and Michael Woodford (2002). "Optimal Interest-Rate Rules: I. General Theory," NBER Working Paper Series 9419. Cambridge, Mass.: National Bureau of Economic Research, December, http://www.nber.org/papers/w9419.

Greenspan, Alan (2004). "Risk and Uncertainty in Monetary Policy," speech delivered at the Meetings of the American Economic Association, San Diego, January 3, https://www.federalreserve.gov/boarddocs/speeches/2004/20040103/default.htm.

Gust, Christopher J., Benjamin K. Johannsen, and David López-Salido (2015). "Monetary Policy, Incomplete Information, and the Zero Lower Bound," Finance and Economics Discussion Series 2015-099. Washington: Board of Governors of the Federal Reserve System, November, http://dx.doi.org/10.17016/FEDS.2015.099.

Gust, Christopher J., and David López-Salido (2024). "Optimal Monetary Policy with Uncertain Private Sector Foresight," Finance and Economics Discussion Series 2024-059. Washington: Board of Governors of the Federal Reserve System, July, https://doi.org/10.17016/FEDS.2024.059.

Gust, Christopher J., David López-Salido, and Steve Meyer (2017). "Asymmetric Monetary Policy and the Effective Lower Bound," *Research in Economics,* vol. 71 (September), pp 441–51.

Hebden, James, Edward P. Herbst, Jenny Tang, Giorgio Topa, and Fabian Winkler (2020). "How Robust Are Makeup Strategies to Key Alternative Assumptions?" Finance and Economics Discussion Series 2020-069. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2020.069.

Karadi, Peter, Anton Nakov, Galo Nuño, Ernesto Pasten, and Dominik Thaler (2024). "Strike While the Iron Is Hot: Optimal Monetary Policy under State-Dependent Pricing," CEPR Discussion Paper 19339. London: Centre for Economic Policy Research, August 9, https://cepr.org/publications/dp19339.

Kay, John, and Mervyn King (2020). *Radical Uncertainty: Decision-Making Beyond the Numbers*. New York: W. W. Norton & Company.

Kiley, Michael T. (2024a). "Monetary Policy, Employment Shortfalls, and the Natural Rate Hypothesis," Finance and Economics Discussion Series 2024-032. Washington: Board of Governors of the Federal Reserve System, May, https://doi.org/10.17016/FEDS.2024.032.

——— (2024b). "Monetary Policy Strategies to Foster Price Stability and a Strong Labor Market," Finance and Economics Discussion Series 2024-033. Washington: Board of Governors of the Federal Reserve System, May, https://doi.org/10.17016/FEDS.2024.033.

Kydland, Finn E., and Edward C. Prescott (1977). "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy,* vol. 85 (June), pp. 473–92.

Levin, Andrew T., Volker Wieland, and John C. Williams (1999). "Robustness of Simple Monetary Policy Rules under Model Uncertainty," in John B. Taylor, ed., *Monetary Policy Rules.* Chicago: University of Chicago Press, pp. 263–318, http://www.nber.org/chapters/c7418.

——— (2003). "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty," *American Economic Review,* vol. 93 (June), pp. 622–45.

McCallum, Bennett T. (2002). "Recent Developments in Monetary Policy Analysis: The Roles of Theory and Evidence," Federal Reserve Bank of Richmond, *Economic Quarterly,* vol. 88 (Winter), pp. 67–96,

https://www.richmondfed.org/~/media/richmondfedorg/publications/research/economic_quarterly/2002/winter/pdf/mccallum.pdf.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016). "The Power of Forward Guidance Revisited," *American Economic Review,* vol. 106 (October), pp. 3133–58.

Nakov, Anton (2008). "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking,* vol. 4 (June), pp. 73–127.

Orphanides, Athanasios, and John C. Williams (2002). "Robust Monetary Policy Rules with Unknown Natural Rates," *Brookings Papers on Economic Activity,* no. 2, pp. 63–118, https://www.brookings.edu/wp-content/uploads/2002/06/2002b_bpea_orphanides.pdf.

——— (2004). "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in Ben S. Bernanke and Michael Woodford, eds., *The Inflation-Targeting Debate*. Chicago: University of Chicago Press, pp. 201–46, https://www.nber.org/system/files/chapters/c9559/c9559.pdf.

Svensson, Lars E.O. (2002). "Inflation Targeting: Should It Be Modeled as an Instrument Rule or a Targeting Rule?" *European Economic Review,* vol. 46 (May), pp. 771–80.

Svensson, Lars E.O., and Michael Woodford (2003). "Indicator Variables for Optimal Policy," *Journal of Monetary Economics,* vol. 50 (April), pp. 691–720.

Taylor, John B. (1999). "A Historical Analysis of Monetary Policy Rules," in John B. Taylor, ed., *Monetary Policy Rules.* Chicago: University of Chicago Press, pp. 319–48, https://www.nber.org/system/files/chapters/c7419/c7419.pdf.

Taylor, John B., and John C. Williams (2011). "Simple and Robust Rules for Monetary Policy," in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics, Volume 3B*. Amsterdam: Elsevier, pp. 829–59, https://doi.org/10.1016/B978-0-444-53454-5.00003-7.

Wieland, Volker (2006). "Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism," *Contributions in Macroeconomics,* vol. 6 (March), pp. 1–34, https://doi.org/10.2202/1534-6013.1288.

Williams, John (2025). "Uncertainty and Robust Monetary Policy," remarks at the Reykjavík Economic Conference, Reykjavík, Iceland, May 9, https://www.newyorkfed.org/newsevents/speeches/2025/wil250509.

Yellen, Janet L. (2016). "The Outlook, Uncertainty, and Monetary Policy," speech delivered at the Economic Club of New York, New York, March 29, https://www.federalreserve.gov/newsevents/speech/yellen20160329a.htm.

FRBSF Economic Letter
2022-15 | June 21, 2022
Research from the Federal Reserve Bank of San Francisco

How Much Do Supply and Demand Drive Inflation?

Adam Hale Shapiro

Inflation has remained at levels well above the Federal Reserve's inflation goal of 2% for over a year. Separating the underlying data from the personal consumption expenditures price index into supply- versus demand-driven categories reveals that supply factors explain about half of the run-up in current inflation levels. Demand factors are responsible for about one-third, with the remainder resulting from ambiguous factors. While supply disruptions are widely expected to ease this year, this outcome is highly uncertain.

Inflation declined rapidly at the onset of the pandemic in the spring of 2020 before taking a dramatic turn upward in early 2021, rising to levels that remain well above the Federal Reserve's longer-run goal of 2% on average. Researchers and policymakers have pointed to both supply and demand factors as being responsible for elevated inflation. For instance, Barnichon and Shapiro (2022) showed the impact of supply-related factors such as labor shortages, while Barnichon, Oliveira, and Shapiro (2021) and Jordà et al. (2022) demonstrated the importance of heightened demand stemming from pandemic-related fiscal relief. The extent to which either supply or demand factors are responsible for higher inflation levels has important implications for monetary policy. As Fed Chair Powell stated in a recent interview, "What [the Fed] can control is demand, we can't really affect supply with our policies…so the question whether we can execute a soft landing or not, it may actually depend on factors that we don't control" (Marketplace 2022).

In this *Economic Letter*, I quantify and track the impact of supply- and demand-related factors on personal consumption expenditures (PCE) inflation. Similar to the methodology introduced in Mahedy and Shapiro (2017), and outlined in Shapiro (2022), I assess inflation rates by spending category. I divide categories in the PCE basket into supply- and demand-driven groups. Demand-driven categories are identified as those where an unexpected change in price moves in the same direction as the unexpected change in quantity in a given month; supply-driven categories are identified as those where unexpected changes in price and quantity move in opposite directions. This methodology accounts for the evolving impact of supply- versus demand-driven factors on inflation from month to month. To help monitor these changes, the San Francisco Fed has launched a new Supply- and Demand-Driven PCE Inflation data page with monthly data updates.

My analysis highlights that both supply and demand factors are responsible for current elevated inflation levels. Supply factors explain about half of the difference between current 12-month PCE inflation and pre-pandemic inflation levels, and the effects appear to be rising more recently. Demand factors are responsible for about a third of the difference, and those effects appear to be diminishing more recently. The remainder is due to factors that cannot be definitively labeled as supply or demand. The large impact

of supply factors implies that inflationary pressures will not completely subside until labor shortages, production constraints, and shipping delays are resolved. Although supply disruptions are widely expected to ease this year, this outcome is highly uncertain.

Separating supply and demand drivers

This analysis uses the more than 100 goods and services categories in the PCE index. For each month of data, I separate the categories where prices moved due to a change in demand from those where prices moved due to a change in supply. To do so, I rely on simple microeconomic theory: Shifts in demand move both prices and quantities in the same direction along the upward-sloping supply curve, meaning prices rise as demand increases. Shifts in supply move prices and quantities in opposite directions along the downward-sloping demand curve, meaning prices rise when supplies decline.

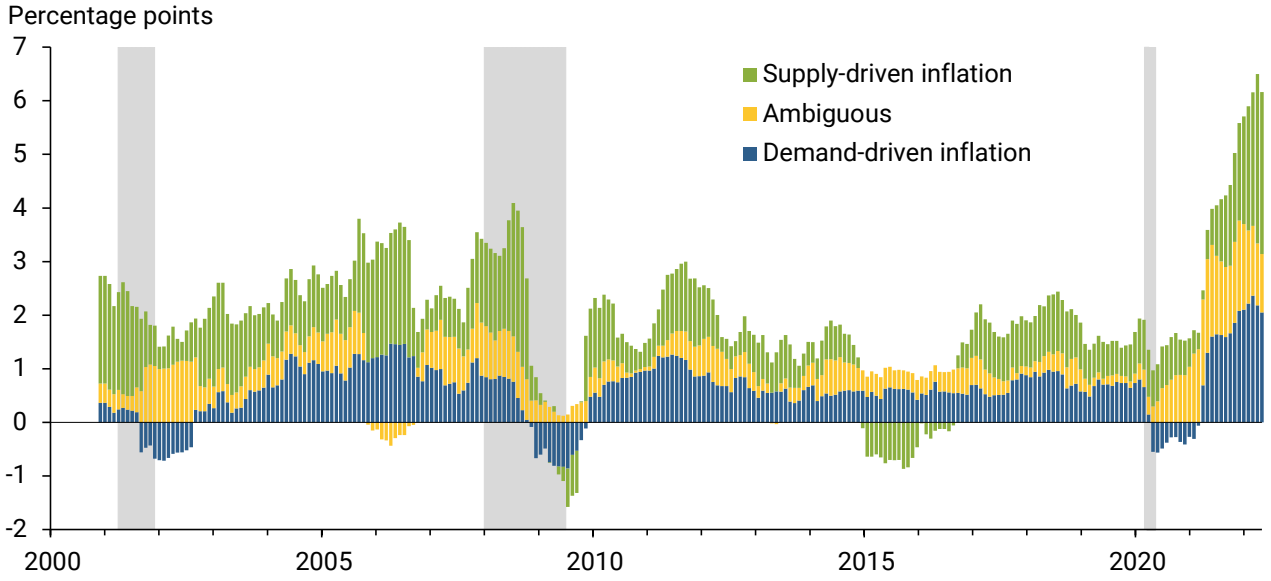
Each month I estimate the changes in the price level and quantity level of each category. It is important to isolate the *unexpected* components of the changes in prices and quantities, as opposed to the simple change itself. This is because prices and quantities generally continue an existing trend. These expected trend components are not likely to represent a shift in demand or supply, but instead reflect longer-run factors such as technological improvements, cost-of-living adjustments to wages, or demographic changes like population aging.

Extracting the unexpected components of the monthly changes in price and quantity for each category is an iterative process. I run 10-year-window rolling regressions for both price and quantity. For example, the first window begins in January 1988, the first period PCE data are available at the detailed level, and ends in December 1997. This generates predicted values for price and quantity in January 1998, which I then compare to the actual values of price and quantity in that same month. If the actual values of price and quantity are both above or both below their predicted values, the category is labeled as “demand-driven” in January 1998. If the difference between the actual and predicted values are of opposite signs, the category is labeled as “supply-driven” in that month. If either of the actual values is close enough to its predicted value that the difference is statistically indistinguishable from zero, the category is labeled as “ambiguous” in that month. I then roll the data window forward one month and repeat the process. I iterate this process for each month until I reach the last window of data, which for this *Letter* begins in May 2012 and ends in April 2022.

Categories that experience frequent supply-driven price changes include food and household products such as dishes, linens, and household paper items. Categories that experience frequent demand-driven price changes include motor vehicle-related products, used cars, and electricity. However, assessing the past two years of data reveals some changes in which categories experienced demand and supply shocks in the post-pandemic period. Specifically, categories with extraordinarily frequent supply-driven price changes in 2021 and 2022 include products with known supply constraints during the pandemic, such as new automobiles, fuel, and repair services. Categories with extraordinarily frequent demand-driven price changes during this period include many goods consumed at home—for example, furniture, clothing, toys, video equipment, and cookware—as well as services related to reopening from pandemic-related closures, such as restaurants and museums.

Figure 1

Supply-driven and demand-driven contributions to year-over-year PCE inflation



Note: Data available at Supply- and Demand-Driven PCE Inflation. Gray shading indicates NBER recession dates.

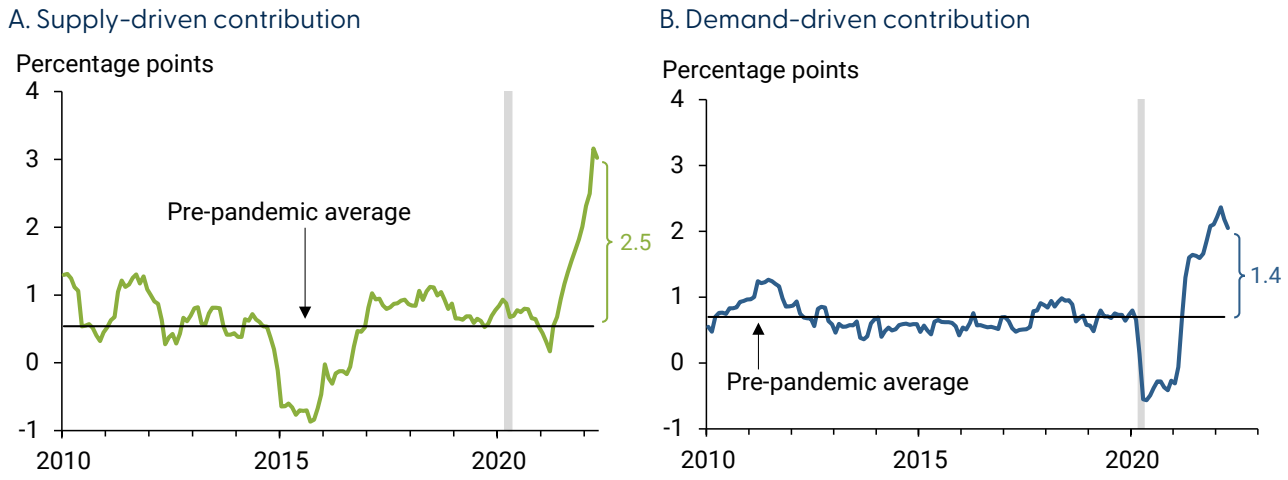
Figure 1 shows the contributions to 12-month headline PCE inflation from supply- (green) and demand-driven (blue) inflation between 1998 and 2022. Yellow sections show the ambiguous portion of PCE inflation that is not labeled as either demand or supply driven. I calculate the monthly contributions to inflation by multiplying the change in price for a given category by its most recent respective weight in the PCE index. The monthly contributions for demand-driven, supply-driven, and ambiguous categories are then calculated as the sum of contributions from all items within that category. The contributions for demand-driven and supply-driven inflation are then calculated as the 12-month trailing sum of their respective monthly contributions. This generates a series that can be directly compared to the year-over-year PCE inflation rate.

The changes in patterns over time show some intuitive dynamics. The contribution of demand-driven factors to headline PCE inflation declines during recessions. The collapse in airline travel immediately after September 11, 2001, is labeled as demand-driven and the sharp energy price declines in 2014 and 2015 are identified as supply-driven.

Inflation drivers over the pandemic period

To get a sense of whether supply or demand factors are responsible for current elevated inflation levels, I compare current supply- and demand-driven inflation to their average levels from the 10 years before the pandemic. During the pre-pandemic period, PCE inflation averaged 1.5%, considerably below the April 2022 rate of 6.3%. Figure 2 compares the most recent contributions of supply-driven inflation in panel A and demand-driven inflation in panel B against their respective 2010–2019 averages. Supply-driven inflation is currently contributing 2.5 percentage points (pp) more than its pre-pandemic average, while

Figure 2
Contributions of supply- and demand-driven factors to headline PCE inflation



Note: Gray shading indicates NBER recession dates.

demand-driven inflation is currently contributing 1.4pp more. Thus, supply-driven inflation explains a little more than half of the 4.8pp gap between current levels of year-over-year PCE inflation and its pre-pandemic average level. Demand factors explain a smaller share of elevated inflation levels, accounting for about one-third of the difference. The ambiguous category, which is not shown, explains the remainder of the difference.

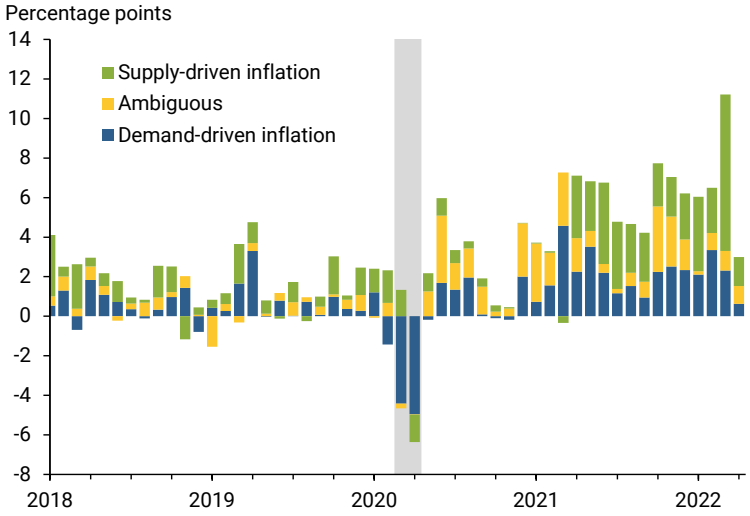
Repeating the same exercise with core PCE inflation, which excludes historically volatile food and energy categories, results in somewhat similar patterns. However, supply and demand factors each explain about half of elevated core inflation levels. Specifically, supply- and demand-driven factors each explain about 45% of the 3.3pp gap between current levels of year-over-year core PCE inflation and its pre-pandemic average level. Monthly results of these contributions to core PCE inflation are available on the data page.

Finally, I assess the contributions to annualized monthly changes in headline PCE inflation. This provides a higher frequency depiction of how supply and demand factors affected inflation over the pandemic period. Figure 3 breaks down these contributions over past five years, showing the one-month changes used to construct the 12-month trailing sums shown in Figure 1.

The decline in inflation at the onset of the pandemic was driven by a decrease in demand factors, while the surge in inflation in March 2021 was mainly due to the increase in demand-driven factors. During this period the economy began to reopen from pandemic-related public health policies, and the American Rescue Plan enacted in March 2021 further stimulated demand factors. These factors began to slow in the summer of 2021 with an increase COVID-19 infections associated with the Delta variant, but demand reemerged in the fall as the Delta wave subsided.

Meanwhile, supply factors began to arise in April 2021, indicating a slightly delayed response from the economy reopening. Supply-driven inflation has remained elevated since then and has accelerated more recently. This acceleration is attributable to food and energy supply disruptions, including those associated with the invasion of Ukraine.

Figure 3
Contributions to annualized monthly changes in inflation



Note: Gray shading indicates NBER recession dates.

Conclusion

Analysis in this *Letter* shows that supply factors are responsible for more than half of the current elevated level of 12-month PCE inflation. This in part reflects supply constraints from continued labor shortages and global supply disruptions related to the pandemic and the war in Ukraine. While demand factors played a large role in the spring of 2021, they explain only about a third of recent elevated inflation levels. Factors that cannot be labeled as either demand or supply are also playing a nontrivial role.

These results showing that factors other than demand account for about two-thirds of recent elevated inflation highlight some risks for the economy. Because supply shocks raise prices and suppress economic activity, the prevalence of supply-related factors raises the risk of entering a period of low growth and elevated inflation levels. This risk depends crucially on how long labor shortages and global supply disruptions persist. While supply disruptions are widely expected to ease this year, this outcome is highly uncertain.

Adam Hale Shapiro

Vice President, Economic Research Department, Federal Reserve Bank of San Francisco

References

Jordà, Òscar, Celeste Liu, Fernanda Nechio, and Fabian Rivera-Reyes. 2022. "Why Is U.S. Inflation Higher than in Other Countries?" *FRBSF Economic Letter* 2022-07 (March 28).

Barnichon, Regis, Luiz E. Oliveira, and Adam Hale Shapiro. 2021. "Is the American Rescue Plan Taking Us Back to the'60s?" *FRBSF Economic Letter* 2021-27 (October 18).

Barnichon, Regis, and Adam Hale Shapiro. 2022. "What's the Best Measure of Economic Slack?" *FRBSF Economic Letter* 2022-04 (February 22).

Mahedy, Tim, and Adam Hale Shapiro. 2017. "What's Down with Inflation?" *FRBSF Economic Letter* 2017-35 (November 27).

Powell, Jerome. 2022. "Inflation, Soft Landings and the Federal Reserve." Interview by Kai Ryssdal, Marketplace Business News Podcast, NPR, May 12. Audio, 27:38.

Shapiro, Adam Hale. 2022. "A Simple Framework to Monitor Inflation." FRB San Francisco Working Paper 2020-29.

Opinions expressed in *FRBSF Economic Letter* do not necessarily reflect the views of the management of the Federal Reserve Bank of San Francisco or of the Board of Governors of the Federal Reserve System. This publication is edited by Anita Todd and Karen Barnes. Permission to reprint portions of articles or whole articles must be obtained in writing. Please send editorial comments and requests for reprint permission to research.library@sf.frb.org

Recent issues of *FRBSF Economic Letter* are available at

https://www.frbsf.org/economic-research/publications/economic-letter/

2022-14	Bok / Petrosky-Nadeau	Estimating Natural Rates of Unemployment https://www.frbsf.org/economic-research/publications/economic-letter/2022/may/estimating-natural-rates-of-unemployment/
2022-13	Lansing	Untangling Persistent versus Transitory Shocks to Inflation? https://www.frbsf.org/economic-research/publications/economic-letter/2022/may/untangling-persistent-versus-transitory-shocks-to-inflation/
2022-12	Duzhak	Pandemic Unemployment Effects across Demographic Groups https://www.frbsf.org/economic-research/publications/economic-letter/2022/may/pandemic-unemployment-effects-across-demographic-groups/
2022-11	Bauer / Mertens	Current Recession Risk According to the Yield Curve https://www.frbsf.org/economic-research/publications/economic-letter/2022/may/current-recession-risk-according-to-yield-curve/
2022-10	Daly	Steering Toward Sustainable Growth https://www.frbsf.org/economic-research/publications/economic-letter/2022/april/steering-toward-sustainable-growth-speech/
2022-09	Albert /Lofton / Petrosky-Nadeau Valletta	Unemployment Insurance Withdrawal https://www.frbsf.org/economic-research/publications/economic-letter/2022/april/unemployment-insurance-withdrawal/
2022-08	Hobijn	"Great Resignations" Are Common During Fast Recoveries https://www.frbsf.org/economic-research/publications/economic-letter/2022/april/great-resignations-are-common-during-fast-recoveries/
2022-07	Jordà / Liu / Nechio / Rivera-Reyes	Why Is U.S. Inflation Higher than in Other Countries? https://www.frbsf.org/economic-research/publications/economic-letter/2022/march/why-is-us-inflation-higher-than-in-other-countries/

NBER WORKING PAPER SERIES

WHAT CAUSED THE US PANDEMIC-ERA INFLATION?

Olivier J. Blanchard
Ben S. Bernanke

Working Paper 31417
http://www.nber.org/papers/w31417

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2023

This research was funded by the Hutchins Center on Fiscal and Monetary Policy at Brookings and the Peterson Institute for International Economics. Neither institution has a financial, ideological, or political stake in its findings. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Olivier J. Blanchard and Ben S. Bernanke. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Caused the US Pandemic-Era Inflation?
Olivier J. Blanchard and Ben S. Bernanke
NBER Working Paper No. 31417
June 2023
JEL No. E31

ABSTRACT

We answer the question posed by the title by specifying and estimating a simple dynamic model of prices, wages, and short-run and long-run inflation expectations. The estimated model allows us to analyze the direct and indirect effects of product-market and labor-market shocks on prices and nominal wages and to quantify the sources of U.S. pandemic-era inflation and wage growth. We find that, contrary to early concerns that inflation would be spurred by overheated labor markets, most of the inflation surge that began in 2021 was the result of shocks to prices given wages. These shocks included sharp increases in commodity prices, reflecting strong aggregate demand, and sectoral price spikes, resulting from changes in the level and sectoral composition of demand together with constraints on sectoral supply. However, although tight labor markets have thus far not been the primary driver of inflation, we find that the effects of overheated labor markets on nominal wage growth and inflation are more persistent than the effects of product-market shocks. Controlling inflation will thus ultimately require achieving a better balance between labor demand and labor supply.

Olivier J. Blanchard
Peterson Institute for International Economics
1750 Massachusetts Avenue, NW
Washington, DC 20036
and NBER
oblanchard@piie.com

Ben S. Bernanke
The Brookings Institution
1775 Massachusetts Avenue, NW
Washington, DC 20036-2103
bensbernanke75@brookings.edu

Replication files are available at http://www.nber.org/data-appendix/w31417

Introduction

Central bankers and most outside economists failed to predict the sharp rise in inflation that began in 2021, and policymakers, both in the United States and in other advanced economies, were accordingly slow to react. Action was further delayed by the view that the inflation burst would be temporary.

Misjudgments of the economic effects of covid-era fiscal programs are one potential explanation of the failure to forecast the subsequent inflation. Federal Open Market Committee (FOMC) forecasts of inflation changed little between the September 2020 and June 2021 iterations of the *Survey of Economic Projections*, and the Committee continued to expect that inflation would return close to the Fed's 2 percent target by 2023, despite congressional passage of a bill in December 2020 that included $900 billion for covid relief and the $1.9 trillion American Rescue Plan signed by President Biden in March 2021. These two programs came on top of the $2.2 trillion CARES Act, passed in March 2020 and signed by President Trump, which had already strengthened firms' and households' balance sheets and increased future ability to spend. Overall, as a share of GDP, the headline costs of these three covid-era fiscal packages were about 4-1/2 times the size of the American Recovery and Reinvestment Act (ARRA), enacted in response to the 2008 financial crisis and the ensuing recession.[1]

According to conventional economic theory, expansionary fiscal policies can stoke inflation if they cause labor markets to become overheated and output to exceed the economy's potential. Indeed, some early analyses of the American Rescue Plan using standard fiscal multipliers concluded that the additional federal spending would indeed overheat the economy, possibly leading to higher inflation (Blanchard, 2021; Edelberg and Sheiner, 2021). These studies acknowledged uncertainty about magnitudes, as the fiscal impacts on employment, output, and inflation would depend on many factors, including the effects of the pandemic on potential output, the size of the initial output gap, households' propensities to consume out of the one-time checks authorized by the bill, and the relation between labor market tightness and inflation.

Inflation optimists (Reifschneider and Wilcox 2022) argued that, even if the new fiscal spending were to drive down unemployment by more than expected, a large burst of inflation was still unlikely, as many studies had concluded that the Phillips curve is quite flat (that is, inflation is relatively insensitive

[1] The headline cost of the CARES Act and the December supplemental appropriations package together totaled about 15 percent of gross domestic product (GDP) in 2020, and the American Rescue Plan was roughly 8 percent of 2021 GDP. By comparison, the ARRA of 2009 was slightly more than 5 percent of that year's GDP.

to labor market tightness), and—in light of more than three decades of low inflation in the United States—inflation expectations were likely to remain well anchored.

In contrast, some economists (including one of the authors) argued that wage inflation, and consequently price inflation, could rise much more than predicted by conventional calculations predicated on a flat Phillips curve (Summers, 2021; Blanchard, 2021). Their concerns were that the increase in aggregate demand likely to result from the unprecedentedly large fiscal transfers, together with the cumulative effects of the easing of monetary policy begun in March 2020, could cause more overheating of an already-tight labor market than the optimists expected. An extremely low unemployment rate might in turn cause the Phillips curve to steepen. Moreover, higher and, consequently, more psychologically salient levels of inflation might lead inflation expectations to de-anchor, raising the potential for a wage-price spiral.

The critics' forecasts of higher inflation would prove to be correct—indeed, even too optimistic—but, in substantial part, the sources of the inflation, at least in its early stages, would prove to be different from those they warned about. The labor market did tighten significantly in 2021 and 2022, as reflected in several indicators, including unsustainably high rates of job creation, an increasing ratio of job openings to unemployed workers, and low levels of quits.[2] But, as this paper will show, the behavior of nominal wages in 2021 and 2022 was broadly consistent with the relations the Fed had relied on, and medium-term inflation expectations rose modestly by some measures but did not de-anchor. Labor market overheating would indeed ultimately prove to be a source of persistent inflation, as the critics had expected, but the traditional wage Phillips curve mechanism was not the main event, at least not until recently.

In retrospect, the failure to forecast the inflation burst reflected in large part the fact that, in focusing on the labor market, both the Fed and its critics underestimated the inflationary potential of developments in goods markets, that is, from increases in prices given wages. The shocks to prices came from several sources. First, like other forecasters (including participants in commodity futures markets), most economists did not anticipate either the magnitude or the duration of the rise in commodity prices that began in 2021. Beyond the more-familiar (but unexpectedly large and sustained) shocks to food and energy prices, the pandemic itself led to unusual distortions in key product markets, such as those for new and used vehicles. Strong aggregate demand and shifts in the composition of consumer spending (e.g., from services to durable goods) during the pandemic combined with constraints on supply in some

[2] By contrast, the unemployment rate, which remained close to FOMC projections, showed fewer signs of labor market tightening.

sectors to create shortages and sectoral price increases not offset by decreases in other sectors (Guerrieri et al., 2022; di Giovanni et al., 2023). These sectoral mismatches between demand and supply proved more intractable and longer-lasting than many had expected. Together, these shocks to prices given wages would prove to be the critical triggers of the rise in inflation.

This is the story we develop in this paper.

We start with a simple analytical model, focusing on the behavior of wages, prices, and short and long-run inflation expectations, taking labor market slack and shocks to prices as given. We show the very different dynamic effects on wage and price inflation of overheating in the labor market and price shocks in goods markets. Appreciating these differences is essential for understanding the changing relative importance of the two sources of inflation, and by implication, the evolution of inflation over time.

We then estimate an empirical version of the model, allowing for a generous lag structure to accommodate flexible dynamics and estimating it on the pre-covid period. We examine the stability of the wage, price, and short and long-run inflation expectations relations betweem the pre-covid and covid periods, taking into account the roles of labor market tightness, energy and food price shocks, and sectoral shortages. We use the estimated model to decompose the sources of the pandemic-era inflation into the effects of the various shocks. Our decomposition captures general equilibrium effects, lags, and linkages among alternative sources of inflation operating through inflation expectations and other channels. We conclude that labor market tightness made at most a modest contribution to inflation early on, resolving the puzzle of how inflation could rise so much despite a flat wage Phillips curve. Instead, most of the early action in inflation came from the goods market, in the form of sharp increases in some relative prices, including commodity prices and prices in sectors in which strong demand confronted limits on supply. To the extent that commodity prices stabilize and sectoral shortages moderate—processes that both seem well advanced as of this writing (June 2023)—the goods market component of inflation is likely to decrease in importance, and the labor market component to become more dominant. Looking forward, with labor market slack still below sustainable levels and inflation expectations modestly higher, we conclude that the Fed has to slow the economy to return inflation to target. The extent of that slowing will depend however on the evolution of certain structural features of the labor market, notably the efficiency of the process of matching workers with jobs.

I. A Brief Literature Review

Much recent work has provided empirical analyses of the pandemic-era inflation. Ball et al (2022) is the closest to ours, and we learned from it. They propose a decomposition of inflation into weighted median inflation (instead of the conventional core measure of inflation) and a residual, equal to headline inflation less median inflation. Their use of median inflation rather the more conventional measure of core inflation is intended to remove large price spikes, due not only to the prices of energy and food, which core inflation does, but also those due to sectoral shortages, which have made this inflation episode different from earlier ones. They find that the increase in overall inflation resulted primarily from shocks to the residual (headline less median inflation) but that labor market tightness has been an important and increasing contributor to median inflation. These conclusions are broadly consistent with ours, but differences in sample period and model specification (e.g., their specification implies that the effects of labor market tightness are felt only gradually) lead them to be generally more pessimistic than we are about the costs of reversing the inflation surge.

Cecchetti et al (2023) study the recent inflation and the prospects for disinflation in a historical context. As we do, they provide a stripped-down model of the inflation process. However, in contrast to the largely unconstrained dynamics that we allow in the empirical model presented below, they impose a tight New Keynesian structure. They argue that nonlinearities in the Phillips curve are important for explaining inflation (see also Benigno and Eggertsson 2023, Gagnon and Sarsenbayev 2022). In contrast, we find that a linear relation between labor market tightness and wages, estimated on pre-covid data, fits the covid-era data reasonably well and is adequate to explain the recent behavior of wages. Cecchetti et al also conclude that, based on the historical record, disinflation is likely to be costly, and they fault the Fed for failing to tighten policy pre-emptively.

Gagliardone and Gertler (2023) focus on the monetary-policy origins of the inflation. Using a calibrated DSGE model to structure their analysis, they argue that the Fed's commitment to high employment in the face of inflationary oil price shocks explains the sharp rise in the pace of price increases. Gagliardone and Gertler's attention to oil price shocks parallels our own focus on factors directly affecting prices given wages. Their main contributions concern the role of the Fed's reaction function and policy framework in the over-expansion of aggregate demand during the recovery from the pandemic—an important topic but not the focus of our own paper.

Koch and Noureldin (2023) of the International Monetary Fund conduct a post-mortem of their institution's failure to forecast the (global) inflation. Like other papers in the literature, their work finds

an important role for increased aggregate demand and, especially, pandemic-induced supply constraints in setting off the inflation. They also point to over-estimation of output gaps as a factor leading to excessive stimulus. However, they pay limited attention to dynamics, and in particular they place less weight than we would on commodity price shocks and their effects on other prices and inflation expectations over time.

II. A Simple Model of Wage-Price Determination

To help interpret recent inflation developments without imposing too much a priori structure, we use a bare-bones but flexible model of aggregate wage-price determination. Our approach to inflation is aggregative, or "top down"; with a partial exception discussed in the next section, we do not consider the microeconomic determinants of sectoral price changes, e.g., factors affecting the behavior of rents or the differential impact of the pandemic on different industries. Our model reduces to four equations, which jointly determine nominal wages, prices, and short-run and long-run inflation expectations. As already noted, when we take the bare-bones model to the data in the next section we incorporate a more flexible lag structure, to allow for richer dynamics, and include measures of key shocks to product and labor markets to account more explicitly for the forces affecting inflation.

The wage equation

We start with an aggregate wage equation. We make the standard assumptions that nominal wages depend on the expected price level and the degree of labor market slack. In a departure from the standard specification, we pursue the question of whether workers accept losses in purchasing power resulting from unanticipated price shocks, or whether instead, in their wage bargaining, they try to "catch up" with prior inflation.[3]

Specifically, we assume that the nominal wage *w* in each quarter depends on the expected price for that quarter p^e, an aspiration real wage ω^A (see below), and an indicator of the tightness of the labor market x:

(Eq 1) $w = p^e + \omega^A + \beta x$

[3] Kamin and Roberts (2023) analyze how alternative forms of real-wage catch-up might affect the course of the economy.

Wages, prices, and expected prices are in log-levels, so differences should be interpreted as growth rates. We use ω to designate the real aspiration wage as a reminder that that variable is in real terms, in distinction to the nominal wage *w*. We discuss empirical measures of labor market tightness x in the next section, only noting here the sign convention that we interpret higher values of x as corresponding to a tighter labor market. We suppress the time subscript for current-quarter variables.

We include the aspiration real wage ω^A in the model as a device for modeling a possible catch-up effect, in which workers seek to make up for past losses of purchasing power, implying a degree of real-wage rigidity. The aspiration wage may be thought of as the real wage that workers believe is achievable—a focal point in wage bargaining, perhaps—given recent developments in wages, prices, and other factors. Alternatively, the aspiration real wage might be thought of as representing the long-run equilibrium real wage, toward which the actual real wage gradually adjusts. Adopting the notation $q(-n)$ to stand for the variable q lagged n quarters, we assume that the aspiration real wage in each quarter is a weighted average of the previous quarter's aspiration wage and realized real wage, plus a shock term z_ω, standing for all the other factors that affect wage determination:

(Eq 2) $$\omega^A = \alpha\omega^A(-1) + (1-\alpha)\big(w(-1) - p(-1)\big) + z_\omega$$

When estimating the model, we also allow wage growth to depend on an exogenous trend in productivity growth. For simplicity of exposition, we ignore that term here.

Putting the two equations above together to eliminate ω^A yields:

(Eq 3) $$w - w(-1) = \big(p^e - p(-1)\big) + \alpha\big(p(-1) - p^e(-1)\big) + \beta\big(x - \alpha x(-1)\big) + z_w$$

Examination of these equations suggests two cases of interest, depending on the value of the "catch-up" parameter α. If $\alpha = 0$ then, from (Eq 2), the aspiration real wage in each quarter is just the real wage in the previous quarter (adjusted for trend productivity growth, in the empirical version of the model later) plus an error term. Further, from (Eq 3), in the special case of $\alpha = 0$, the change in the nominal wage equals the expected rate of inflation in the previous quarter, $p^e - p(-1)$, plus a term that depends on the degree of labor-market tightness x—in short, a standard expectations-augmented wage Phillips curve.

If $\alpha \neq 0$, on the other hand, then (Eq 3) implies that among the factors determining nominal wage growth is the difference in the previous period's price level and the price level that had been

expected for that period, $p(-1) - p^e(-1)$. We call this term the catch-up term. The existence of a catch-up term implies that, in their bargaining with employers, workers seek to be compensated for last period's unexpected inflation. In other words, to the extent that workers' expectations or focal points (or those of their institutional representatives) influence the wage bargaining process, the economy exhibits a degree of real-wage rigidity. Note also from (Eq 3) that, if $a \neq 0$, the term relating to labor-market tightness can be rewritten as $(1-a)x + a(x - x(-1))$, implying that wages potentially depend on both the level and change of labor-market tightness.

The price equation

Given wages, we assume that the price level p depends on the level of nominal wages plus a shock term z_p that captures the relative costs of nonlabor inputs, variations in markups, and other factors affecting price-setting:

$$p = w + z_p$$

Or, in first differences:

(Eq 4) $$p - p(-1) = (w - w(-1)) + (z_p - z_p(-1))$$

In the empirical model, we will flesh out the shock term by including variables such as commodity price shocks and supply-chain problems that increase prices given wages. Since prices should depend on unit labor costs rather than wages per se, in the estimation we also include the productivity trend in the price equation. We will not treat producer markups as an independent variable but note that markups will endogenously increase in our model when strong sectoral demand confronts limited supply (see below).

Inflation expectations equations

The model is closed by equations that describe the behavior of the public's short-run and long-run inflation expectations. Short-run inflation expectations are a weighted average of long-run inflation expectations, π^*, and last period's inflation:

(Eq 5) $$p^e - p(-1) = \delta\pi^* + (1-\delta)(p(-1) - p(-2))$$

Long-run inflation expectations in turn evolve as a weighted average of last period's long-run inflation expectations and actual inflation:

(Eq 6) $$\pi^* = \gamma\pi^*(-1) + (1-\gamma)(p(-1) - p(-2))$$

The parameters δ and γ capture the degree of anchoring of short and long-run inflation expectations. If both δ and γ are close to 1, expectations are well anchored. Both δ and γ play important though different roles in determining the dynamic effects of shocks on inflation. The parameter δ affects short run dynamics, and the degree to which price inflation affects wage inflation: the lower δ, the more persistent the dynamic effects of a price shock. The parameter γ affects long run dynamics, and in particular the effect of transitory shocks on long run inflation.

Inspection of the equations above shows that, in the long-run steady state, both inflation expectations and the inflation rate consistent with full employment are indeterminate. For any value of long-run expected inflation π^*, assuming that $x = 0$, the model solution implies that short-run inflation expectations, wage inflation, and price inflation all equal the arbitrarily chosen value of π^*. Put another way, the value of π^* at any point in time is determined by the history of inflation. An episode of higher inflation leads to higher steady-state inflation; the longer the episode, or the lower γ, the stronger the effect on steady-state inflation. One can think of the "de-anchoring" that pessimists worried about at the start of the inflation episode as decreases in either δ or γ. Other things equal, such de-anchoring would lead to stronger and more persistent inflation following an inflationary shock.

When the aspiration wage is substituted out, as above, the model determines the evolution of four endogenous variables: π^*, p^e, w , and p. Despite its simplicity, the basic model displays interesting dynamic properties.

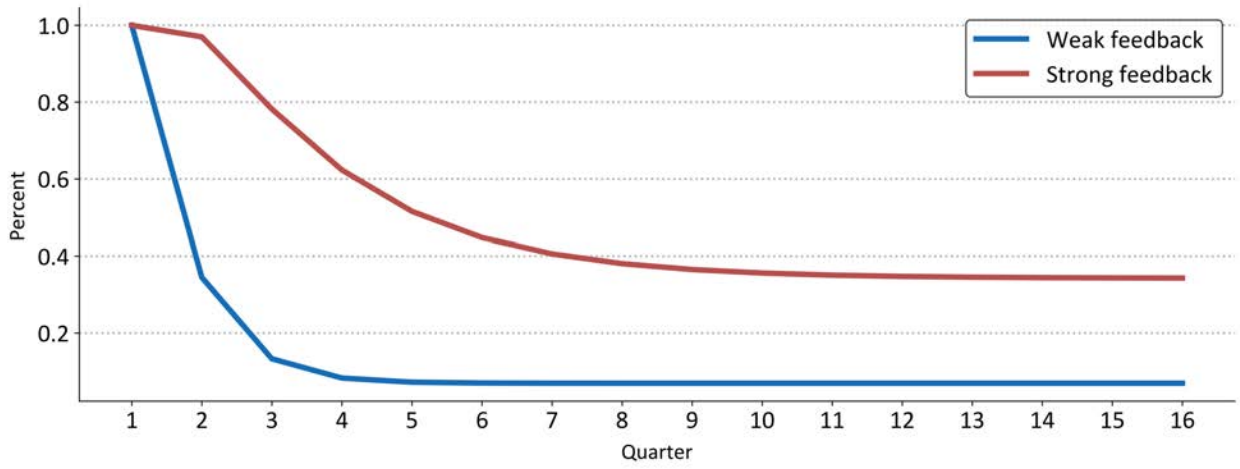
To illustrate, Figure 1 below shows, for alternative parameter assumptions, the model-implied dynamic responses of inflation to a one-time shock to the price equation. Specifically, we simulate the full model under the assumption that the price shock z_p rises permanently by one unit in period 1. (Note that a permanent shock to the price level is a one-period shock to price inflation). We consider two cases: In the first, we assume a limited catch-up effect and stable inflation expectations ($\alpha = 0.2, \delta = 0.9, \gamma = 0.95$).[4] In this case, as Figure 1 shows, the shock to the price level leads to a sharp increase in inflation that is almost completely reversed after a few periods (see the curve marked "weak

[4] We do not need to assign a value to β as doing so would not affect these simulations.

feedback"). The low persistence of inflation in this case reflects the weak catch-up effect (workers are unable to recoup the effects of unexpected inflation on their real wages) and well-anchored inflation expectations. Because of the (small) effect of the inflation shock on long-run expectations, however, inflation ends up permanently slightly higher, by 0.06 percent in this example. With a limited catch-up effect and well-anchored expectations, the conventional wisdom that monetary policymakers can "look through" temporary supply shocks is justified.

The response of inflation to a one-time price shock is quite different when the catch-up effect is stronger and inflation expectations are less well anchored (say, $\alpha = 0.6, \delta = 0.7, \gamma = 0.9$). In this second case the inflation response to a one-time price shock dies down only slowly (see the curve marked "strong feedback" in Figure 1). A stronger catch-up effect implies that a "price-wage spiral" is operating, with inflation triggering higher nominal wage demands, which in turn imply higher prices and greater inflation persistence. In this case, the simulated long-run effect on inflation of a temporary inflation shock, 0.34 percent, is no longer negligible. Again, the simulation is consistent with conventional wisdom at central banks, which holds that poorly-anchored inflation expectations lead to more extended episodes of inflation and to higher inflation in the long run, thereby requiring a strong policy response.

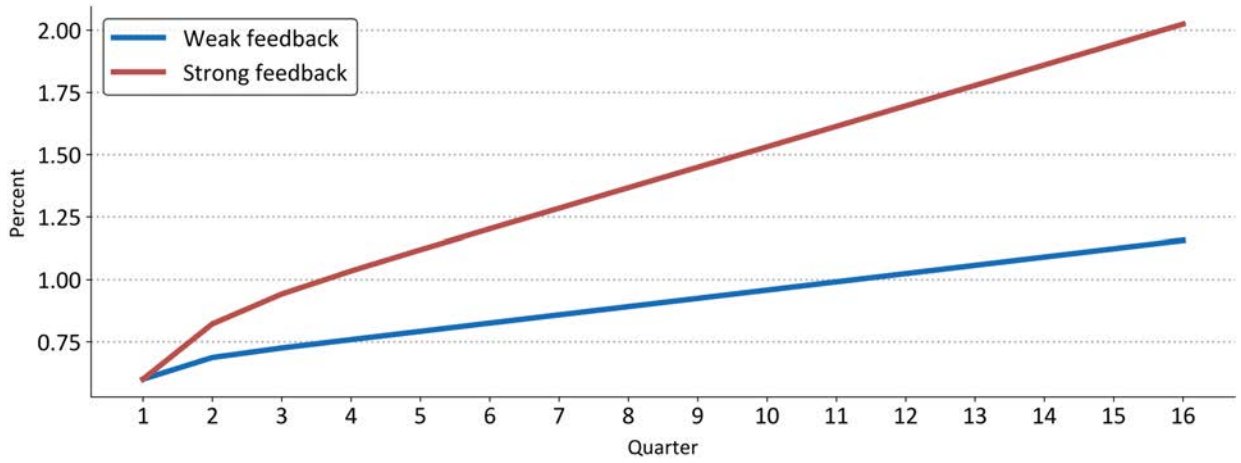
FIGURE 1. RESPONSES OF INFLATION TO A PRICE SHOCK FOR ALTERNATIVE PARAMETER CHOICES



Note: The figure shows inflation in alternative simulations of the model, starting from steady state and assuming a permanent shock to the price level in period 1. Parameter choices that imply less wage catch-up and well-anchored inflation expectations lead to less persistent inflation (weak feedback). Parameters implying greater catch-up and less well-anchored expectations lead to more persistent inflation (strong feedback).

Inflation can also originate in the labor market in our model. For the same two parameter configurations as in Figure 1, Figure 2 shows the simulated behavior of inflation following a one-time, permanent one-unit increase in labor market tightness *x*. After an initial increase, inflation continues to increase over time as the labor market remains tight. The rate at which inflation increases depends on the degree of anchoring of expectations, and to a lesser degree, the strength of the catch-up effect. With well-anchored expectations and little catch-up effect (case 1), a tight labor market results in a relatively slow increase in price inflation (see the curve labeled "weak feedback" in Figure 2.) However, real-wage rigidity and loosely-anchored inflation expectations (case 2) magnify the effects of labor market tightness on wages and, in turn, on prices (the curve marked "strong feedback" in Figure 2). Thus, the longer the overheating episode, the stronger the catch-up effect, and the weaker the anchoring of expectations, the larger is the effect of labor market tightness on inflation, and, implicitly, the stronger the eventual monetary contraction needed to return inflation to target, all else equal.

FIGURE 2. RESPONSES OF INFLATION TO A PERMANENT INCREASE IN LABOR MARKET TIGHTNESS FOR ALTERNATIVE PARAMETER CHOICES



Note: The figure shows inflation in alternative simulations of the model, starting from steady state and assuming a permanent one-unit shock to labor market tightness in period 1. Parameter choices that imply less catch-up and well-anchored inflation expectations lead to a weaker effect of labor market tightness on inflation, while parameters consistent with stronger catch-up effects and less well-anchored inflation expectations lead to a stonger effect.

Figures 1 and 2 yield important conclusions. If, in our model, an economy experiences a large, positive, permanent price shock, it will experience a strong but mostly temporary increase in (headline) inflation. The persistence of that increase depends on the anchoring of expectations and the degree of real-wage rigidity. If, on the other hand, the economy experiences sustained overheating in the labor

market, it will see a slow but steady increase in inflation, with the rate of increase reflecting not only the sensitivity of wages to labor market tightness, but, as in the case of a price shock, on the degree to which inflation expectations are anchored and the extent of real wage rigidity. If the two types of shocks hit the economy simultaneously, the shock to prices will dominate inflation initially. However, as those effects fade, higher underlying inflation emanating from the labor market will become increasingly important.

III. Empirical Implementation of the Model

In this section we discuss the empirical implementation of the model, including our choices of exogenous driving variables, and we summarize the estimation results for each of the four equations. More complete estimation results and results for alternative specifications are presented in the Appendix. The model is estimated using quarterly data. For data availability reasons, and because there is evidence of a break in inflation dynamics around 1990 (e.g., Blanchard, 2016), we begin our estimation in 1990Q1, using a data set that, to accommodate four quarterly lags, begins in 1989Q1.

For our baseline results, in order to see whether pre-covid relations can explain what has happened since covid struck, we end the sample in 2019Q4, thus excluding the covid period (2020Q1-2023Q1), with one exception: In order to include meaningful variation in the effects of sectoral shortages, we estimate the baseline price equation over the full sample, including the covid period. Generally, the results obtained for the wage and inflation expectations equations (as well as for the price equation, excluding the sectoral shortage term) are very similar whether we use the pre-covid period or the whole sample. This similarity suggests that, given the shocks, the conditional process generating prices, wages, and expectations during the pre-covid period remained largely stable during the covid era.

Our estimation strategy is a hybrid approach that approximates a structural vector autoregression (SVAR) with added exogenous variables. In the spirit of SVAR analysis, we imposed restrictions on the contemporaneous relationships among the variables based on our simple model but left the coefficients on lagged variables largely unrestricted.[5] Specifically, in each estimated equation we

[5] In the reported results, we imposed a homogeneity restriction on the estimates that implies that, in the long run, a given change in inflation implies an equal change in short-run and long-run inflation expectations. Alternatively, the restriction can be interpreted as requiring that the long-run Phillips curve be vertical. Under either interpretation, the restriction seems economically justified. The restriction is not rejected at the 5 percent level for

included four lags of each of the endogenous and exogenous variables included in that equation. (We included only one lag of the productivity trend, which is already constructed as an eight-quarter moving average, in the price and wage equations.) Identification of the wage equation is achieved by assuming that wages respond to other variables with a lag of one quarter. Wage inflation then affects inflation, and by implication, inflation expectations contemporaneously.

We turn now to estimation of the model equations. With the estimated parameters in hand, we will be able to evaluate the drivers and dynamics of any of the endogenous variables, including wage and price inflation. For reference, Table 1 below lists the mnemonics and data sources of the endogenous and exogenous variables in the empirical model.

Table 1. Variables and mnemonics used in estimation of the model

Endogenous variables

- *gp* Price inflation, as measured by quarterly annualized rates of change in the Consumer Price Index (CPI), Bureau of Labor Statistics (BLS)
- *gw* Rate of growth of nominal wages, quarterly and annualized, as measure by the rate of change in the Employment Cost Index (ECI), BLS
- *cf1* Short-term inflation expectations, measured by one-year inflation expectations as constructed by the Federal Reserve Bank of Cleveland
- *cf10* Long-term inflation expectations, measured by the Cleveland Fed's ten-year inflation expectations series
- *catch-up* Losses to workers' purchasing power due to inflation, measured by the four-quarter average of CPI inflation minus the one-year inflation expectation four quarters earlier. Catch-up is a linear combination of past *gp* and *cf1*.

Exogenous variables

- *grpe* Rate of growth of the relative price of energy, quarterly and annualized, measured as the rate of change of the ratio of CPI energy prices to the ECI
- *grpf* Rate of growth of the relative price of food, quarterly and annualized, measured as the rate of change of the ratio of CPI food prices to the ECI

the wage and price equations but is rejected at the 2 percent level for the two expectations equations. However, even in the latter case, the estimates under the restriction are very close to the unrestricted estimates. Consequently, the restriction has no material effect on the main results of the paper.

- *v/u* Ratio of job vacancies to unemployment, from the BLS job openings and labor turnover survey (JOLTS) and the BLS Employment Report. Pre-2001 data from Barnichon (2010)
- *shortage* An index of supply chain problems based on Google searches; described below
- *gpty* Trend productivity growth, measured by the eight-quarter moving average of log changes of nonfarm business value added divided by nonfarm employee hours, from the BLS

The wage equation

Following our simple model, we posit that nominal wage growth depends on labor market tightness, short-run inflation expectations, the catch-up effect (the tendency of wages to rise to compensate for past unexpected inflation), and trend productivity growth.

We choose the Employment Cost Index (ECI), wages and salaries component, as our measure of the nominal wage. The ECI is generally thought to be a better measure of wages than the main alternative, average hourly earnings as collected by a monthly BLS survey of employers, because it corrects for changes in average earnings due to changes in the composition of employment. The ECI is available only quarterly, but that is not a problem here since we are using only quarterly data in estimation. Official ECI data begin in 2001Q1. To allow our estimation to start in 1990, we extended the series backward using an older ECI series based on a different BLS survey.[6]

We need a measure of the tightness of the labor market, denoted x in the previous section. Traditionally, the state of the labor market has been measured by the unemployment rate u, or sometimes by the difference between the unemployment rate and the "natural" or sustainable unemployment rate, u^*. The perceived importance of unemployment as a labor-market indicator is reflected in the fact that it is one of only four macroeconomic variables forecast by the FOMC in its *Summary of Economic Projections*.

A conceptual weakness of the unemployment rate as an indicator is that it is based solely on information about the status of households and does not directly incorporate information on the hiring needs of employers. But the labor market involves costly search by employers for workers as well as by workers for jobs. In principle at least, a given unemployment rate could be consistent with either a

[6] The older survey was discontinued in 2005. That series is an index for total compensation for all civilian workers, that is including benefits, seasonally adjusted. The BLS mnemonic is ECS10001I. We thank David Reifschneider for calling this series to our attention. We used a dummy variable to correct for a level shift between the pre-2001 ECI data and the currently official post-2001 series.

strong labor market, with upward pressure on wages, or a weak labor market and low wage pressure, depending on whether job openings are plentiful or scarce.

An alternative indicator of labor market slack that uses information from both workers and employers is the ratio of the vacancy rate (job listings reported by employers divided by the labor force) to the civilian unemployment rate, v/u. Normally u and v/u usually move closely (inversely) together—their correlation in our pre-covid sample is -0.88—and they perform similarly in wage Phillips curves estimated over the pre-covid period (Ball et al., 2022; Cecchetti et al., 2023). However, as has been widely noted, including in Federal Reserve communications, during the covid era these two indicators diverged dramatically. Following a spike during the pandemic lockdowns, the unemployment rate declined gradually. By the end of 2021 it appeared to stabilize at a level modestly below 4 percent, implying—under the FOMC's maintained assumption of a 4 percent natural unemployment rate—that labor market conditions, though moderately tight, were not becoming significantly tighter. Moreover, the unemployment rate in late 2021 was comparable to or slightly higher than that just prior to the pandemic, a period without evident inflation pressures. The signal from the unemployment rate was consequently that labor market overheating was not an imminent threat to price stability.

The picture portrayed by the vacancy-to-unemployment ratio during the covid era was quite different. That ratio rose from less than 1.0 in April 2021 (that is, less than one open job per unemployed worker) to a historically high level of about 1.9 (nearly two job openings per unemployed worker) a year later. This increase implied an exceptionally rapid pace of labor-market tightening, to a level considerably greater than during the immediate pre-covid period and, indeed, any period since these data have been collected. As of this writing, v/u has declined only moderately from its peak, consistent with a labor market that remains overheated.

Which indicator of labor market tightness is better? As their high correlation and comparable performance in Phillips curves suggests, in normal times the two indicators give very similar assessments of labor market conditions. However, the relationship between vacancies and unemployment changed materially during the covid period, with more vacancies (more search effort by employers) needed to achieve a given rate of hiring, holding constant the number of unemployed (Blanchard, Domash, and Summers 2022). Equivalently, the Beveridge curve, the downward relation between vacancies and unemployment, has shifted upward since 2020.[7] If labor market tightness, as reflected in workers'

[7] Many factors no doubt contributed to this shift. A high rate of firm closures implied that many laid-off employees had to find new jobs rather than return to their old ones. Shifts in consumer demand patterns and worker preferences induced a reallocation of jobs among firms and sectors, slowing the matching process. Changes in

bargaining power, depends on the number of jobs available relative to the number of people looking for work, then the ratio of vacancies to unemployment remains a good measure of tightness even when employers must search more intensively to find potential hires. In contrast, given the number of unemployed, workers' bargaining power and labor market tightness is much greater when there are many vacancies to be filled than when vacancies are scarce. We conclude that v/u is a more reliable indicator of labor market conditions in times, as during the pandemic and its aftermath, in which the efficiency of the employer-worker matching process has changed materially.[8] For this reason, and following the direction taken in much recent research and in Federal Reserve policy communications, we use v/u as our empirical measure of labor market pressure.[9]

To capture short-term inflation expectations, we use one-year inflation expectations as constructed by the Federal Reserve Bank of Cleveland. This expectations series is partly retrospective and model-based, which is a drawback, but it has the advantage of combining several types of real-time information, including data from financial markets and from surveys. We have also estimated our model using forecasts of the Survey of Professional Forecasters (SPF) as proxies for inflation expectations (see the Appendix for the results).[10] When we use this alternative measure, the differences for both the wage equation and the other equations are minor, and the conclusions we draw in the text do not change. We comment on some results for alternative data choices and sample periods below.

We construct our measure of the wage catch-up term as the difference between realized CPI inflation over the current and past three quarters and the one-year expectation of inflation as of four quarters earlier. The catch-up variable thus equals unexpected inflation, which we hypothesize may affect subsequent wage bargains if workers are focused on maintaining the real buying power of their wages. As has been noted, our measure of labor productivity is an eight-quarter moving average of nonfarm productivity growth.

---

family arrangements, illness, and other effects of the pandemic lockdowns may have reduced the (unobserved) intensity of search effort of some unemployed workers. And employers may have failed to appreciate the extent to which, because of concerns about returning to jobs involving personal contact or perhaps because of reassessments of work-life balance, workers had increased their reservation wages.

[8] Like the vacancy-to-unemployment ratio, the quit rate also reflects both labor demand and labor supply, in that the willingness of workers to quit their current jobs depends on their perceptions of labor-market conditions. We found however that the quit rate does not perform as well as v/u in our wage equation.

[9] For more on the Beveridge curve, and a formalization of the argument that the vacancy-to-unemployment ratio should be used as a measure of labor market tightness, see Blanchard and Diamond (1989).

[10] The SPF forecasts are made in real time, an advantage. However, they raise some issues of timing, e.g., ten-year forecasts in each quarter are for the subsequent decade including the current year, rather than for the ten years beginning in the forecast quarter. This feature leads to discontinuities between fourth-quarter and first-quarter ten-year forecasts.

As also discussed above, we want to be guided by our formal model while maintaining flexible lag structures. We make the identifying assumption that wages respond only to lagged values of the independent variables, and we regress wage growth *gw* on a constant, four (quarterly) lags of itself, four lags of short-run inflation expectations *cf1*, four lags of the unexpected inflation term *catch-up*, four lags of our proxy for labor market tightness v/u, and one lag of trend productivity growth *gpty*.

Table 2 below reports, for each right-side variable in the wage equation, the included lags, the sum of the estimated coefficients for each variable and two p-values. The first p-value, denoted *p-stat (sum)*, is the probability of rejection of the null hypothesis that the sum of current (if applicable) and lag coefficients for a given variable is zero. The second p-value, *p-stat (joint)*, tests the joint hypothesis that each of the current and lag coefficients is separately zero. To ensure that the long-run wage Phillips curve is vertical, we impose a homogeneity assumption (see footnote 5 for more discussion), which here implies that the sum of coefficients on lagged nominal wage growth and lagged short expectations is equal to one.

Table 2. Wage growth regression: Dependent variable = *gw*

Independent variable	gw	v/u	catch-up	cf1	gpty
Lags	-1 to -4	-1 to -4	-1 to -4	-1 to -4	-1
Sum of coefficients	0.460	0.693	-0.024	0.540	0.031
p-stat (sum)	0.008	0.030	0.765	0.002	0.608
p-stat (joint)	0.071	0.023	0.994	0.022	0.608
R-squared			0.578		
No. observations			120		

Notes: Sample is 1990Q1 to 2019Q4. *p-stat (sum)* is the p-value for the null hypothesis that the sum of coefficients is zero, *p-stat (joint)* is the p-value for the joint hypothesis that each of the lag coefficients separately equals zero.

Table 2 suggests that, when v/u is used as the measure of slack, the wage Phillips curve remains alive and well, in that the estimated effect of labor market conditions on wage growth is significant, both statistically and economically. Based on the sum of estimated coefficients for the four quarterly lags of the vacancy-to-unemployment ratio, and the sum of the coefficients on lagged nominal wage growth, an increase in v/u from, say, 1.0 to 1.5, holding constant other factors, has a long-term effect on nominal wage growth of 0.5 x 0.693/(1-.460), or 0.64 percentage points.[11] When we use the SPF 1-

[11] Note the importance of "holding constant other factors". As we saw in the simple model, and as is the case in the empirical model as well, in general equilibrium, letting other factors adjust leads to steadily increasing inflation.

year forecast as a proxy for short-run inflation expectations, the effect of v/u is marginally larger, 0.75, and highly statistically significant.

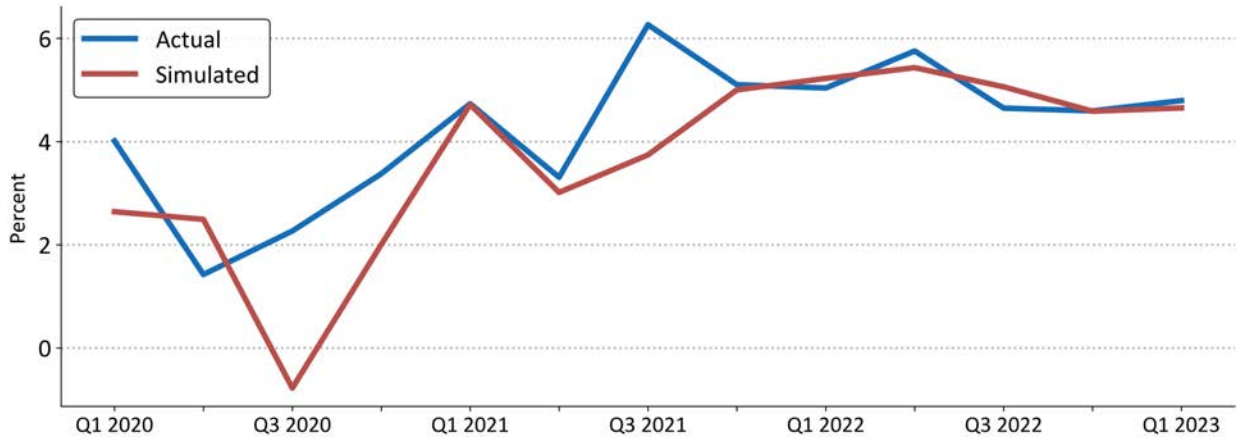
The sum of the coefficients on short-run inflation expectations is high relative to the sum of coefficients on lagged nominal wage growth (under the homogeneity assumption, the sums of the coefficients on the two variables add to one). This estimate suggests that short-run inflation expectations feed fairly quickly into wages, although, as we shall see below, short-run inflation expectations react slowly to actual price inflation.

There is no evidence of a catch-up effect in Table 2, a result that appears robust. For example, when we use SPF 1-year forecast to measure short-run inflation expectations, the catch-up effect is estimated to be small and positive, though not statistically significant, with sum of coefficients equal to 0.1. When we estimate the equation over the full sample, for either measure of short-term expectations, the estimated catch-up effect remains small, and is again statistically insignificant.

The moving-average trend in productivity growth has no statistically significant effect on wage growth, suggesting that some of the effect of productivity growth may be captured by the constant term of the equation.

As a check on our specification, we can see how the wage equation estimated on the pre-covid sample fits wage data during the covid period (2020Q1-2023Q1). As Figure 3 shows, the estimated wage equation predicts wage behavior after 2019 reasonably well, including the upward trend in wage growth over the period. The fitted wage equation's largest miss is for 2020Q3, when the equation predicts a larger decline in wage growth than actually occurred. Given the tremendous deterioration in labor market conditions in the early stages of the pandemic, it is not surprising that the equation predicts a commensurate decline in wages in that quarter. Downward nominal rigidity in wages and pervasive measurement issues likely explain most of the miss.

FIGURE 3. WAGE GROWTH, 2020Q1-2023Q1, COMPARED TO PREDICTIONS OF WAGE EQUATION



Note: The figure shows actual and predicted nominal wage growth, 2020Q1-2023Q1, based on the wage equation estimated on the pre-covid sample.

The price equation

We turn now to the determinants of price inflation. In our model, price inflation depends on nominal wage growth and on other factors that affect prices given wages, such as growth in other input costs. Our estimated equation includes four lags of inflation and the current value and four lags of wage growth. We use the more familiar and timelier Consumer Price Index (CPI) as our measure of the price level. The Appendix shows that measuring prices using the personal consumption expenditures (PCE) deflator leads to very similar results. We continue to use the ECI as our measure of nominal wages.

A substantial part of the pandemic-era inflation surprise reflected rapid, largely unanticipated increases in food and energy prices. We consequently included in our specification contemporaneous values and four lags each of the growth rates in food (*grpf*) and energy prices (*grpe*), as defined by the CPI, with both measured relative to the contemporaneous growth rate in nominal wages. Shocks to these variables should thus be interpreted as shocks to relative rather than absolute prices. Since lags of nominal wage growth are included in the estimated equation, however, the fit of the equation would be the same independent of whether the growth of input prices enters in relative or absolute terms.

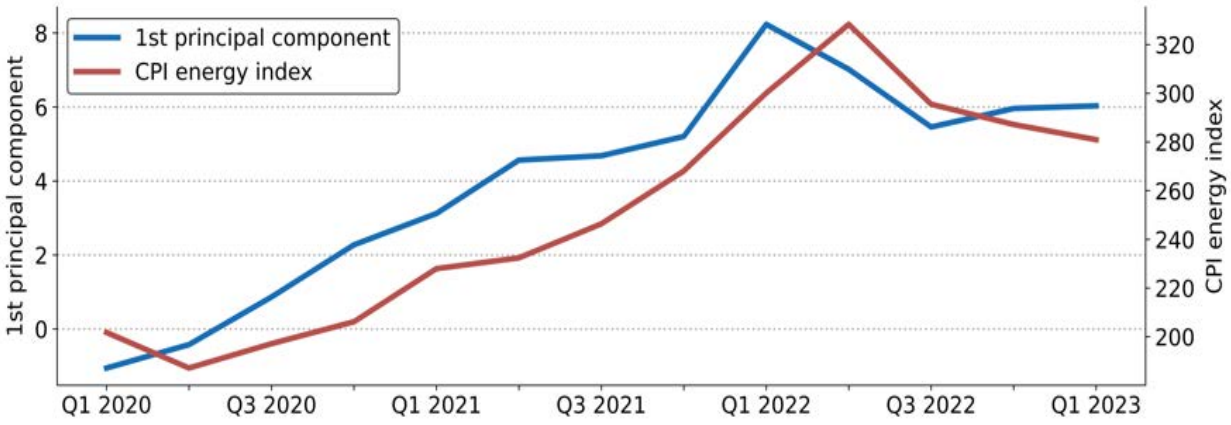
Simply treating relative food and energy prices as exogenous in the price equation, per conventional practice, is unsatisfying, as it leaves unanswered the question of why these prices moved up so sharply and, apparently, in a roughly synchronized way. To gain some insight into this question, we performed the following experiment.

We hypothesized that common movements in the prices of a range of disparate commodities were more likely to reflect the influence of global demand changes (the result of global macroeconomic

policies, reopening from the pandemic, and large geopolitical events like the Ukraine war) rather than changes in supply, as conditions affecting the supplies of diverse commodities, such as local weather conditions or the opening of new mines, are likely to be largely idiosyncratic. To find the common trend in commodity prices, we used a principal components approach.

Specifically, we collected monthly prices for 1990 to the present of the 19 commodities included in the familiar Commodity Research Bureau (CRB) commodity price index.[12] We standardized each commodity price series to have mean zero and a standard deviation of one. We calculated the first principal component of these price series, finding that it explained 66.0 percent of the overall variance. The two figures below show the estimated principal component (converted to a quarterly frequency), together with the CPI energy (Figure 4) and food price (Figure 5) components since 2020Q1.[13]

FIGURE 4. CRB PRICES PRINCIPAL COMPONENT AND ENERGY PRICES, 2020Q1-2023Q1

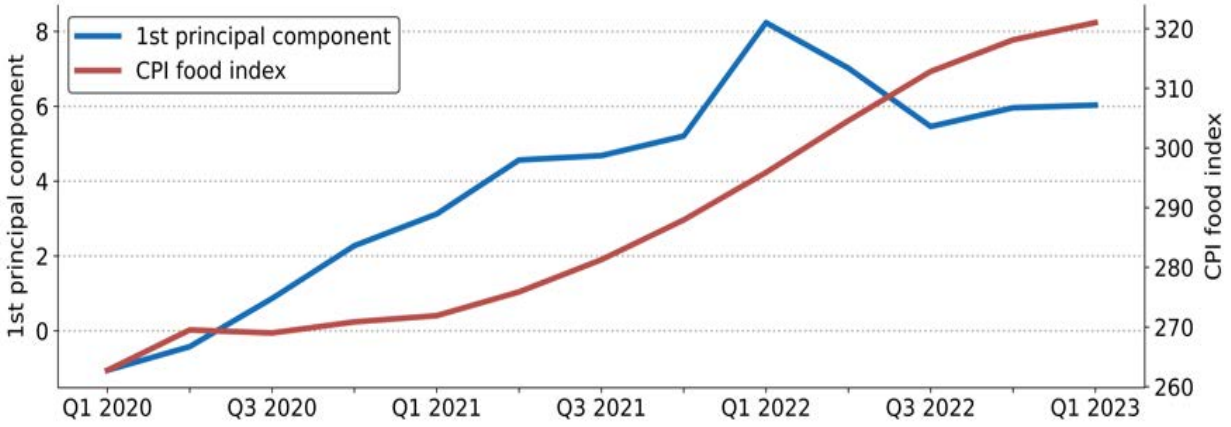


Note: The figure shows the first principal component (common trend) of the prices of 19 commodities currently included in the CRB commodity price index, together with the CPI price of energy, in levels (1982-84 = 100).

[12] The CRB index currently comprises prices of 19 commodities, including energy, metals, and agricultural commodities. We gathered prices of commodities currently included in the index, ignoring changes since 1990 in the composition of the index. Commodity prices were weighted equally in our principal components exercise. R^2s of regressions of each commodity price on the principal component, using monthly data from 2020M1, were high for energy, food, and metals, with a few exceptions, notably gold ($R^2 = 0.08$), cocoa ($R^2 = 0.0$), and orange juice ($R^2 = 0.41$).

[13] As noted in the text, Figures 4 and 5 show the principal component of all the commodity prices in the CRB index, including energy and food prices. To avoid the possibility of a spurious correlation between CPI energy and food prices and the principal component, we repeated the exercise excluding from the principal components analysis, in sequence, (i) all energy and food commodities (ii) all energy commodities, and (iii) all food commodities. The first principal components extracted from these reduced lists of commodity prices looked quite similar to that shown in Figures 4 and 5.

FIGURE 5. CRB PRICES PRINCIPAL COMPONENT AND FOOD PRICES, 2020Q1-2023Q1



Note: The figure shows the first principal component (common trend) of the prices of 19 commodities currently included in the CRB commodity price index, together with the CPI price of food, in levels (1982-84 = 100).

Figures 4 and 5 present a mixed picture. Energy prices generally moved in line with the common component, rising for a period, spiking in early 2022, and then declining. Food prices followed the common component upward as well but continued to rise after the common component peaked and began to decline. Our interpretation is that in the earlier part of the pandemic, energy and food prices were indeed driven primarily by demand factors, including stimulative U.S. fiscal and monetary policies and the reopening of many economies from covid lockdowns. The Russian invasion of Ukraine that began on February 24, 2022, caused a further spike in commodity prices, which waned as the perceived threat to global commodity supplies moderated. That story is consistent with the behavior of energy prices. However, the continued increase in food prices following the sharp decline in the prices of raw commodities, a phenomenon not confined to the United States, bears further investigation. Since retail food prices reflect many inputs beyond raw commodities, including labor, changes in non-commodity input prices and wages could be one reason for this divergence. Non-competitive behavior is another possibility, as the food price index includes distribution margins, and the retail food industry has become more concentrated in recent decades (Zeballos, Dong, and Islamaj, 2023).

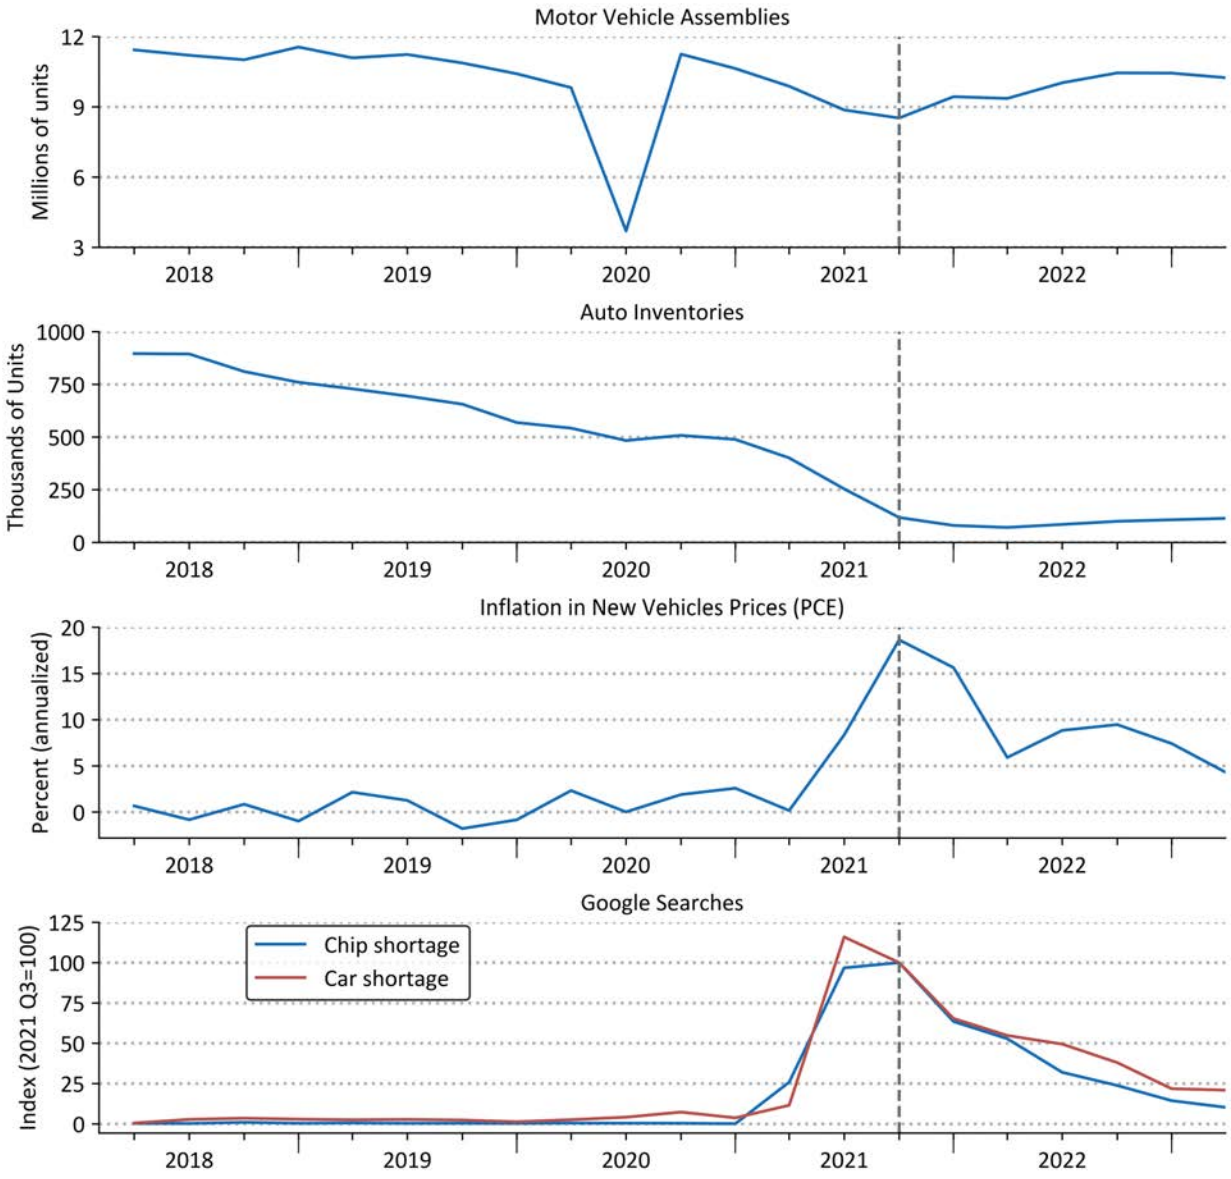
Besides commodity prices, pervasive shortages of certain goods have often been cited as a source of inflation during the pandemic period (Comin, Jones, and Johnson 2023). These shortages and disruptions have reflected developments in both demand and supply. On the demand side, lockdowns and fears of infection led to a shift in consumer spending from in-person services to goods, especially durable goods. The share of nominal consumer spending devoted to durable goods rose from about 10

percent in the years before the pandemic to about 13 percent once the pandemic began, an increase of nearly a third, and it remained high. This sectoral demand shift was augmented by increases in overall consumer spending driven by supportive monetary and fiscal policies. On the supply side, global supply chains and transportation networks were disrupted by pandemic-induced shutdowns, constraining production in key sectors, and even those industries not affected by such problems were in many instances pushed to the limits of their capacity. The combination of an increased demand for goods and limited supply resulted, not surprisingly, in sharp rises in prices in a number of sectors.

One might argue that increases in relative demand in some sectors, and associated higher prices, should have been largely offset by decreases in relative demand in other sectors, and associated lower prices. Shortages however introduce a fundamental asymmetry. If we think for example of firms facing a flat marginal cost up to some physical constraint where supply becomes vertical, then relative demand increases, if large enough, will move prices along the vertical supply curve, but the corresponding relative demand decreases will move prices along the horizontal portion of the supply curve. The increase in prices in some sectors will not be offset by a decrease in prices in the others. This effect will be stronger, the larger the variation in relative demand or in relative supply constraints. Both factors seemed to have played a role in the past three years.

The automobile sector, an important contributor to the early surge in inflation, illustrates the role of demand and of supply constraints (see Figure 6 below). U.S. vehicle production plunged during the initial lockdowns then recovered as factories reopened. Assemblies ran at about an 11.7 million annual rate in July 2020, comparable to or slightly higher than pre-pandemic production rates (top panel of Figure 6). However, persistent shortages of critical chips and other key components subsequently restrained vehicle production, which fell below an annual rate of 9 million units—a decline of about a quarter—by the fall of 2021. The other panels of the figure show that the reduced production did not reflect a lack of potential customers. Inventories of new cars for sale fell to record low levels as dealers could not keep up with demand (second panel of Figure 6) and, consequently, new automobile prices spiked (third panel). The bottom panel of Figure 6 shows an index of U.S.-based Google searches for "chip shortage" and "car shortage," both of which rose sharply in 2021. The vertical dashed line in the figure at 2021Q3 illustrates the close relationship between these developments. That quarter saw the trough in automobile production, the decline in inventories to record low levels, the spike in auto prices, and (a quarter earlier) the peak in Google searches, consistent with a shortage-induced price jump.

FIGURE 6. PRODUCTION AND SHORTAGES IN THE U.S. AUTOMOBILE INDUSTRY, 2018Q1-2022Q4



Note: Motor vehicle assemblies data pertain to the United States, inventories represent inventories held in the U.S. of vehicles assembled in North America. Google searches are measured by an index, 2021Q3=100. Sources include the Federal Reserve Board of Governors, Bureau of Economic Analysis, Google Trends.

As the figure suggests, the Google search variables appear to be good proxies for the sectoral demand shifts and supply constraints, at least in one very important industry. (Shortages of new cars spilled over to the prices of used cars, which at times were also important drivers of inflation.) This led us to use an index of the quarterly average of the number of U.S.-based Google searches for the term

"shortage" in our price equation (the underlying series is weekly, and is normalized by Google to equal 100 in the week when it reaches its highest value, May 2 to 9, 2021) as a proxy for the combined effects of the shift in demand towards durable goods and the disruptions of and limits on supply.[14] We experimented with several other measures of supply chain and transportation problems used in the literature, including the Federal Reserve Bank of New York's supply chain index, the Baltic Dry Index, supplier delivery times as reported by the Institute for Supply Management, the ratio of durable goods to services spending, and others, but found that the Google search variables provided, by far, the most explanatory power.

The collision of high demand and limited supply in some sectors can account for at least some of the increase in markups observed during the pandemic period. While other factors no doubt have influenced markups, including for example the fiscal transfers that directly affected demand in product markets, at least in this simple specification we do not find that including these factors is needed to explain the behavior of pandemic-era inflation (see Figure 7 below).

A summary of estimation results for the price equation is given in Table 3. As before, the Appendix provides more complete results, as well as estimation results using the PCE deflator rather than the CPI as the measure of prices.

Table 3. Price inflation regression (dependent variable = *gp*)

Independent variable	gp	gw	grpe	grpf	shortage	gpty
Lags	-1 to -4	0 to -4	0 to -4	0 to -4	0 to -4	-1
Sum of coefficients	0.335	0. 665	0.066	0.126	0.018	-0.143
p-stat (sum)	0.037	0.000	0.000	0.050	0.281	0.026
p-stat (joint)	0.066	0.000	0.000	0.050	0.000	0.026
R-squared			0.947			
No. observations			133			

Note: Sample is 1990Q1 to 2023Q1. *p-stat (sum)* is the p-value for the null hypothesis that the sum of coefficients is zero, *p-stat (joint)* is the p-value for the joint hypothesis that each of the current and lag coefficients separately equals zero.

[14] The normalization based on weekly data is different from the quarterly normalization we use in Figure 6 (this difference has no consequence for the estimation results). The Google data start in 2004Q1. For earlier observations, to avoid an artificial discontinuity, we assume that *shortage* takes a constant index value of 5, approximately its pre-pandemic average.

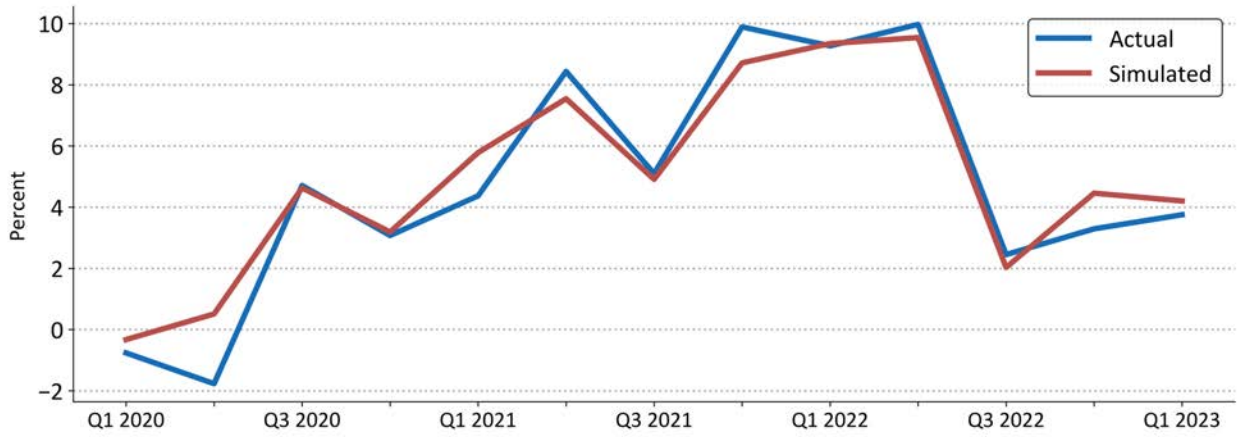
The table reports the results of a regression of price inflation on a constant, four lags of itself, the current value and four lags of nominal wage growth, and the current values and four lags each of inflation in the relative prices of energy and food. Also included as independent variables were the current value and four lags of Google searches for *shortage*, and one lag of trend productivity growth, which should affect unit labor costs and thus prices. Recall that, in order to include the *shortage* variable, which is roughly constant and small pre-covid, the inflation equation, unlike the other equations in our model, is estimated for a sample that includes the covid period. Estimation of the price equation in the pre-covid sample, excluding *shortage*, produced similar estimates for the coefficients of the other variables in the equation. The estimated equation reflects the imposed homogeneity assumption that the sum of the coefficients on past price and on wage inflation sum to one.

In Table 3, the large sum of coefficients on nominal wage growth relative to the sum of coefficients on lagged inflation implies a fairly rapid passthrough of wages to prices. The relative prices of energy and food also affect inflation, as expected, though the sum of the coefficients on the relative price of food is only marginally statistically significant. The estimated long-run effect of an energy price shock on the price level, other things equal, is 0.066/(1-0.335) = 10 percent, so about a third larger than energy's share in the CPI basket (7.4 percent). The estimated long-run effect of a food price shock is equal to 0.126/(1-0.335) = 19 percent, so about 30 percent larger than the share of food in the basket (14.4 percent). Since wages are controlled for, these findings suggest the presence of some second-round price-price effects of energy and food price shocks on the prices of other goods and services over time.

The sum of the coefficients on the shortage variable is not statistically significant, though the joint hypothesis that the lag coefficients are each zero is strongly rejected. We interpret these findings as suggesting that shortages have a strong but temporary effect on the price level. Trend productivity growth enters with the expected negative sign (higher productivity reduces unit labor costs, given wages), with an estimated coefficient that is statistically significant at the 5 percent level.

As we did for wage growth above, we can use this estimated equation to predict the path of inflation over the covid period. As Figure 7 shows, the predicted inflation series is quite close to the actual covid-era data.

FIGURE 7. INFLATION, 2020Q1-2023Q1, COMPARED TO PREDICTION OF PRICE EQUATION



Note: The figure shows predicted and actual inflation, 2020Q1-2023Q1, based on the price equation estimated over the full sample.

Inflation expectations equations

The final equations to be estimated are those determining short-run and long-run inflation expectations. Inflation expectations can be measured in many ways, including through financial-market indicators, consumer and business surveys, and the expectations of professional forecasters. For our baseline results we use the one-year and ten-year inflation expectations series constructed by the Cleveland Fed. The Appendix discusses results obtained when we use measures of inflation expectations based on the Survey of Professional Forecasters.

Begin with the equation for short-term inflation expectations (*cf1*). Our model posits that short-run inflation expectations in each period are a weighted average of long-term inflation expectations and realized inflation. In the application we also include lags of short-term inflation expectations, which allows for the possibility of gradual adjustment. We accordingly estimate an equation that has short-run inflation expectations depending on four lags of itself, the current value and four lags of our measure of long-run inflation expectations, *cf10*, and current and four lags of actual inflation. Table 4 below summarizes the results.

Table 4. Short-run inflation expectations (dependent variable = *cf1*)

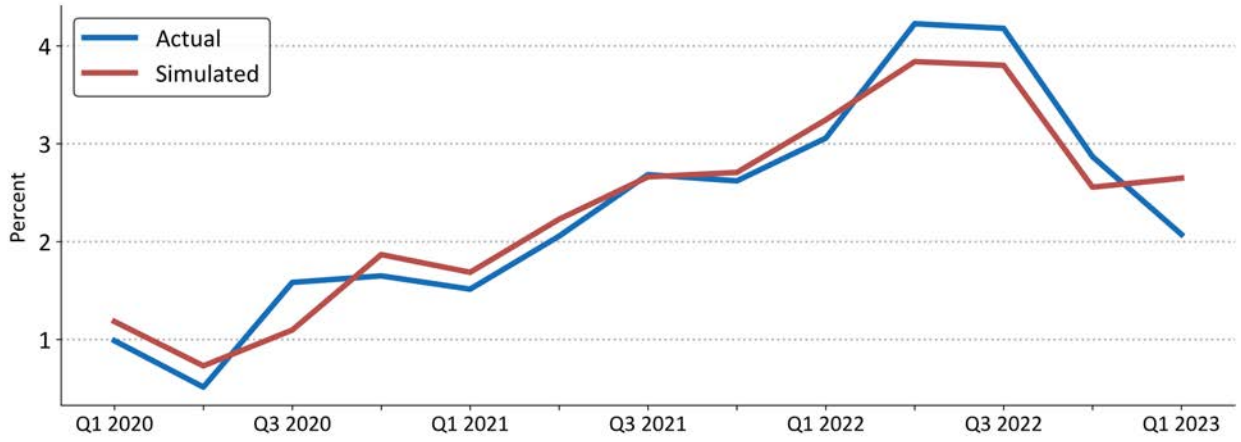
Independent variable	cf1	cf10	gp
Lags	-1 to -4	0 to -4	0 to -4
Sum of coefficients	0.369	0.506	0.124
p-stat (sum)	0.014	0.000	0.001
p-stat (joint)	0.001	0.000	0.000
R-squared		0.908	
No. observations		120	

Note: Sample is 1990Q1 to 2019Q4. *p-stat (sum)* is the p-value for the null hypothesis that the sum of coefficients is zero, *p-stat (joint)* is the p-value for the joint hypothesis that each of the current and lag coefficients separately equals zero.

We again impose homogeneity, which in this case implies that the sum of the coefficients of the three variables equal one. (The restriction has little effect, as the unconstrained sum of coefficients is 0.99.) As noted, this homogeneity restriction, together with a similar restriction on the long-run inflation expectations equation, ensures that, in the long run, a sustained increase in inflation will raise inflation expectations by the same amount. In the short run, however, that relation need not hold, and the fact that current and lagged inflation enter this equation with a small sum of estimated lag coefficients suggests that short-term inflation expectations are relatively well-anchored—that is, all else equal, they are only modestly affected by near-term changes in inflation. However, given the high autocorrelation of short-run inflation expectations, the estimates also imply that sustained inflation gets increasingly reflected in short-run expectations over time.

Like the previous equations, the estimated equation for short-run inflation expectations does a good job of predicting pandemic-era inflation expectations out of sample, as shown by Figure 8. There is no evidence of de-anchoring, defined here as an increase in the short-run effect of inflation on expected inflation: If estimated over the whole sample, including the covid period, estimated coefficients are very similar. For example, the sum of coefficients on inflation is 0.133 in the full sample, compared to 0.124 in the pre-covid sample alone.

FIGURE 8. SHORT-RUN INFLATION EXPECTATIONS, 2020Q1-2023Q1, COMPARED TO PREDICTION OF SHORT-RUN EXPECTATIONS EQUATION



Note: The figure shows predicted and actual short-run inflation expectations, 2020Q1-2023Q1, with predicted values based on the associated equation estimated on the pre-covid sample.

Finally, Table 5 shows our estimated equation for long-run inflation expectations (*cf10*). Following our modeling assumption that long-run inflation expectations evolve as a weighted average of the previous level of long-run expectations and actual inflation, in the estimated equation we allow long-run inflation expectations to depend on four lags of itself and the current value and four lags of realized inflation. The results echo those for short-run inflation expectations. We impose the homogeneity restriction that the sum of coefficients is equal to one. The unconstrained sum of coefficients is 0.99, so once again the homogeneity restriction makes no substantial difference to the results.

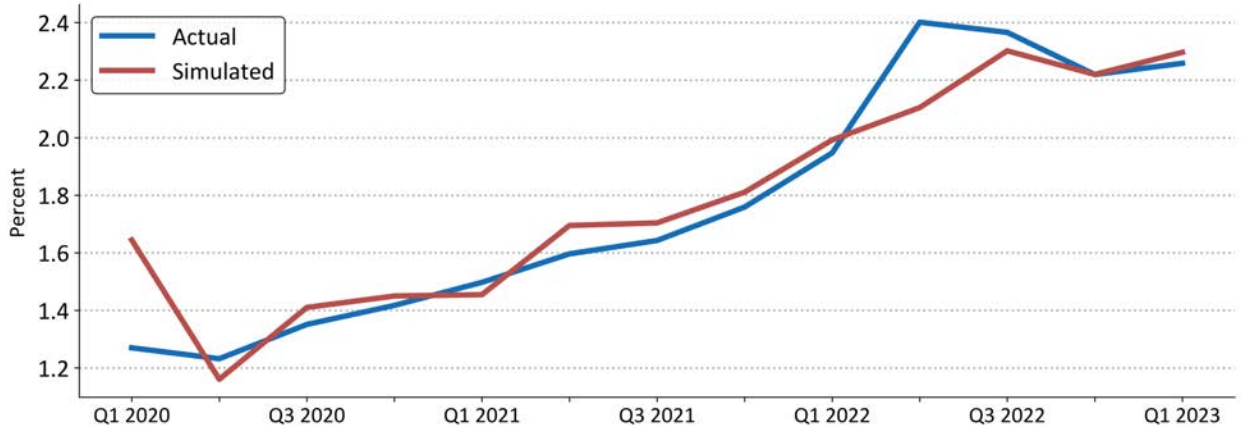
The response of long-run inflation expectations to short-run movements in prices is small, suggesting that long-run expectations are well anchored. In particular, there is no evidence that long-run expectations became de-anchored during the covid period. The sum of coefficients on current and lagged inflation were not much changed, being estimated as 0.031 in the whole sample, compared to 0.025 in the pre-covid sample. The high estimated autocorrelation of long-run inflation expectations implies that a sustained increase in inflation will lead to slowly but steadily rising long-term inflation expectations, which matches the pattern in the covid-era data. See Figure 9.

Table 5. Long-run inflation expectations (dependent variable = *cf10*)

Independent variable	cf10	gp
Lags	-1 to -4	0 to -4
Sum of coefficients	0.975	0.025
p-stat (sum)	0.000	0.208
p-stat (joint)	0.000	0.004
R-squared	0.931	
No. observations	120	

Note: Sample is 1990Q1 to 2019Q1. *p-stat (sum)* is the p-value for the null hypothesis that the sum of coefficients is zero, *p-stat (joint)* is the p-value for the joint hypothesis that each of the current and lag coefficients separately equals zero.

FIGURE 9. LONG-RUN INFLATION EXPECTATIONS, 2020Q1-2023Q1, COMPARED TO PREDICTION OF LONG-RUN EXPECTATIONS EQUATION



Note: The figure shows predicted and actual long-run inflation expectations, 2020Q1-2023Q1, with the prediction based on the associated equation estimated over the pre-covid sample.

Impulse Response Functions

The estimates of this section provide parameter values for the empirical version of our model. The model determines four endogenous variables (price inflation, wage inflation, and short-run and long-run inflation expectations) given initial conditions and five exogenous variables (the vacancy-to-unemployment ratio, the relative prices of food and energy, the shortage indicator, and trend productivity.)

The estimated model can be used to calculate impulse response functions (IRFs) that take into account both the simultaneity and lagged relationships among wages, prices, and inflation expectations. For example, we can estimate the dynamic effects of a shock to a particular equation or of a change in any of the exogenous variables, taking a general equilibrium and dynamic perspective. Estimated impulse responses using the full model capture not only the direct effects of a given shock (e.g., as

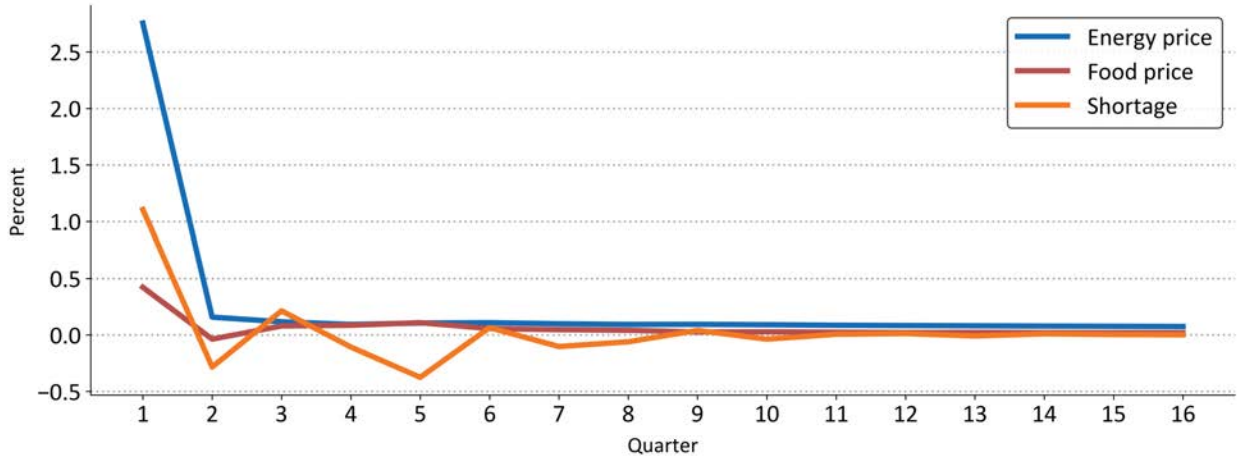
reflected in the estimated coefficients of the shocked variable in each of the equations) but also the knock-on effects in all subsequent periods. For example, an increase in the relative price of energy affects price inflation directly, but it may also motivate workers to try to maintain their real wages (the catch-up effect) and lead to increases in both short-run and long-run inflation expectations, each of which feeds into wage inflation and subsequently into price inflation over time.

In Section 2, using the simple model and starting from the steady state, we showed the responses of price inflation to (i) a positive, permanent shock to the price equation (Figure 1) and (ii) a positive, permanent shock to labor market tightness (Figure 2). We conducted these exercises using two sets of hypothetical parameters, in the process confirming the intuition that the persistence of inflation in the face of a shock is typically less when the catch-up effect is small (workers are unable to regain much of the purchasing power lost to unexpected inflation) and inflation expectations are well anchored. With the full model, now including estimated parameters and unrestricted lags, we can repeat those exercises. As before, we start from a steady state and consider one-standard-deviation positive shocks to each of three price shock variables—the relative price of energy, the relative price of food, and the shortage index—and to the vacancy-to-unemployment ratio. The standard deviations used for the shocks are calculated using data from 2020Q1 to 2023Q1.[15] The results are given in Figure 10 below for the three price shocks and in Figure 11 for the vacancy-to-unemployment ratio.

Figure 10 is striking, in that the effects of the three price shocks on inflation are quite short-lived, even more so than in the simulation of the analytical model under weak feedback parameters (Figure 1). The effects mostly disappear in the quarter following the shock, although the short burst of inflation together with not fully anchored inflation expectations yields, in the case of an increase in the price of energy, a small sustained increase in inflation of about 0.1 percent. This long-term effect reflects the fact that, as discussed earlier, inflation is indeterminate in this model, and its steady-state level depends on the behavior of inflation in the past.

[15] The standard deviations are respectively 32 percent for the quartely rate of change of the relative price of energy (which varied for 45 percent to -52 percent during the period) at an annual rate, 3.2 percent for the rate of change of the relative price of food (which varied from 8.9 percent to -1.3 percent), 11.6 for the shortage index (which varied from 10 to 48), and 0.5 for the vacancy-to-unemployment ratio (which varied from 0.26 to 1.91).

FIGURE 10. RESPONSES OF INFLATION TO SHOCKS TO THE RELATIVE PRICE OF ENERGY, THE RELATIVE PRICE OF FOOD, AND THE SHORTAGE INDEX



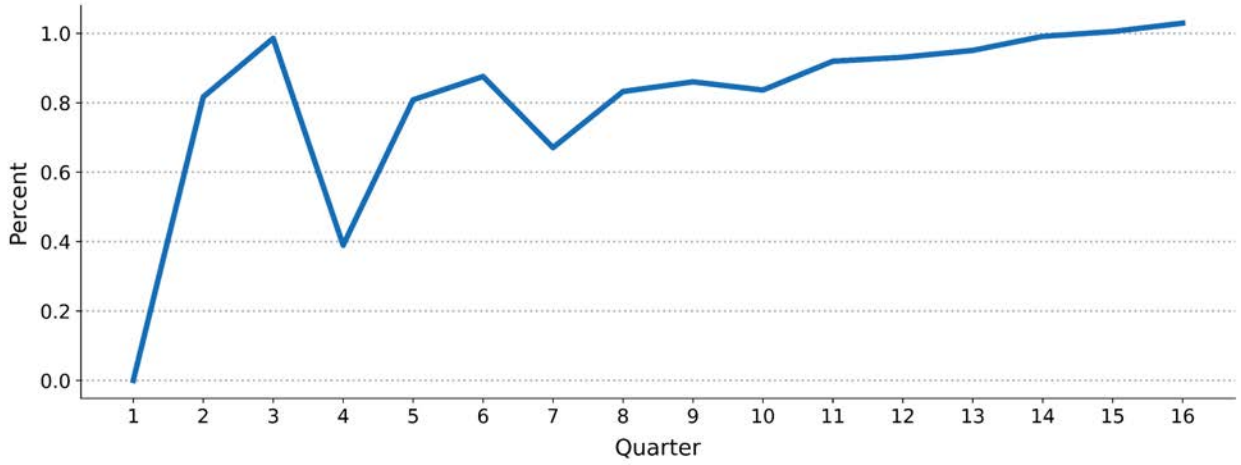
Note: The figure shows the full-model, dynamic responses of inflation to a one-standard-deviation positive shock to relative energy prices, relative food prices, and the shortage variable. It may be compared to the hypothetical responses of inflation under alternative parameter assumptions in Figure 1.

The short-lived responses of inflation to various price shocks shown in Figure 10 reflect our findings that real wage rigidity (the catch-up effect) is limited, and that short-run and long-run inflation expectations appear to be reasonably well anchored, adjusting very slowly in response to actual inflation. This anchoring likely reflects the credibility of the Fed's commitment to return inflation to its target level over time. These results are good news for policy makers, as they suggest that, as price shocks fade, their effects on inflation quickly (although not fully) disappear. These results are also in sharp contrast to what happened in the 1970s, a time when inflation expectations were much less anchored, and because of automatic cost-of-living adjustments and other mechanisms, catch-up effects were presumably much stronger. As a result, inflation due to price shocks (notably oil and food price shocks, in the 1970s) was very persistent, despite the fact that on the whole the labor market was not particularly tight during that decade.

Figure 11 shows a full-model simulation of how inflation responds over time to a persistently tight labor market, modeled as a permanent one-standard deviation increase in the ratio of vacancies to unemployment. The jagged initial response reflects the fact that, in the empirical wage equation, both the level and the change in the vacancy-to-unemployment ratio affect inflation. Once these short-run effects fade, inflation increases slowly but steadily, reflecting the effects of a tight labor market on wages (which in turn affect prices) and the feedback of realized inflation to short-run and long-run

inflation expectations. In this simulation, inflation has risen about 0.7 percent after 6 quarters, 1.0 percent after 12 quarters, and keeps increasing at about 0.15 percent per year from then on. This pattern is very similar, even quantitatively, to the weak feedback case we showed in Figure 2. In general, the takeaway is that persistent labor market pressure leads to ever-increasing inflation. The rate of increase in our estimated model is relatively low, reflecting a weak estimated catch-up effect and well-anchored inflation .expectations. Yet, as inflation expectations are not fully anchored, extended labor market tightness can lead to significant additional inflation.

FIGURE 11. RESPONSE OF INFLATION TO SHOCKS TO THE VACANCY TO UNEMPLOYMENT RATIO.

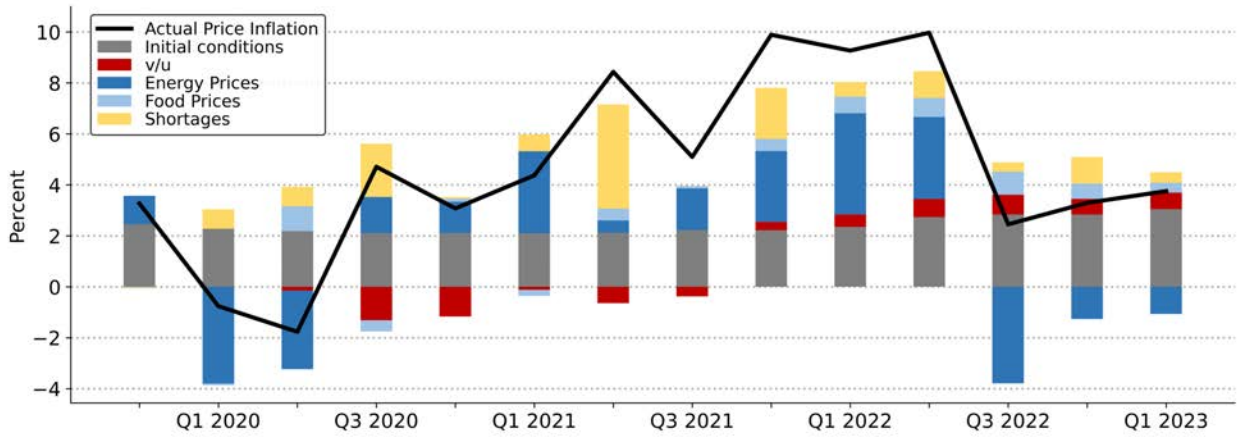


Note: The figure shows the full-model, dynamic responses of inflation to a one-standard-deviation positive shock to vacancy-to-unemployment ratio. It may be compared to the hypothetical responses of inflation under alternative parameter assumptions in Figure 2.

Quantifying the sources of the pandemic-era inflation

With the estimated model in hand, we can calculate the impulse response functions of any endogenous variable to shocks to both (i) the residuals associated with the four structural equations and (ii) the exogenous variables. As we have seen, these estimated impulse responses capture both the direct and indirect (general equilibrium) effects on each endogenous variable, in both the current period and in the future, of any exogenous shocks. Moreover, because the model is linear, the IRFs of each endogenous variable, including the responses to the equation residuals, sum exactly to realized value of the endogenous variable at each date. Consequently, the estimated model and the calculated IRFs for historical shocks can be used to decompose the variation in any endogenous variable into its sources, taking fully into account dynamic and general equilibrium effects.

FIGURE 12. THE SOURCES OF PRICE INFLATION, 2020Q1 to 2023Q1



Note: The figure shows a decomposition of the sources of inflation, 2020Q1 to 2023Q1, based on the solution of the full model and the implied impulse response functions. The continuous line shows actual inflation, and the total net heights of the bars are the model's forecast of inflation in each period, given initial conditions through 2019Q4 and excluding the effects of equation residuals. The grey portion of each bar shows the contribution of pre-2020 data (and also include the contributions of productivity shocks). Colored segments of each bar show the general equilibrium, fully dynamic contribution of each exogenous variable to inflation in that period, as implied by the estimated model.

Using this approach, Figure 12 shows the estimated sources of inflation from 2020Q1 to 2023Q1. For comparison, the continuous line shows actual inflation. (The inflation data are quarterly, at annualized rates, and thus more jagged than the annualized series often reported.) Figure 12 does not show the contributions to inflation of shocks to the residuals of the model equations, both to avoid clutter and because the residuals by definition do not have a clear structural interpretation. Consequently, the sum of the estimated sources of inflation in each quarter (the net heights of the bars) does not exactly equal realized inflation, the difference being the omitted effects of the equation residuals.

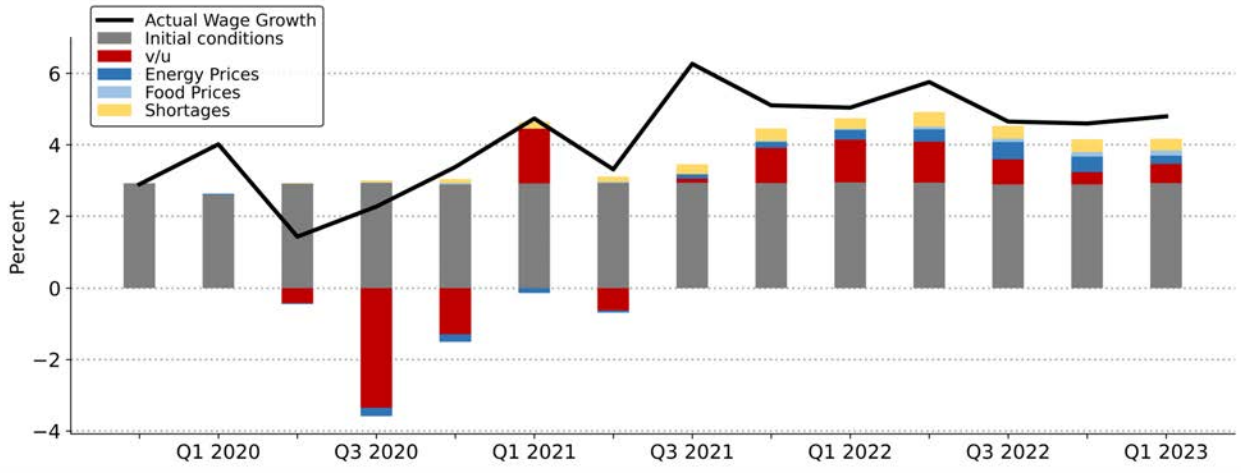
The decomposition underlying Figure 12 takes as given data prior to 2020Q1. The implications of the constants in the equation, plus the dynamic effects of pre-covid shocks, plus the actual behavior of the productivity variable (whose dynamic effects we do not analyze) are put together in an "initial conditions" category, shown in the figure by grey bars. The figure shows the contribution of initial conditions to be roughly constant over time, suggesting that, absent the pandemic-era shocks considered here, inflation would have likely remained stable at around 2 percent or a bit higher into 2023.

Initial conditions aside, the colored components of the vertical bars in Figure 12 show the full dynamic contributions to inflation, in each quarter, of each of the four exogenous variables. Shocks to the rate of change of the relative price of energy and food are constructed as deviations of the values of those variables from zero. Shocks to the shortage and the vacancy-to-unemployment ratio variables are constructed as the actual values of those variables minus their values in 2019Q4.

Figure 12 yields several conclusions. First, the contributions of food and (especially) energy price shocks to the pandemic-era inflation were large. Energy price shocks in particular account for much of the rise of overall inflation in late 2021 and the first half of 2022, and the for the decline in inflation in the second half of 2022. The large and extended contributions of commodity price shocks to inflation shown in Figure 12 are not inconsistent with our earlier finding that price shocks tend to have transitory effects on inflation, as both energy and food prices rose continuously over much of the period (see Figures 4 and 5), which our procedure interprets as a series of positive shocks. Second, sectoral shortages were the dominant source of inflation in 2021Q2, and the effects of shortages on inflation, both direct and indirect, remained significant through the end of our sample period. Third, and importantly, the contribution to inflation of tight labor-market conditions—the leading concern of many early critics of U.S. monetary and fiscal policies—was quite small early on, and indeed was negative in 2020 and early 2021 as labor markets suffered from the effects of the pandemic recession. However, over time, as the labor market has remained tight, the traditional wage Phillips curve effect has begun to assert itself, with the high vacancy-to-unemployment ratio becoming an increasingly important, though by no means dominant, source of inflation.

To get a better understanding of the role of labor market tightness, we can use the same approach to decompose wage inflation into its sources (Figure 13). Again, the contribution of “initial conditions” is very stable, at around 3 percent per year, roughly the rate of nominal wage growth consistent with 2 percent price inflation. Price shocks (energy, food, shortages) contribute less to wage growth than to price inflation, but even at the beginning of 2023 the direct and indirect effects of shocks to prices still explain a greater share than does the vacancy-to-unemployment ratio of the excess of wage growth over that predicted by the initial conditions.

FIGURE 13. THE SOURCES OF WAGE INFLATION, 2020Q1 to 2023Q1



Note: The figure shows a decomposition of the sources of nominal wage growth, 2020Q1 to 2023Q1, based on the solution of the full model and the implied impulse response functions. The continuous line shows actual wage growth, and the total (net) heights of the bars are the model's forecast of wage growth in each period, given initial conditions through 2019Q4 and excluding the effects of equation residuals. The grey bar shows the contribution of pre-2020 data (and also includes the contributions of productivity shocks). Colored segments of each bar show the general equilibrium, fully dynamic contribution of each exogenous variable to wage growth in that period, as implied by the estimated model.

The retrospective evaluation of how policymakers assessed the economy in early 2021, around the time of the passage of the American Rescue Plan, is thus more complicated than sometimes portrayed. At least for the short run, policymakers were initially more right than wrong in arguing that tight labor markets would not generate high inflation, although in focusing on unemployment rates and total employment rather than the vacancy-to-unemployment ratio they did underestimate labor market tightness. As we have seen, the initial sources of the inflation were primarily shocks to prices given wages, including large energy price shocks and the surprisingly persistent shortages of cars and other durable goods. By raising inflation expectations, and to a smaller extent by creating a catch-up effect, these shocks ultimately affected wages as well as prices. However, over time a very tight labor market has begun to exert increasing pressure on inflation, pressure which our model predicts will continue to grow. Over the longer run, then, the critics who worried about an overheated economy leading to inflation may be able to claim some vindication.

Model-based projections

Some insight into the future, as well as into the workings of the model, can be obtained by projecting the model forward, under specified assumptions about the future paths of the four

exogenous variables. We briefly look at some projections in this section. However, we do so very cautiously, for several reasons.

First, beyond the usual statistical uncertainty surrounding the parameter estimates and the specification of the model itself, there is also a great deal of uncertainty about the future course of the exogenous variables in the model. A full analysis would consider alternative scenarios with differing assumptions about energy prices and other variables, which we do not attempt here. Instead, we interpret the simulations described below as conditional projections, not unconditional forecasts.

Second, all else equal, our model is designed to project inflation and wage growth conditional on the state of the labor market, as measured by the vacancy-to-unemployment ratio (v/u). Since we know the past realizations of v/u, we were able to use the model in the previous section to estimate the portion of historical price and wage inflation stemming from tight labor markets. However, the model takes v/u as given and provides no insight about the particular fiscal and monetary policies that would be needed to achieve a given state of the labor market in the future. In addition, we know that the Beveridge curve, the downward-sloping relation between job vacancies and unemployment, shifted dramatically during the pandemic (Blanchard, Domash, and Summers 2022), obscuring the relationship of the vacancy-to-unemployment ratio and other measures of labor market tightness, including unemployment. We return to this point below.

Finally, for the purposes of projection the "natural" or steady-state value of the vacancy-to-unemployment ratio, v/u^*, is important. Specifically, absent new shocks, in our model bringing down wage inflation requires a period in which v/u is below its steady-state value. The steady state of v/u implied by our model depends on the parameters of the wage and price equations and is not tightly estimated. For consistency with the decomposition exercises of the previous section, and following the assumption made by Blanchard, Domash, and Summers, for these projections we assume that v/u^* equals 1.2, its value in 2019Q4, the last full quarter before the pandemic, a period when inflation appeared stable.[16]

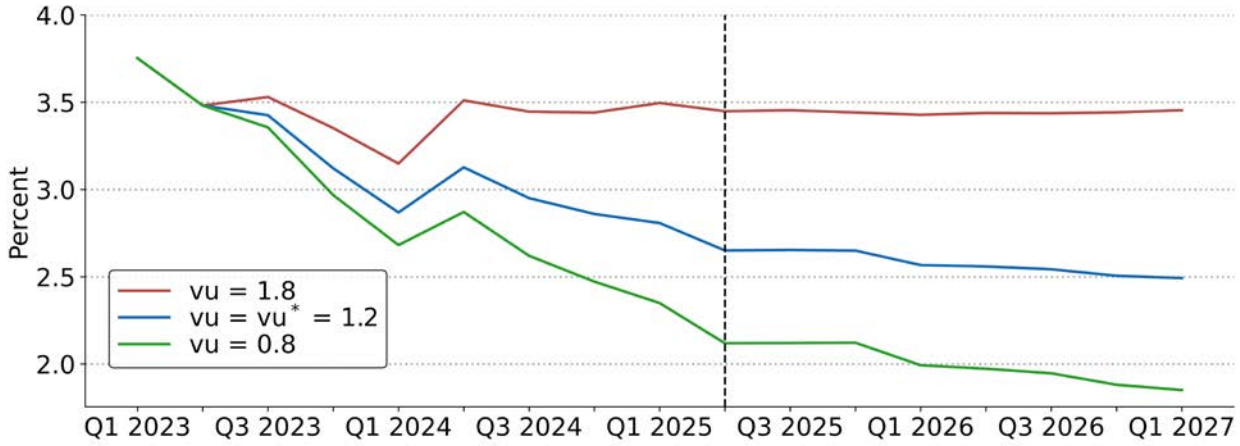
With these substantial caveats, we can use the model to project the course of price inflation and wage growth, starting in 2023Q2, with initial conditions determined by the actual data for the prior four quarters. We assume that food and energy prices grow for the next few years at a rate equal to nominal wage growth, so the growth rates of the relative prices *grpe* and *grpf* are both equal to zero; that

[16] We adjust the constant in the estimated wage equation so that the model delivers a steady-state value of v/u^* equal to 1.2.

demand and supply at the sectoral level have largely normalized, so that the shortage variable takes an index value of 10 (on a scale in which, recall, its highest weekly value is 100);[17] and that productivity growth is 1 percent, close to its in-sample average. Equation residuals are set to zero. As in the rest of the paper, we use the price equation estimated over the full sample, including the covid period, and base the estimates for the other three equations on pre-covid data.

We consider three projections, corresponding to three assumed trajectories of the labor market. Specifically, in our simulations we assume that macro policy brings down v/u over eight quarters, beginning in 2023Q2, to one of three terminal values: 1.8 (approximately the current value of v/u); 1.2 (the value that prevailed immediately before the pandemic, which we have assumed is the steady-state or "natural" value of v/u); and 0.8. Obviously, lower values of v/u require tighter monetary and fiscal policies, all else equal. The value of v/u after eight quarters is not important for the points we want to make, so we simply assume that it remains constant thereafter. The results of the projections are shown in Figure 14. The three curves in the figure show the projected path of inflation under each of the assumed terminal points for the vacancy-to-unemployment ratio (the first observation in each trajectory is actual inflation in 2023Q1). As before, we use the CPI to measure inflation (not the PCE deflator, which the Fed targets) and the ECI to measure the nominal wage.

FIGURE 14. INFLATION PROJECTIONS UNDER ALTERNATIVE PATHS FOR v/u



Note: The figure shows projected inflation in the model if v/u is brought down linearly to the specified terminal level over eight quarters. The value for 2023Q1 reflects actual data.

[17] The realized values of *shortage* in the four quarters preceding the start of the simulation are 33, 18, 24, 22.

The qualitative lesson of Figure 14 is straightforward and would generalize for alternative assumptions about the steady-state value of v/u. All else equal, and not surprisingly, allowing v/u to remain near current levels does not bring inflation down in our projections. Indeed, because an extended period of inflation raises long-term inflation expectations, it leads to slowly increasing inflation. In contrast, in our projections, reducing the vacancy-to-unemployment ratio to its assumed steady-state level of 1.2 over the next eight quarters brings CPI inflation down to about 2.7 percent, which, adjusting for historical differences between CPI and PCE inflation, is within striking distance of the Fed's target. A more aggressive policy, which brings v/u down to 0.8 in eight quarters, is projected to reduce CPI inflation to close to 2 percent by the end of that period. In the latter case, the projections show continuing declines in inflation for a constant level of v/u; the implication is that this more restrictive policy should ultimately be reversed.

One would like to know the unemployment rates associated with the three alternative paths. Unfortunately, that depends on a factor outside the model, namely, the future evolution of the position of the Beveridge curve (see the debate between Blanchard, Domash, and Summers 2022 and Figura and Waller 2022). As we have discussed, during the pandemic the Beveridge curve shifted upward, implying that, for reasons not fully understood, the job-worker matching process had become less efficient. If that shift were to be fully reversed, then a decline in v/u to a level below its natural rate of 1.2, together with an unemployment rate above the implied natural rate of 3.6 percent (the actual unemployment rate in late 2019) could bring down inflation over time with a limited increase in unemployment (not quite an "immaculate disinflation", as unemployment must exceed the natural rate for some time). If, in the other extreme case, the Beveridge curve were to remain at its position as of early 2023, then the unemployment rate associated with v/u equal to 1.2—that is, the implied natural rate of unemployment—would be about 4.3 percent.[18] As implied by the figure, reducing inflation to the 2 percent target in that case would require an unemployment rate above that level for a time.

IV. Conclusion

[18] Blanchard, Domash, and Summers show that the unemployment rate associated with a given v/u is proportional to $\left(\frac{v}{u}\right)^{-0.6}$. We calculate the natural rate by inferring the coefficient of proportionality from current data on unemployment and vacancy-to-unemployment rates and then multiplying by our estimated v/u^* of 1.2, raised to a power of -0.6.

The pandemic-era inflation has been a complicated phenomenon that involved both multiple sources and complex dynamic interactions. Ultimately, as many have recognized, the inflation largely reflected strong aggregate demand, the product of easy fiscal and monetary policies, excess savings accumulated during the pandemic, and the reopening of locked-down economies. The strength of aggregate demand led to a (belatedly recognized) tightening of the labor market, as reflected in record high ratios of job openings to unemployed workers. However, at least initially, the arguments that the tighter labor market would not create much inflation, so long as the labor market tightness was temporary, were in fact correct. Although the wage Phillips curve was operative, inflation expectations did not de-anchor and there was little evidence of a wage-price spiral, in that workers did not achieve nominal wage gains sufficient to compensate them for unexpected price increases (a weak catch-up effect). Our decomposition of inflation shows that tight labor markets alone made only a modest contribution to inflation early on.

Instead, initially, the inflation reflected primarily developments in product markets. Excessive aggregate demand set off the inflation but did so primarily by raising prices given wages, through its contribution to higher commodity prices (which reflected other forces as well, including, from 2022 on, the war in Ukraine) and by increasing the demand for goods for which supply was inelastic.

The good news for policymakers from our analysis is that the indirect effects of the price shocks on inflation were more short-lived than might have been expected. In particular, inflation expectations, especially short-run expectations, rose enough to put upward pressure on wages and prices but generally remained well anchored. Although shocks to prices given wages played a central role in the pandemic-era inflation, in the absence of strong indirect effects or renewed shocks, the portion of inflation accounted for by these shocks should largely (although not fully) dissipate over time, even without policy action.

However, even as the effects of price shocks have waned, the effects of tight labor markets have begun to cumulate. Our decomposition shows that, as of early 2023, tight labor market conditions still accounted for a minority share of excess inflation. But according to our analysis, if current labor market conditions persist, that share is likely to grow and will not subside on its own. The portion of inflation which traces its origin to overheating of labor markets can only be reversed by policy actions that bring labor demand and supply into better balance. The effects of such policies on unemployment depend on the extent to which matching efficiency in the labor market normalizes (that is, on whether the Beveridge curve shifts down to its pre-pandemic position).

What are the lessons of all this for macroeconomics? The pandemic was extraordinarily disruptive, making forecasting and policymaking difficult. Nevertheless, we don't think that the recent experience justifies throwing out existing models of wage-price dynamics. We were able to simulate the pandemic-era inflation using a simple and rather traditional model of the interactions of wages, prices, and inflation expectations. That said, our results show that inflation can be a complicated phenomenon. Policymakers must be alert to the possibility that inflationary pressures can come from product markets as well as labor markets, for example, through unexpected changes in input costs or shifts in demand that collide with inelastic sectoral supply curves. Moreover, through various mechanisms, inflationary pressures from product and labor markets can interact and be mutually reinforcing. The good news is that, empirically, we find the second-round and indirect effects of product-market shocks to be relatively modest, implying that, as the early critics warned, labor-market balance should ultimately be the primary concern for central banks attempting to maintain price stability.

Appendix. Model estimates and robustness checks

Estimates and selected statistics for the four equations of the model using alternative measures of the endogenous variables are given immediately below (Tables A2 and A3). Table A4 provides full estimation results, including estimated coefficients for all lagged values, for the base specification and some specifications using alternative data. Price equations are estimated over the sample 1990Q1 to 2023Q1, which includes the covid period, while other equations are estimated on the pre-covid sample, ending in 2019Q4. P-values are given for both (i) the hypothesis that the sum of a set of coefficients is zero (*p-stat (sum)*) and (ii) the joint hypothesis that each of the coefficients applying to a given right-side variable is zero (*p-stat (joint)*). Numbers in parentheses indicate included lags, e.g., (0 to -4) says that the contemporaneous value and four lags of the variable are included in the regression. Variable mnemonics used in the reporting of regression results are given in Table A1 below (see also Table 1 in the text).

TABLE A1. VARIABLE MNEMONICS USED IN THIS APPENDIX

WAGE EQUATION
gw = change in the log nominal wage (ECI)
cf1 = one-year inflation expectations (Federal Reserve Bank of Cleveland)
spf1 = one-year inflation expectation (Survey of Professional Forecasters)
v/u = vacancy-to-unemployment ratio
catch-up = the excess of the actual log price level over the expected log price level
gpty = 8-quarter moving average productivity trend

PRICE EQUATION (additional variables)
gp = gcpi = change in the log nominal price level (CPI)
gpce = change in the log nominal price level (PCE)
grpe = change in the log of CPI energy prices relative to the aggregate wage
grpf = change in the log of CPI food prices relative to the aggregate wage
shortage = index of Google searches for “shortage”

INFLATION EXPECTATIONS EQUATIONS (additional variables)
cf10 = ten-year inflation expectations (Federal Reserve Bank of Cleveland)
spf10 = long-run inflation expectations (ten-year inflation forecast from Survey of Professional Forecasters)

TABLE A2. ESTIMATION WITH PCE INFLATION RATHER THAN CPI INFLATION

Dependent variable: gw

Independent variable	gw	v/u	catch-up	cf1	gpty
Lags	-1 to -4	-1 to -4	-1 to -4	-1 to -4	-1
Sum of coefficients	0.345	0.774	0.008	0.655	0.068
p-stat (sum)	0.050	0.016	0.092	0.000	0.258
p-stat (joint)	0.226	0.037	0.404	0.003	0.258
R-squared			0.579		
No. observations			120		

Dependent variable: gpce

Independent variable	gpce	gw	grpe	grpf	shortage	gpty
Lags	-1 to -4	0 to -4	0 to -4	0 to -4	0 to -4	-1
Sum of coefficients	0.501	0.499	0.038	0.086	0.014	-0.104
p-stat (sum)	0.001	0.001	0.002	0.133	0.442	0.116
p-stat (joint)	0.001	0.004	0.000	0.237	0.018	0.116
R-squared			0.844			
No. observations			133			

Dependent variable: cf1

Independent variable	cf1	cf10	gp
Lags	-1 to -4	0 to -4	0 to -4
Sum of coefficients	0.375	0.527	0.098
p-stat (sum)	0.031	0.001	0.016
p-stat (joint)	0.001	0.000	0.000
R-squared		0.884	
No. observations		120	

Dependent variable: cf10

Independent variable	cf10	gp
Lags	-1 to -4	0 to -4
Sum of coefficients	0.971	0.029
p-stat (sum)	0.000	0.129
p-stat (joint)	0.000	0.006
R-squared	0.929	
No. observations	120	

TABLE A3. ESTIMATION WITH SURVEY OF PROFESSIONAL FORECASTERS INSTEAD OF CLEVELAND FED FORECASTS

Dependent variable: gw

Independent variable	gw	v/u	catch-up	spf1	gpty
Lags	-1 to -4	-1 to -4	-1 to -4	-1 to -4	-1
Sum of coefficients	0.233	1.154	0.087	0.767	0.154
p-stat (sum)	0.205	0.002	0.285	0.000	0.024
p-stat (joint)	0.494	0.001	0.441	0.001	0.024
R-squared			0.599		
No. observations			120		

Dependent variable: gp

Independent variable	gp	gw	grpe	grpf	shortage	gpty
Lags	-1 to -4	0 to -4	0 to -4	0 to -4	0 to -4	-1
Sum of coefficients	0.335	0.665	0.066	0.126	0.018	-0.143
p-stat (sum)	0.037	0.000	0.000	0.050	0.281	0.027
p-stat (joint)	0.066	0.000	0.000	0.050	0.000	0.027
R-squared			0.947			
No. observations			133			

Dependent variable: spf1

Independent variable	spf1	spf10	gp
Lags	-1 to -4	0 to -4	0 to -4
Sum of coefficients	0.854	0.113	0.033
p-stat (sum)	0.000	0.031	0.067
p-stat (joint)	0.000	0.000	0.000
R-squared		0.952	
No. observations		120	

Dependent variable: spf10

Independent variable	spf10	gp
Lags	-1 to -4	0 to -4
Sum of coefficients	0.969	0.031
p-stat (sum)	0.000	0.002
p-stat (joint)	0.000	0.029
R-squared	0.978	
No. observations	120	

TABLE A4. FULL ESTIMATION RESULTS FOR ALTERNATIVE MEASURES OF INFLATION AND INFLATION EXPECTATIONS

Dependent variable: gw

Using gcpi		Using gpce		Using spf	
gw (-1)	0.16	gw (-1)	0.14	gw (-1)	0.08
	(0.10)		(0.10)		(0.10)
gw (-2)	0.12	gw (-2)	0.08	gw (-2)	0.10
	(0.10)		(0.10)		(0.10)
gw (-3)	0.17*	gw (-3)	0.15	gw (-3)	0.12
	(0.10)		(0.10)		(0.10)
gw (-4)	0.00	gw (-4)	-0.03	gw (-4)	-0.06
	(0.10)		(0.10)		(0.10)
cf1 (-1)	0.34**	cf1 (-1)	0.38***	spf1 (-1)	-0.11
	(0.14)		(0.13)		(0.40)
cf1 (-2)	-0.03	cf1 (-2)	0.00	spf1 (-2)	0.22
	(0.15)		(0.14)		(0.46)
cf1 (-3)	0.20	cf1 (-3)	0.27*	spf1 (-3)	0.35
	(0.16)		(0.15)		(0.45)
cf1(-4)	0.03	cf1(-4)	-0.01	spf1(-4)	0.30
	(0.16)		(0.15)		(0.39)
magpty (-1)	0.03	magpty (-1)	0.07	magpty (-1)	0.15**
	(0.06)		(0.06)		(0.07)
vu(-1)	3.77*	vu(-1)	3.31*	vu(-1)	3.38*
	(1.98)		(1.94)		(1.96)
vu (-2)	-1.81	vu (-2)	-1.69	vu (-2)	-0.67
	(3.63)		(3.58)		(3.57)
vu. (-3)	-3.64	vu. (-3)	-3.58	vu. (-3)	-3.77
	(3.66)		(3.61)		(3.60)
vu (-4)	2.38	vu (-4)	2.74	vu (-4)	2.21
	(2.10)		(2.07)		(2.10)
catch-up (-1)	-0.01	catch-up (-1)	-0.11	catch-up (-1)	0.12
	(0.08)		(0.13)		(0.08)
catch-up (-2)	0.01	catch-up (-2)	0.13	catch-up (-2)	0.04
	(0.08)		(0.15)		(0.09)
catch-up (-3)	-0.00	catch-up (-3)	-0.05	catch-up (-3)	-0.07
	(0.08)		(0.14)		(0.09)
catch-up (-4)	-0.02	catch-up (-4)	0.04	catch-up (-4)	-0.00
	(0.07)		(0.11)		(0.08)
constant	-0.27	constant	-0.93*	constant	-0.75***
	(0.20)		(0.50)		(0.24)
R-squared	0.578	R-squared	0. 579	R-squared	0. 599
No. observations	120	No. observations	120	No. observations	120

Dependent variable: gp = (gcpi or gpce)

Using gcpi		Using gpce		Using spf	
gcpi (-1)	0.04	gpce (-1)	0.03	gcpi (-1)	0.04
	(0.11)		(0.09)		(0.11)
gcpi (-2)	0.18*	gpce (-2)	0.23**	gcpi (-2)	0.18*
	(0.10)		(0.09)		(0.10)
gcpi (-3)	0.23**	gpce (-3)	0.29***	gcpi (-3)	0.23**
	(0.11)		(0.09)		(0.11)
gcpi (-4)	-0.11	gpce (-4)	-0.04	gcpi (-4)	-0.11
	(0.10)		(0.08)		(0.10)
gw	0.37***	gw	0.21**	gw	0.37***
	(0.09)		(0.10)		(0.09)
gw (-1)	0.18	gw (-1)	0.08	gw (-1)	0.18
	(0.11)		(0.11)		(0.11)
gw (-2)	0.04	gw (-2)	0.24**	gw (-2)	0.04
	(0.11)		(0.11)		(0.11)
gw (-3)	0.04	gw (-3)	-0.04	gw (-3)	0.04
	(0.11)		(0.11)		(0.11)
gw (-4)	0.05	gw (-4)	0.00	gw (-4)	0.05
	(0.10)		(0.11)		(0.10)
magpty	-0.14**	magpty	-0.10	magpty	-0.14**
	(0.06)		(0.07)		(0.06)
grpe	0.09***	grpe	0.04***	grpe	0.09***
	(0.00)		(0.00)		(0.00)
grpe (-1)	0.00	grpe (-1)	0.03***	grpe (-1)	0.00
	(0.01)		(0.00)		(0.01)
grpe (-2)	-0.01	grpe (-2)	-0.01	grpe (-2)	-0.01
	(0.01)		(0.01)		(0.01)
grpe (-3)	-0.02**	grpe (-3)	-0.02***	grpe (-3)	-0.02**
	(0.01)		(0.01)		(0.01)
grpe (-4)	0.01	grpe (-4)	-0.01	grpe (-4)	0.01
	(0.01)		(0.00)		(0.01)
grpf	0.11***	grpf	0.03	grpf	0.11***
	(0.04)		(0.03)		(0.04)
grpf (-1)	-0.02	grpf (-1)	0.03	grpf (-1)	-0.02
	(0.04)		(0.03)		(0.04)
grpf (-2)	-0.00	grpf (-2)	-0.04	grpf (-2)	-0.00
	(0.04)		(0.03)		(0.04)
grpf (-3)	-0.00	grpf (-3)	-0.01	grpf (-3)	-0.00

	(0.04)		(0.03)		(0.04)
grpf (-4)	0.04	grpf (-4)	0.07**	grpf (-4)	0.04
	(0.04)		(0.03)		(0.04)
shortage	0.11***	shortage	0.04*	shortage	0.11***
	(0.02)		(0.02)		(0.02)
shortage (-1)	-0.03	shortage (-1)	0.04**	shortage (-1)	-0.03
	(0.02)		(0.02)		(0.02)
shortage (-2)	0.00	shortage (-2)	0.01	shortage (-2)	0.00
	(0.02)		(0.02)		(0.02)
shortage (-3)	-0.03	shortage (-3)	-0.04**	shortage (-3)	-0.03
	(0.02)		(0.02)		(0.02)
shortage (-4)	-0.02	shortage (-4)	-0.04*	shortage (-4)	-0.02
	(0.02)		(0.02)		(0.02)
constant	-0.11	constant	-0.27	constant	-0.11
	(0.16)		(0.20)		(0.16)
R-squared	0.947	R-squared	0.844	R-squared	0.947
No. observations	133	No. observations	133	No. observations	133

Dependent variable: cf1/spf1

Using gcpi		Using gpce		Using spf	
cf1 (-1)	0.30***	cf1 (-1)	0.25***	spf1 (-1)	0.63***
	(0.09)		(0.10)		(0.09)
cf1 (-2)	-0.23**	cf1 (-2)	-0.21**	spf1 (-2)	0.17
	(0.10)		(0.10)		(0.11)
cf1 (-3)	0.18*	cf1 (-3)	0.15	spf1 (-3)	0.21*
	(0.10)		(0.10)		(0.11)
cf1 (-4)	0.12	cf1 (-4)	0.19**	spf1 (-4)	-0.15*
	(0.07)		(0.09)		(0.09)
cf10	1.20***	cf10	1.23***	spf10	0.67***
	(0.14)		(0.16)		(0.15)
cf10 (-1)	-0.47**	cf10 (-1)	-0.40	spf10 (-1)	-0.04
	(0.23)		(0.25)		(0.18)
cf10 (-2)	0.02	cf10 (-2)	-0.14	spf10 (-2)	-0.58***
	(0.23)		(0.26)		(0.19)
cf10 (-3)	0.14	cf10 (-3)	0.20	spf10 (-3)	-0.11
	(0.23)		(0.25)		(0.18)
cf10 (-4)	-0.38**	cf10 (-4)	-0.36*	spf10 (-4)	0.18
	(0.17)		(0.20)		(0.15)
gcpi	0.04***	gpce	0.09***	gcpi	0.03***
	(0.01)		(0.02)		(0.01)
gcpi (-1)	0.10***	gpce (-1)	0.08***	gcpi (-1)	0.02***
	(0.01)		(0.03)		(0.01)
gcpi (-2)	0.00	gpce (-2)	-0.01	gcpi (-2)	-0.01
	(0.02)		(0.03)		(0.01)
gcpi (-3)	0.01	gpce (-3)	-0.01	gcpi (-3)	-0.00

	(0.02)		(0.03)		(0.01)
gcpi (-4)	-0.02	gpce (-4)	-0.05**	gcpi (-4)	-0.01
	(0.02)		(0.03)		(0.01)
R-squared	0.908	R-squared	0.929	R-squared	0.952
No. observations	120	No. observations	120	No. observations	120

Dependent variable: cf10/spf10

Using gcpi		Using gpce		Using spf	
cf10 (-1)	0.85***	cf10 (-1)	0.82***	spf10 (-1)	0.70***
	(0.09)		(0.09)		(0.09)
cf10 (-2)	-0.05	cf10 (-2)	-0.06	spf10 (-2)	0.26**
	(0.12)		(0.12)		(0.11)
cf10 (-3)	0.20	cf10 (-3)	0.25**	spf10 (-3)	0.00
	(0.12)		(0.12)		(0.11)
cf10 (-4)	-0.03	cf10 (-4)	-0.05	spf10 (-4)	-0.00
	(0.09)		(0.09)		(0.09)
gcpi	0.03***	gpce	0.05***	gcpi	0.01
	(0.01)		(0.01)		(0.00)
gcpi (-1)	0.01	gpce (-1)	-0.02	gcpi (-1)	0.01**
	(0.01)		(0.02)		(0.00)
gcpi (-2)	-0.01	gpce (-2)	-0.00	gcpi (-2)	0.01*
	(0.01)		(0.01)		(0.00)
gcpi (-3)	-0.00	gpce (-3)	-0.01	gcpi (-3)	0.00
	(0.01)		(0.01)		(0.00)
gcpi (-4)	-0.00	gpce (-4)	0.01	gcpi (-4)	0.01*
	(0.01)		(0.01)		(0.00)
R-squared	0. 931	R-squared	0.932	R-squared	0.978
No. observations	120	No. observations	120	No. observations	120

REFERENCES

Ball, Laurence, Daniel Leigh, and Prachi Mishra (2022). "Understanding U.S. Inflation During the COVID Era," presented at the Brookings Papers on Economic Activity fall conference, Washington DC, September.

Barnichon, Regis (2010). "Building a Composite Help-Wanted Index", Economics Letters, vol. 109, no. 5, December, 175-178.

Benigno, Pierpaolo and Gauti Eggertson (2023). "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve," NBER working paper 31197, April.

Blanchard, Olivier (2016). "The Phillips Curve: Back to the '60s?," American Economic Review, vol. 106, no. 5, May, 31-34.

Blanchard, Olivier (2021). "In Defense of Concerns Over the $1.9 Trillion Relief Plan," Peterson Institute for International Economics, February 18.

Blanchard, Olivier and Peter Diamond (1989). "The Beveridge Curve," Brookings Papers on Economic Activity, vol. 20, Spring, 1-76.

Blanchard, Olivier, Alex Domash, and Lawrence Summers (2022). "Bad News for the Fed from the Beveridge Space," Peterson Institute Policy Brief 22-7, July.

Cecchetti, Stephen, Michael Feroli, Peter Hooper, Frederic Mishkin, and Kermit Schoenholtz (2023). "Managing Disinflations," U.S. Monetary Policy Forum, University of Chicago Booth School of Business, February 24.

Comin, Diego, Robert Johnson, and Callum Jones, (2023). "Supply Chain Constraints and Inflation," NBER working paper 31179, April.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, Muhammed Yıldırım (2023). "Quantifying the Inflationary Impact of Fiscal Stimulus under Supply Constraints," American Economic Review Papers and Proceedings, forthcoming.

Edelberg, Wendy and Louise Sheiner (2021). "The Macroeconomic Implications of Biden's $1.9 Trillion Fiscal Package," Brookings Institution, January 28.

Figura, Andrew, and Christopher Waller (2022). "What Does the Beveridge Curve Tell Us About the Likelihood of a Soft Landing?" FEDS Notes, Federal Reserve System, July 29.

Gagnon, Joseph and Madi Sarsenbayev (2022), "25 Years of Excess Unemployment in Advanced Economies: Lessons for Monetary Policy", working paper 22-17, Peterson Institute for International Economics, October.

Gagliardone, Luca, and Mark Gertler (2023). “Oil Prices, Monetary Policy and Inflation Surges,” Manuscript, New York University.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning (2022). “Macroeconomic Implications of Covid-19: Can Negative Supply Shocks Cause Demand Shortages?” American Economic Review, May, vol. 112, no. 5, 1437-74.

Kamin, Steven and John Roberts (2023). “Will a Recovery of Real Wages Obstruct Progress Toward Disinflation?” AEI Economic Policy Working Paper Series, May 25.

Koch, Christoffer, and Diaa Noureldin (2023). “How We Missed the Recent Inflation Surge.” Finance and Development, International Monetary Fund, March, 21-23.

Reifschneider, David and David Wilcox (2022). “The Case for a Cautiously Optimistic Outlook for US Inflation,” Peterson Institute for International Economics, Policy Brief 22-3, March.

Summers, Lawrence (2021). “The Biden Stimulus Is Admirably Ambitious. But It Brings Some Big Risks, Too,” Washington Post, February 4.

Zeballos, Eliana, Xiao Dong, and Ergys Islamaj (2023). “A Disaggregated View of Market Concentration in the Food Retail Industry,” U.S. Department of Agriculture, Economic Research Report ERR-314, January.

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

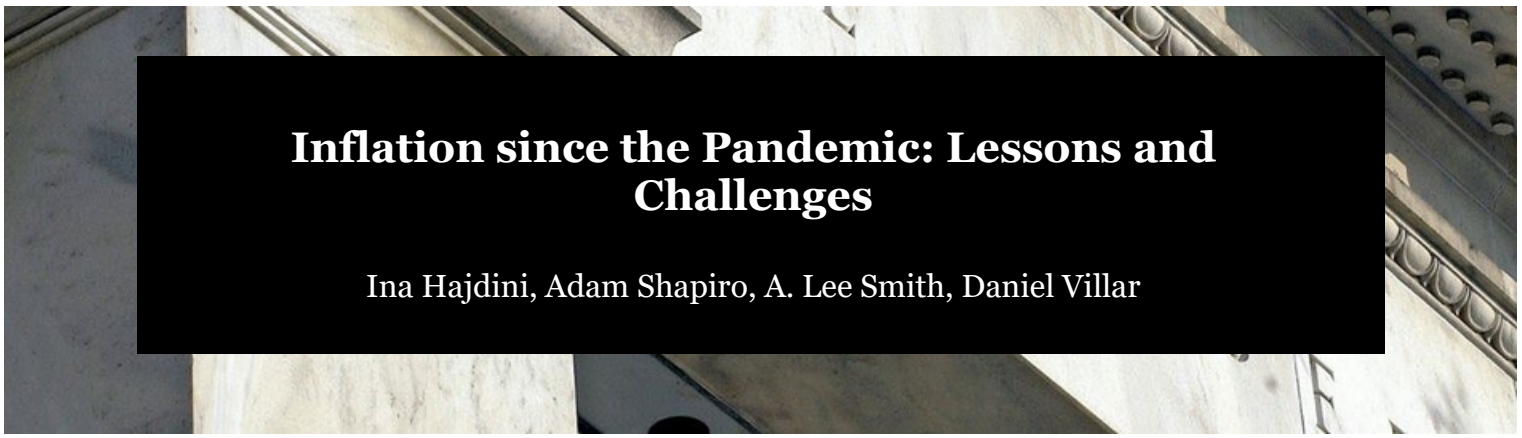
Inflation since the Pandemic: Lessons and Challenges

Ina Hajdini, Adam Shapiro, A. Lee Smith, Daniel Villar

2025-070

Please cite this paper as:
Hajdini, Ina, Adam Shapiro, A. Lee Smith, and Daniel Villar (2025). "Inflation since the Pandemic: Lessons and Challenges," Finance and Economics Discussion Series 2025-070. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.070.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.



Inflation since the Pandemic: Lessons and Challenges

Ina Hajdini, Adam Shapiro, A. Lee Smith, Daniel Villar

The analysis in this paper was presented to the Federal Open Market Committee as background for its discussion of the Federal Reserve's 2025 review of its monetary policy strategy, tools, and communications.

Abstract: This paper reviews the drivers of the post-pandemic U.S. inflation surge and subsequent decline, including the behavior and role of inflation expectations. The sharp rise in inflation reflected severe imbalances between supply and demand stemming from the shocks of the pandemic and the policy response. Measures of short-term inflation expectations increased alongside realized inflation, especially those of households and firms, which may have contributed to inflation's persistence through price- and wage-setting behavior. However, measures of longer-term inflation expectations remained generally well anchored, which likely prevented a larger or more lasting increase in inflation. The stability of longer-term inflation expectations, together with easing supply and demand imbalances, allowed inflation to fall from its peak in mid-2022 without a large increase in unemployment. We conclude by reviewing some lessons learned from this episode as well as potential risks to inflation going forward.

JEL Classification: E31, E52, E58, E70.

Keywords: Inflation, inflation expectations, COVID-19, monetary policy.

Note: Authors' affiliations are Federal Reserve Bank of Cleveland (Hajdini), Federal Reserve Bank of San Francisco (Shapiro), Federal Reserve Bank of Kansas City (Smith), and Board of Governors of the Federal Reserve System (Villar), respectively. The views expressed are those of the authors and do not necessarily reflect the views of anyone else affiliated with the Federal Reserve System. The authors thank Andrew Figura, Ed Knotek, Sylvain Leduc, Anna Paulson, Trevor Reeve, Sam Schulhofer-Wohl, and Stacey Tevlin for valuable feedback and suggestions. On earlier drafts, the authors also benefited from comments by Stephanie Aaronson, Glenn Follette, Kurt Lewis, Ekaterina Peneva, Jeremy Rudd, and William Wascher. The authors thank Lucas Moyon for excellent research assistance.

1. Introduction and overview

U.S. inflation surged in the aftermath of the pandemic after being subdued for more than two decades. This paper reviews the key drivers of U.S. inflation over this period as well as the behavior and role of inflation expectations and highlights some lessons learned that could be informative for understanding and monitoring inflation dynamics going forward. Our key takeaways are the following:

- High inflation reflected the interaction of large and persistent shifts in supply and demand across the U.S. economy. These shifts were driven by the pandemic itself, as well as stabilization efforts, including strong fiscal stimulus and accommodative monetary policy. The resulting imbalances were amplified by capacity constraints, which limited the ability of supply to increase—often referred to as a nonlinearity.
- Although the labor market eventually reached historically tight levels, supply and demand imbalances extended well beyond the labor market. Even accounting for nonlinearities, most Phillips curve estimates suggest that labor market tightness can explain only part of the rise in inflation in 2021 and 2022.
- Despite increases in shorter-term inflation expectations, monetary policy helped to keep longer-term inflation expectations generally well anchored. The stability of longer-term inflation expectations, in turn, likely helped to contain the inflation surge, allowing inflation and short-term inflation expectations to fall without a large increase in unemployment.
- The events driving the pandemic inflation surge were extreme, which warrants caution against drawing too many conclusions from this episode. Nevertheless, we think there are two lessons that can be carried forward: (1) that the inflation environment can change abruptly amid large shocks and (2) the importance of keeping longer-term inflation expectations well anchored.
- Looking ahead, inflation risks appear to be more balanced now than they appeared five years ago. We highlight three particular considerations: (1) a greater appreciation for supply factors and the importance of capacity constraints, (2) the potential for inflation expectations to be more responsive to price increases after a period of high inflation, and (3) a reassessment of effective lower bound (ELB) risk and the threat it poses to longer-term inflation expectations.

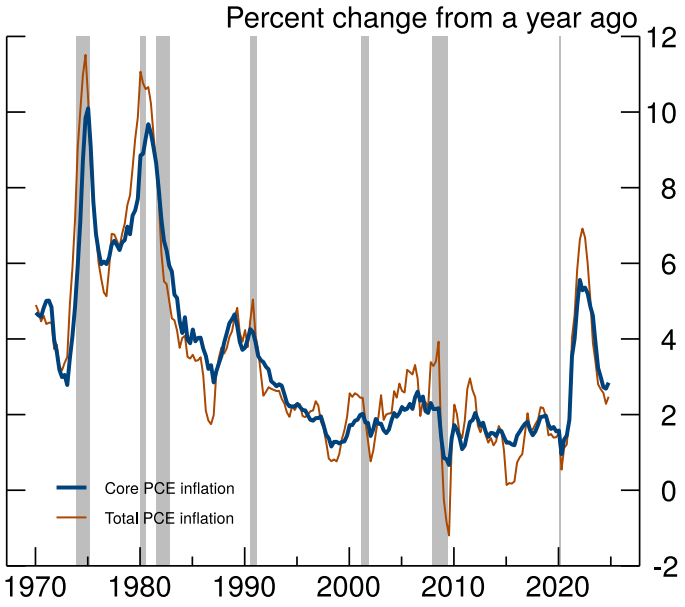
2. The inflation surge: The roles of supply, demand, and nonlinearities

In this section, we review the shocks resulting from the pandemic and subsequent policy responses that led to large shifts in supply and demand. We also examine the role played by supply constraints in amplifying the inflationary effects of these shifts—so-called nonlinearities in the aggregate supply curve. Finally, we discuss the easing of supply–demand imbalances beginning in 2022 that brought about the decline in inflation.

2.1 The post-2020 inflation surge—a stark break from pre-pandemic inflation

The U.S., along with economies across the globe, experienced a surge in inflation in the years following the onset of the COVID-19 pandemic.[1] Inflation, as measured by the 12-month change in the personal consumption expenditures (PCE) price index, reached 7.3 percent in the summer of 2022. As shown in figure 1, this spike represented a stark break from the decade before the pandemic, in which inflation was mostly low and increases in inflation were not persistent.

Figure 1: Headline and core PCE inflation



Source: Bureau of Economic Analysis; National Bureau of Economic Research.

This sharp rise in inflation reflected severe imbalances between supply and demand stemming from the large and persistent shocks unleashed by the pandemic and the policy response. The initial phase of the pandemic was centered around public health risks that prompted production shutdowns, disruptions in global supply chains, and a contraction in labor

[1] For a discussion of the international inflation experience, see Lipinska, Martinez Garcia, and Schwartzman (2025).

supply. Strong monetary and fiscal support came quickly, which kept household balance sheets strong and supported demand throughout 2020. By the spring of 2021, the unemployment rate had recovered significantly, and as shown in figure 2, total consumer spending had returned to its pre-pandemic trend.[2]

Consumer preferences, however, had shifted dramatically away from high-contact services toward goods. Although total consumer spending had recovered, services consumption was approximately 5 percent below its pre-pandemic trend, while goods consumption was approximately 10 percent above its trend. With pandemic-related supply constraints still binding, this change in consumer preferences strained the goods market.[3] This result was evidenced by exhausted inventories, long delivery times, and, as shown in figure 3, increasing goods price inflation.

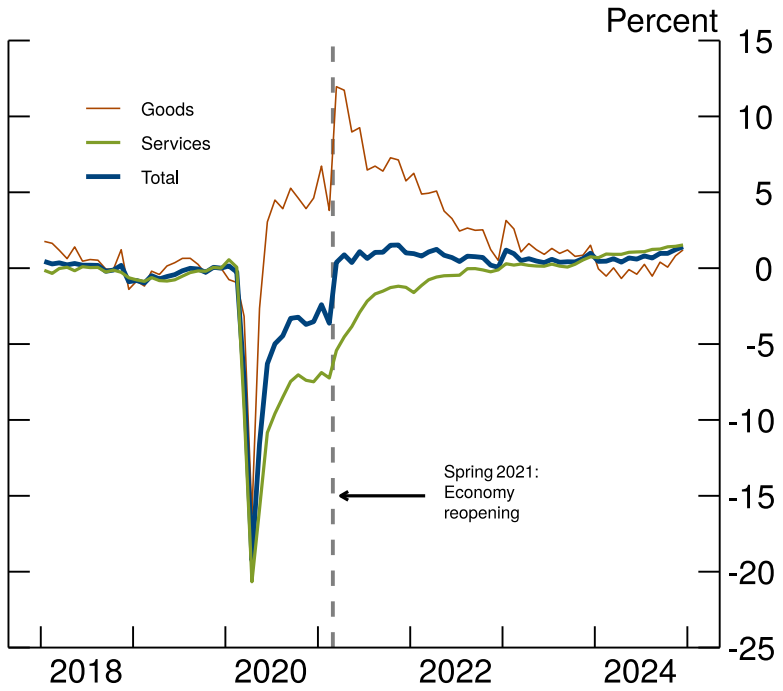
The relative demand for goods and services started normalizing after the economy began reopening in earnest in spring 2021. However, inflation pressures emerged in the services sector alongside a sluggish recovery in labor supply. Supply and demand imbalances in the labor market worsened throughout 2021, indicated by a surge in quits and job vacancies. Nominal wage growth steadily increased, keeping upward pressure on production costs as well as propping up nominal household incomes. As figure 4 shows, wage growth increased substantially among both job switchers and all workers. By early 2022, an array of economic indicators suggested that aggregate demand greatly exceeded the reduced productive capacity of the economy, and total PCE inflation was nearing peak levels.[4]

[2] The trend is constructed over the 2015–19 period.

[3] Ferrante, Graves, and Iacoviello (2023) find that the shift in demand toward goods contributed significantly to the surge in overall inflation because of the constraints faced by producers in expanding production.

[4] For example, the vacancy-to-unemployment (V/U) ratio and quits rate both peaked in March 2022, the share of respondents in the Quarterly Survey of Plant Capacity Utilization (QPC) reporting an insufficient supply of materials peaked at 43 percent in the fourth quarter of 2021 (relative to a pre-pandemic level near 10 percent), and the housing rental vacancy rate bottomed out at 5.6 percent in the second quarter of 2022 (relative to a pre-pandemic level near 7 percent).

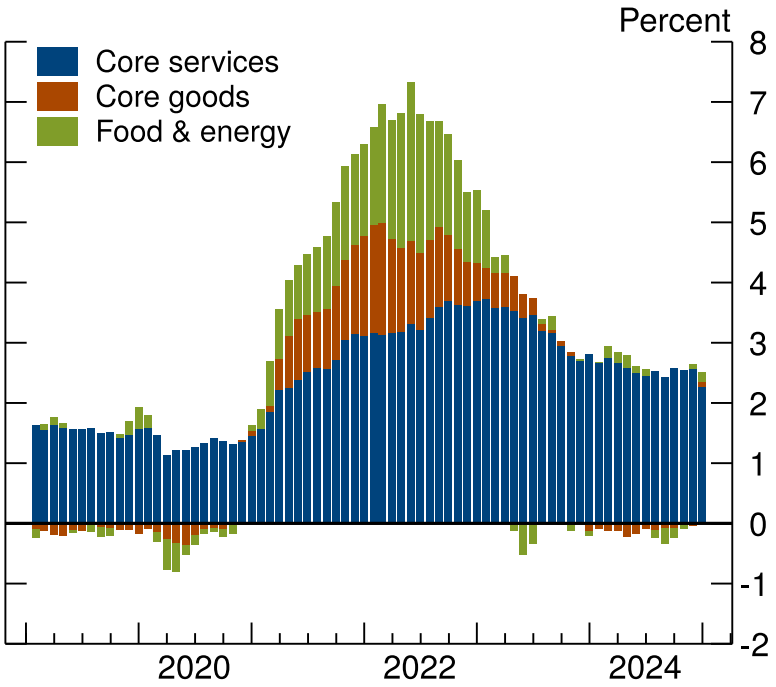
Figure 2: Real consumer spending, percent deviation relative to pre-pandemic trend



Notes: Displayed are the percent deviations of real goods PCE, real services PCE, and total real PCE from their respective 2015–19 trends.

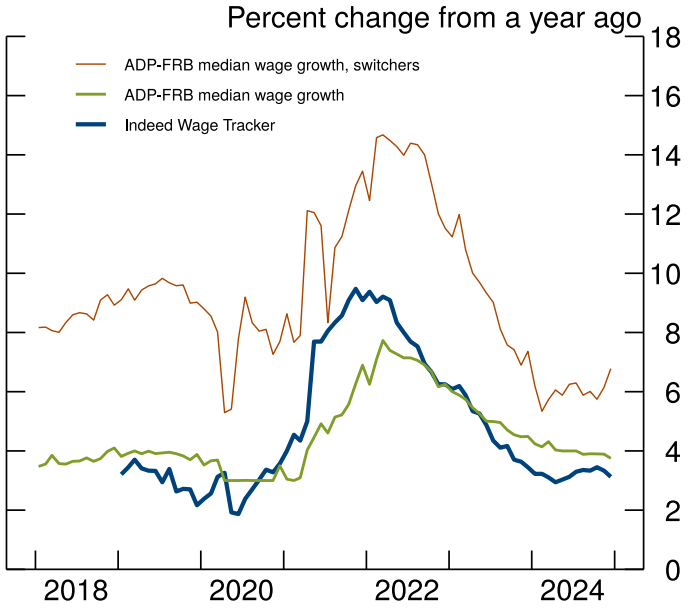
Source: Bureau of Economic Analysis via FRED; authors' calculations.

Figure 3: Contributions to 12-month headline PCE inflation (year-over-year)



Source: Bureau of Economic Analysis via FRED; Federal Reserve Bank of San Franscisco staff calculations.

Figure 4: Measures of wage growth



Source: Indeed, Inc.; Automatic Data Processing, Inc. Payroll Processing Data; Federal Reserve Board staff calculations.

2.2 Large and persistent shifts in supply and demand spurred inflation

Shortages and supply disruptions indicate that reduced supply helped spur inflation in 2021; however, the surge in economic activity, reflected by consumer spending in figure 2, suggests that strong demand was also an important factor. An active line of research attempts to measure the contributions of demand and supply to the inflation surge, and there remains no consensus on the relative importance of each. Some studies report that demand factors were the major driver, while others point to mainly supply factors.[5] Most show that while some mix of supply and demand factors led to heightened inflation, their effects occurred at different times.[6]

[5] Bernanke and Blanchard (2025), for example, find that sectoral and commodity price shocks were especially important contributors in 2021 and 2022. Cuba-Borda and L'Huillier (2025) review the literature and evidence in favor of inflation being primarily a supply phenomenon. In contrast, Giannone and Primiceri (2024) find that high inflation in both the U.S. and the euro area was driven largely by demand forces. Lansing (2025) finds that demand forces were important for inflation during the pandemic era and dominated the influence of supply forces in the U.S.

[6] Di Giovanni, Kalemli-Özcan, Silva, and Yildirim (2023 [rev. 2025]), for instance, find that while supply shocks were important in 2020, demand factors became more relevant in 2021 and 2022. Dupor and Hogan (2025) find that demand conditions began rebounding in 2021 before supply factors began having stronger inflationary effects in 2022

One example of this pattern is the San Francisco Fed's decomposition of 12-month PCE inflation into supply- and demand-driven components shown in figure 5.[7]

This decomposition estimates how supply and demand factors evolved in the years preceding COVID-19 and over the course of the pandemic period. The supply-driven contribution steadily increased following the onset of the pandemic alongside the emergence of severe supply chain disruptions and reductions in labor supply.[8] Specifically, the supply-driven contribution increased 0.7 percentage point through 2020 and an additional 1.3 percentage points through 2021. This steady increase in supply-driven inflation coincided with accelerating input costs across the globe in 2021—namely, commodities such as oil, lumber, plastic, and metals. The overall contribution of supply factors peaked in June 2022, when it reached 3.2 percentage points. This peak was four months after Russia's invasion of Ukraine, which further disrupted global energy and commodity markets.

Demand factors also contributed to the inflation surge, though their dynamics differed from those of supply factors, particularly in the early phase of the pandemic. The demand-driven contribution fell 1.4 percentage points in April 2020, as stay-at-home orders and virus-related risks caused consumers to significantly pull back on spending. Demand-driven inflation then began to reverse course, coinciding with the broader reopening of the economy as well as the implementation of the American Rescue Plan—which boosted government spending and household incomes while monetary policy remained very accommodative.[9] Demand-driven inflation steadily increased thereafter, peaking in September 2022 at 3.1 percentage points—a few months after the peak in supply-driven inflation, but with a similar magnitude. This pattern suggests that strong demand exerted upward pressure on inflation for well over a year following the economy's reopening.

---

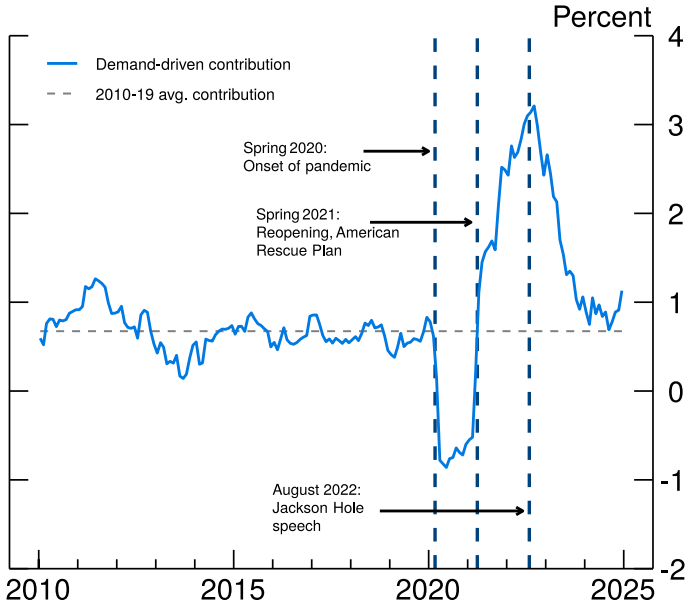
[7] The decomposition, outlined in Shapiro (2025), is based on the price and quantity movements in the categories underlying the PCE price index. Demand-driven categories are identified as those where an unexpected change in price moves in the same direction as the change in quantity in a given month. Supply-driven categories are identified as those where unexpected changes in price and quantity move in opposite directions. One caveat to this type of identification scheme is that it includes the endogenous reactions to aggregate supply and demand shocks. For example, as shown in Shapiro (2025), a negative oil supply shock can cause a positive demand shock for non-oil fuels. The decomposition also includes an "ambiguous" category (not shown), which captures cases where the methodology cannot distinguish between supply and demand factors. For this reason, the supply- and demand-driven contributions in figure 5 do not sum to total PCE inflation.

[8] Labor supply reductions were apparently caused by virus-related health risks, caregiving responsibilities for children unable to attend school, and an increase in early retirements.

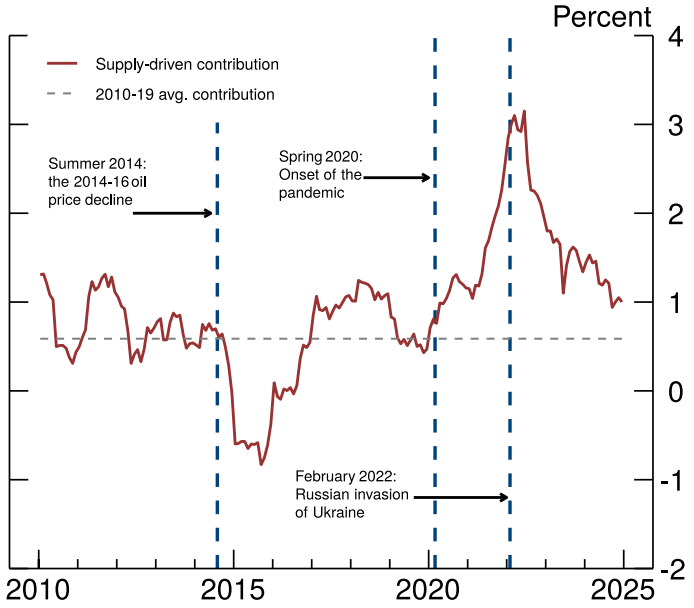
[9] The demand-driven contribution of monthly PCE inflation began surging in March 2021. Recent research continues to debate the contributions of monetary and fiscal policy to the inflation surge, with little apparent consensus. Jorda and Nechio (2023) find that fiscal support contributed 3 percentage points to inflation by the end of 2021, while Barnichon et al. (2021) estimated an effect of only 0.3 percentage point. Bocola et al. (2024) find that the shift to accommodative monetary policy, and specifically the flexible average inflation targeting (FAIT) framework, explains half of the rise in inflation in the post-2020 period. In contrast, Duncan, Martinez Garcia, and Miller (2025) find that FAIT played a much smaller role.

Figure 5: Supply- and demand-driven PCE inflation

Panel A



Panel B



Notes: Panel A displays the demand-driven contribution to 12-month PCE inflation. Panel B displays the supply-driven contribution to 12-month PCE inflation. Not depicted is the ambiguous contribution.

Source: Federal Reserve Bank of San Francisco.

2.3 Amplifying forces and the role of nonlinearities

The previous section highlighted that shifts in both supply and demand contributed to the rise in inflation. Partly underlying the inflationary effect of those shifts was the emergence of capacity constraints. Capacity constraints imply that firms' or industries' supply curves can steepen at higher levels of output—referred to as a nonlinearity.[10] As depicted in panel A of figure 6, it becomes costly for industries to increase their output when production levels are near capacity and, as a result, prices are raised to a greater extent in response to a given increase in demand. As shown in panel B, this effect can be exacerbated by declines in the productive capacity of the industry, which increase the likelihood that the industry is operating on the steep portion of the supply curve. Specifically, a negative supply shock—an inward shift of the supply curve from S_1 to S_2—lowers the point at which the supply curve steepens. A positive demand shock is therefore more likely to cross the steep part of the supply curve for an industry experiencing a negative supply shock. Boehm and Pandalai-Nayar (2022) show that nonlinearities in industry-level supply curves can imply a nonlinearity at the aggregate level. Other studies have shown that more complex interactions of supply and demand shocks across industries, inputs, and countries can have important amplification effects on inflation.[11]

Although identification is challenging, recent studies show evidence of nonlinearities in supply for firms and industries.[12] Furthermore, numerous indicators suggest that capacity constraints occurred throughout 2021 and 2022 across many sectors of the economy. Delivery times for materials spiked, and shortages for critical inputs such as semiconductors arose. The New York Fed's Global Supply Chain Pressure Index (GSCPI), which aggregates a number of supply chain indicators, reached four standard deviations above its average historical value. The QPC showed historically unprecedented constraints on production due to material and labor shortages, as well as logistical problems.[13] Labor shortages were also evidenced by historically

[10] For instance, the estimated demand-driven contribution to inflation seen in figure 5 may look large, in part, because firms were constrained in their ability to meet elevated levels of demand.

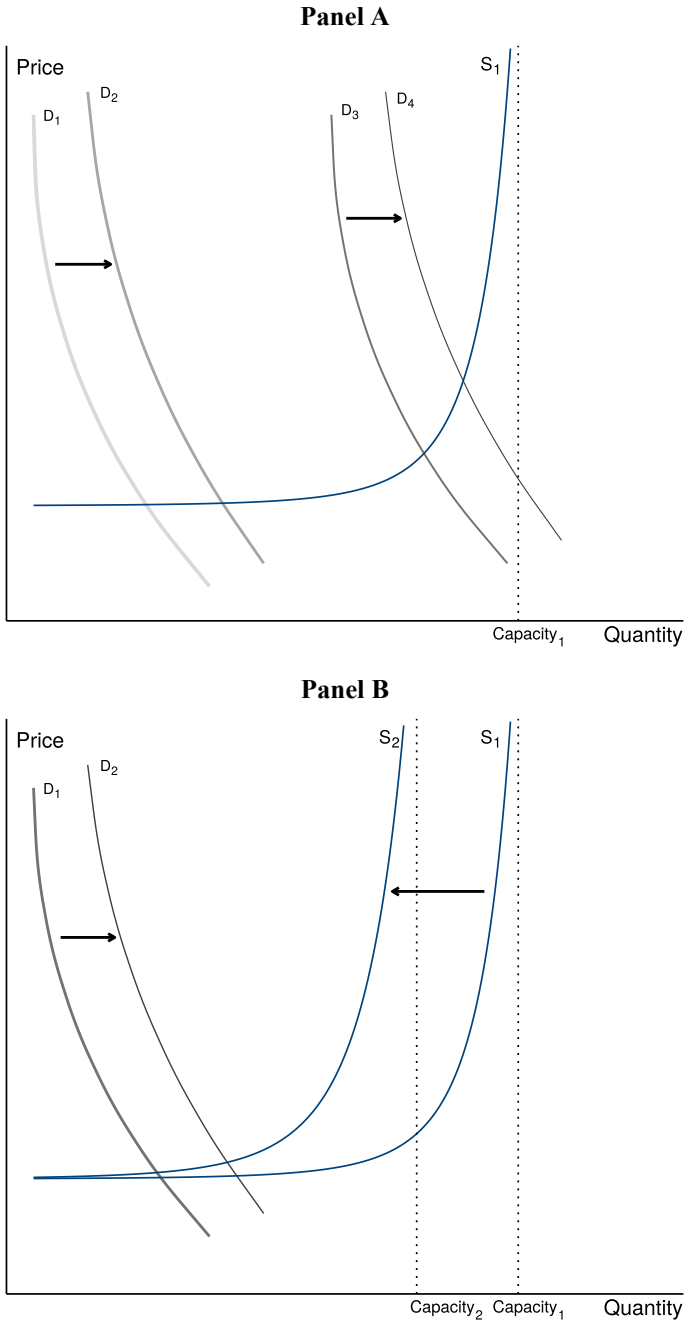
[11] Guerrieri, Lorenzoni, Straub, and Werning (2022) present a model in which a negative supply shock in one sector can reduce demand for products in other sectors. Di Giovanni, Kalemli-Özcan, Silva, and Yildirim (2023 [rev. 2025]) estimate a model in which this effect can occur across countries. Luo and Villar (2023) find evidence that the inflationary effect of sectoral supply shocks was amplified by input-output linkages based on a calibrated multi-sector model. Amiti, Heise, Karahan, and Sahin (2023) show that the simultaneity of supply shocks—a labor disutility shock and input price shock—acted to limit substitution across input types, amplifying the inflationary effect of cross-sectoral shocks.

[12] Boehm and Pandalai-Nayar (2022) find that industries that produce close to capacity have up to five times the price elasticity as those that produce far below capacity. Balleer and Noeller (2024) find that firms that experience shortages of materials raise prices by more than unconstrained firms.

[13] Braun, Flaaen, and Hacioglu Hoke (2024) use QPC data to show that such constraints contributed significantly to inflation in goods.

high levels of the V/U ratio.[14] The consistent pattern across all of these indicators implies that capacity constraints played at least some role in exacerbating the inflation surge.

Figure 6: The role of nonlinearities in amplifying inflation



[14] See Foote et al. (2025) for a discussion of the Beveridge curve in the context of assessing maximum employment.

Nonlinearities have also been studied in the context of the Phillips curve, which captures the relationship between aggregate excess demand—that is, the level of economic activity relative to its available capacity—and price inflation. It is common in the literature to use indicators of labor market tightness—either the unemployment gap or the V/U ratio—as measures of aggregate excess demand, although numerous alternatives have been proposed.[15] In this context, a nonlinear Phillips curve implies that a change in labor market tightness has a larger effect on inflation when the labor market is already tight.

The evidence of a nonlinear Phillips curve is mixed. Some studies—for instance, Benigno and Eggertson (2023)—find that tight labor markets can explain much of the inflation surge when allowing for a nonlinearity. Estimates from other studies imply smaller effects or even that nonlinearities are not present at all.[16] A notable identification issue is that it is challenging to distinguish between changes in the slope and shifts in the Phillips curve.[17] The disagreement across studies underscores the challenges of estimating nonlinearities in a Phillips curve setting, making it unclear to what extent labor market tightness accounts for heightened inflation levels.[18] Overall, the timing, breadth, and global nature of inflation since 2021 suggest that imbalances in the economy extended well beyond the labor market and that labor market tightness was not the sole source of the inflation surge.[19]

One piece of evidence pointing to nonlinear effects can be seen in the frequency at which prices change, which increased substantially in the recent high inflation period, as seen in figure 7.[20] Because firms raise prices more quickly when inflation is high, the inflationary effect of

[15] Barnichon and Shapiro (2022) run a horse race between alternative measures of economic slack in forecasting inflation in a Phillips curve framework. They find that labor market indicators, particularly the V/U ratio, provide superior inflation forecasts for prices and wages.

[16] Barnichon and Shapiro (2024) and Smith, Timmerman, and White (2024) find that the slope of the Phillips curve increases in tight labor markets but still remains relatively small. Doser, Nunes, Rao, and Sheremirov (2023) and Beaudry, Hou, and Portier (2025) argue that the nonlinearity disappears when properly controlling for inflation expectations.

[17] For instance, Comin, Johnson, and Jones (2023) argue that constraints act to shift the Phillips curve—similar to a markup shock.

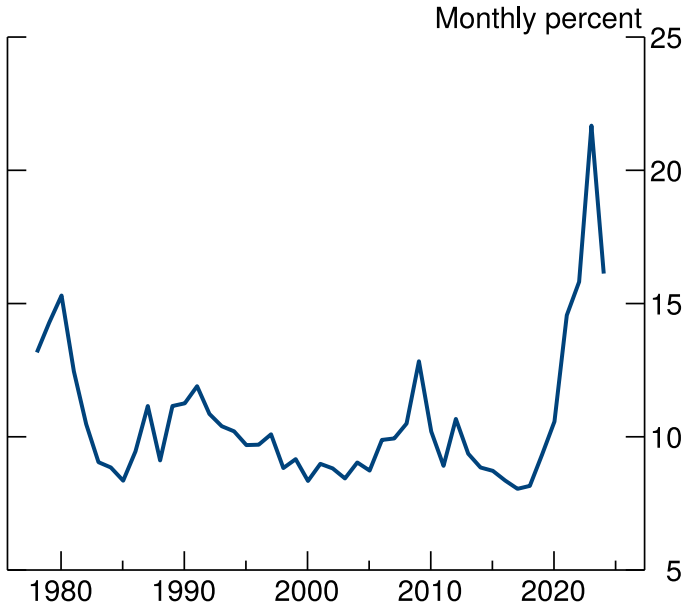
[18] Limited data is an important issue, as there are not many episodes with severely tight labor market conditions. Furthermore, accounting for endogeneity is more difficult in a nonlinear model because alternative instruments are needed. Barnichon and Shapiro (2024) use time fixed effects at the metropolitan statistical area level to help alleviate endogeneity issues and find that the nonlinearity disappears under this specification.

[19] Inflation began rising in the spring of 2021, well before labor markets became unusually tight. Furthermore, the surge was most notable in the goods sector, where inflation did not appear to be related to labor market slack before the pandemic. We also note that wage Phillips curve models based on measures of labor market tightness—and similar to the ones discussed here—do not match the magnitude of the increase in wage growth. This suggests that shocks and imbalances beyond those in the labor market, including possibly households' and firms' near-term inflation expectations, also mattered for wage growth during this period. Finally, see Lipinska, Martinez Garcia, and Schwartzman (2025) for a discussion of the global nature of inflation.

[20] Cavallo, Lippi, and Miyahara (2024) find large increases in the frequency of food and beverage price changes in several countries, while Montag and Villar (2025) find broad-based increases in the frequency in the U.S.

tightness in the economy may be greater and more rapid. In this way, changes in firms' price-setting behavior can contribute to nonlinearities in the Phillips curve.

Figure 7: Median frequency of price change in the U.S. CPI



Source: Bureau of Labor Statistics; Montag and Villar (2025).

The role of price setting is illustrated by "state dependent" models in which firms face costs of adjusting their prices—so-called menu costs. In such models, firms adjust prices more frequently and thus will tend to raise prices more quickly when there is a pick-up in cost or general price inflation. State-dependent models therefore imply that the pass-through of cost shocks to inflation is larger and more rapid for large shocks, which can contribute to a nonlinearity in the Phillips curve.[21] In addition, the recent rise in the frequency of price changes in the U.S. was much larger than many menu cost models would predict, and even larger than the rise seen in previous high inflation periods.[22] This finding suggests that firms faced relatively smaller costs or constraints to changing prices than they did before the pandemic. As a result,

[21] Blanco et al. (2024) and Harding, Linde and Trabandt (2024) present price setting models that imply non-linearities in the Phillips curve.

[22] See Montag and Villar (2025). In addition, most of the macro models used to analyze monetary policy in central banks assume that the frequency of price change is fixed, as in a Calvo (1983)-style model. In such models, the frequency of price change is directly related to the slope of the Phillips curve, so a higher frequency implies a steeper Phillips curve. However, the fact that the frequency changes over time is inconsistent with the assumptions of such models.

shocks to inflation may have been transmitted more rapidly and strongly than in the past, leading to an amplified response to supply–demand imbalances.[23]

2.4 The post-2022 unwinding of supply and demand imbalances

Since mid-2022, an easing of supply and demand imbalances has helped inflation decline to levels much closer to the Committee's target. Monetary policy tightening and a fading of fiscal support reduced demand, while a healing of supply chains, strong productivity growth, and an increase in the labor force helped expand supply. As shown in figure 2, consumer expenditures across goods and services fully normalized by the end of 2022, which alleviated strains on supply chains.

Figure 5 shows that demand-driven inflation began declining in the fall of 2022, as the Committee rapidly tightened monetary policy and fiscal impetus began fading.[24] By 2022, the level of "excess savings", the additional amount of household savings attributable to the pandemic and associated fiscal support, was significantly lower than in 2021.[25] The healing of supply chains and increases in productive capacity also helped reduce inflationary pressure during this time. Materials shortages diminished, production of certain goods, such as motor vehicles, expanded, and there was a surge in multifamily housing construction. As a result, inflation for goods and market rents fell sharply in 2022 and 2023.

Greater supply was also apparent in the labor market, where a surge in immigration and increases in labor-force participation helped expand the labor force.[26] Overall labor-market rebalancing was evidenced by a decline in the V-U ratio and in the quits rate, and by cooling wage growth. By late 2023, measures of wage growth for job switchers and in job postings, which likely better reflect supply and demand conditions in the labor market than the wage growth of workers remaining at the same job, were at or below their pre-pandemic levels (figure 4).[27]

---

[23] The reduction in the costs or constraints around price changes may have been due to firms being more aware of the need to change prices given the magnitude and expected persistence of the changes occurring, or consumers, who may normally dislike seeing prices increase, being more understanding of frequent price changes given the environment.

[24] Najjar and Shapiro (2025) find that core PCE inflation would have been approximately 3 percentage points higher on average between 2022 and 2024 if the FOMC had not raised interest rates in 2022 and 2023. Comin, Johnson, Jones (2023) attribute part of the relaxation in constraints to monetary policy tightening.

[25] See Aladangady et al. (2022), Abdelrahman and Oliveira (2023) and Klitgaard and Higgins (2023). Abdelrahman, Oliveira, and Shapiro (2024) show that the same pattern emerges when measuring household wealth.

[26] Duzhak (2024a) estimates that approximately one-fifth of the easing of labor market tightness in 2023 can be attributed to the increase in immigration. Beyond immigration, there was a large increase in labor force participation among prime-age adults beginning in 2021. Considine et al. (2025) document that much of this increase was due to a rise in labor force participation among women.

[27] See Foote et al. (2025) for a discussion of how wage growth can reflect the balance between demand and supply in the labor market.

Notably, inflation has fallen almost 5 percentage points since mid-2022 to early 2025 without a large increase in the unemployment rate or much softening in GDP growth. Studies have noted that this pattern is consistent with a nonlinear Phillips curve, however, we reiterate that it is difficult to pinpoint the underlying mechanism.[28] For instance, this pattern is also consistent with a shift in the Phillips curve due to a decline in short-term inflation expectations alongside anchored longer-term expectations, which we describe in section 3.

Although inflation has declined substantially, it remains above the Committee's longer-run goal of 2 percent as of early 2025. Persistently elevated inflation is notable in the services sector, especially housing services. While market rent growth returned to pre-pandemic levels in 2022, Consumer Price Index and PCE rent inflation remained stubbornly high for years, reflecting the slow pace at which existing lease contracts catch up to market rates and illustrates how inflation can persist.[29] Disinflation for core non-housing services inflation has also been slow, partly reflecting the lagged adjustment of some prices to large shocks. For instance, insurance premiums and fee schedules for payments to health care providers, which represent about 35 percent of core non-housing services, are typically adjusted once per year and heavily influenced by regulatory factors.[30]

3. The behavior and role of inflation expectations

In this section, we describe the behavior of short- and longer-term inflation expectations across various agents in the economy and the role they played in shaping recent inflation dynamics. As described in the previous section, we think that the recent inflation episode was primarily driven by pandemic-induced supply–demand imbalances. However, recent research suggests that the rise in short-term inflation expectations may have contributed to the *persistence* of the inflation surge.[31] Nonetheless, the evidence suggests that longer-term inflation expectations remained generally well anchored, which likely helped contain the inflation surge and contributed to inflation falling without a large increase in unemployment.

[28]Convexities in the Beveridge curve suggest that small movements in unemployment can be associated with a large disinflation, as explained by Crust, Lansing, and Petrosky-Nadeau (2023). Figura and Waller (2022) note that being on the very steep portion of the Beveridge curve, when vacancies far exceed unemployment, can allow for disinflation without a sharp rise in unemployment.

[29] Gallin et al. (2024) and Loewenstein, Meyer, and Verbrugge (2024) document that continuing-tenant rents adjust more slowly to market conditions than new-tenant rents, which could be due to landlords avoiding moving costs or providing discounts to reliable tenants.

[30] The PCE price index includes payments to healthcare providers from Medicare and Medicaid, which are set based on lagged cost indexes and also subject to adjustments set by the Centers for Medicare and Medicaid Services (CMS). Clemens, Gottlieb, and Shapiro (2014) highlight that private payment rates follow Medicare payment rates but adjust sluggishly.

[31] Research by Armantier et al. (2022) found that consumer short-term inflation expectations were as sensitive to inflation news during the inflation surge as during the pre-pandemic period, implying that households' short-term inflation expectations could have exhibited a similar behavior before and after the pandemic.

3.1 Short-term inflation expectations of consumers and firms likely contributed to inflation persistence

Beginning in 2021, year-ahead inflation expectations measured across consumers, firms, professional forecasters, and financial markets increased alongside realized inflation, peaking in 2022, as visualized by figure 8.[32] After the inflation peak in 2022, inflation expectations began to decline to finally stabilize at close to pre-pandemic values as of the end of 2024. Our key takeaway is that the large and prolonged rise in households' and firms' short-term inflation expectations – following the inflation surge – likely contributed to inflation persistence through their documented influence on wage and price setting dynamics.

Theory suggests that short-term inflation expectations—regardless of how they are formed—play an important role for households' wage negotiations and firms' pricing decisions. Empirically isolating the effect of short-term inflation expectations on inflation is challenging.[33] However, the literature has recently started to use randomized controlled trial (RCT) experiments to quantify the effects of consumers' and firms' expectations on their economic decisions, providing insights into the role of short-term inflation expectations for inflation dynamics. During the recent inflation surge, an increase in U.S. consumers' inflation expectations was found to increase the probability that they would search for a higher paying job, and to have a small but positive effect on their income growth expectations.[34] Hence, the elevated short-term household inflation expectations could have contributed to higher wage growth, perhaps because workers increased job search to find higher wages or used the threat to do so to bargain for higher wages.[35] Other evidence highlights that a positive shock to firms' short-term inflation expectations can cause firms to raise prices, as firms associate higher future inflation with higher prices for raw materials.[36] Moreover, increases in firms' cost expectations positively correlate

[32] Firms' short-term expectations about their individual costs, as measured by the Atlanta Fed Business Inflation Expectations Survey, behaved similarly.

[33] As shown in Hajdini (2023), short-term inflation expectations tend to highly correlate with past inflation, blurring the direct effects of inflation expectations on inflation.

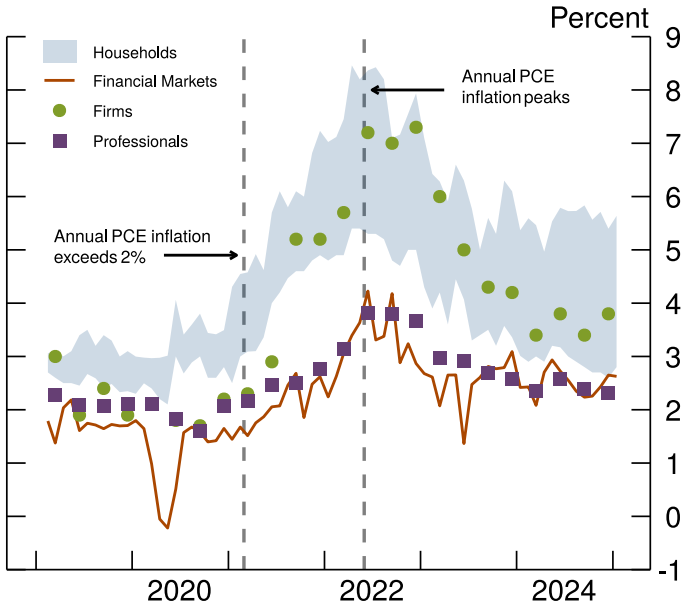
[34] See Hajdini et al. (2022). Pilossoph and Ryngaert (2024) also document that consumers with higher inflation expectations are more likely to search and change jobs.

[35] Corroborating this pattern, Glick and Leduc (2022) find that heightened inflation expectations lead to greater wage growth as workers demanded higher wages. Moreover, as mentioned in section 2, Phillips curve models based only on labor market tightness would have trouble fitting the magnitude of the increase in wage growth seen in recent years. The findings discussed here also help explain the increase in wage growth over this period. Bernanke and Blanchard (2025), for example, find that including short- term inflation expectations in their model helps improve its ability to fit observed wage growth. Importantly, Barlevy et al. (2024) argue that part of the wage growth arising from an increase in on-the-job search results in higher labor costs and hence inflation.

[36] In a study of Italian firms, Coibion, Gorodnichenko, and Ropele (2020) have found that increases in the inflation expectations of firms can cause them to raise prices and lower employment when the ELB is not binding. In another study of Swiss firms, Abberger et al. (2024) find that a rise in firms' short-term inflation expectations causes them to raise prices. Finally, Baumann et al. (2024) find that firms in the euro area plan to increase prices in response to higher inflation expectations.

with their anticipated future price changes.[37] Combined together, the evidence above suggests that the rise in households' and firms' short-term inflation expectations, following the inflation surge, could have contributed to the persistence of inflation, by influencing wage and price setting dynamics.

Figure 8: Evolution of various measures of one-year-ahead annual inflation expectations



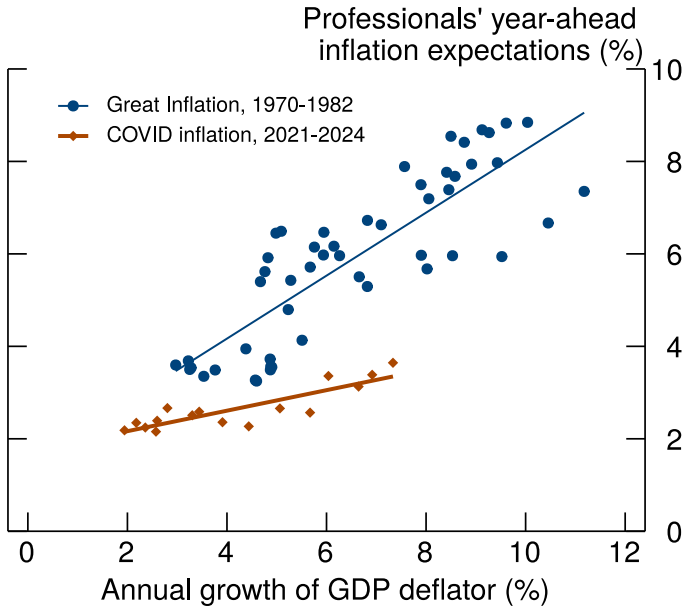
Sources: Bureau of Economic Analysis; Survey of Firms' Inflation Expectations, Federal Reserve Bank of Cleveland; Survey of Consumer Expectations, Federal Reserve Bank of New York; Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia; Indirect Consumer Inflation Expectations, Morning Consult; Survey of Consumers, University of Michigan; authors' calculations.
Note: The range of households' short-term inflation expectations includes several measures of consumers' year-ahead inflation expectations.[38] Professionals' inflation expectations refer to the mean CPI forecasts from the Survey of Professional Forecasters. Financial market expectations refer to the expectations constructed by the Federal Reserve Bank of Cleveland.[39]

[37] See Meyer and Sheng (2025).

[38] These measures include the median for the density and point forecasts from the Federal Reserve Bank of New York Survey of Consumer Expectations, the mean and median forecasts from the University of Michigan Survey of Consumers, and the Indirect Consumer Inflation Expectations. The indirect consumer inflation expectations indicator is constructed from measuring the income growth consumers would need to make up for one-year-ahead expected changes in the price index associated with their current consumption basket. See Hajdini et al. (2024) for more details.

[39] This measure extracts inflation expectations from asset prices using the method of Haubrich, Pennacchi, and Ritchken (2012). Combining a model with various sources of data, such as Treasury yields, inflation data, inflation swaps, and survey-based measures of inflation expectations, the method separates inflation expectations from inflation risk premium and the real risk premium.

Figure 9: The relationship between professional forecasters' year-ahead inflation expectations and annual inflation, measured by the GDP deflator



Sources: Bureau of Economic Analysis; Federal Reserve Bank of Philadelphia (mean GDP deflator forecast from the Survey of Professional Forecasters); authors' calculations.
Note: Data is at a quarterly frequency.

The short-term inflation expectations of professional forecasters and financial markets also rose with inflation, but at a much slower pace than expectations of consumers and firms, and moved toward 2 percent relatively quickly after inflation started falling. The relatively muted response of professionals' year-ahead inflation expectations during the recent inflation surge and their subsequent quick return to values around 2 percent contrasts with their heightened sensitivity to changes in inflation during the Great Inflation in the 1970s and early 1980s, as shown in figure 9. The stark difference in the professionals' inflation expectations between the two postwar inflation episodes could reflect professionals' perception that the recent rise in inflation was transitory, or also perhaps the accumulated credibility of the Federal Reserve in keeping inflation stable in recent decades. Regarding this latter possibility, the lower sensitivity of professionals' short-term inflation expectations is consistent with our assessment that longer-term expectations have remained generally well-anchored, which we will discuss in the following subsection.[40]

3.2 The stabilizing role of longer-term inflation expectations

The post-pandemic inflation surge presented a significant test to the anchoring of longer-term inflation expectations and marked the largest deviation from the Federal Reserve's 2

[40] Carvalho et al. (2023) show that the behavior of inflation and short-term expectations can be used to infer both longer-term inflation expectations and the degree of anchoring.

percent inflation target since it was adopted in 2012. Our reading of the evidence suggests that the inflation target was effective during the recent inflation surge by promoting stability in longer-term inflation expectations. Despite large increases in shorter-term inflation expectations, longer-term inflation expectations remained generally well-anchored. The stability in longer-term inflation expectations, in turn, likely helped to contain the inflation surge and helped inflation to return toward the 2 percent target without a large increase in unemployment.

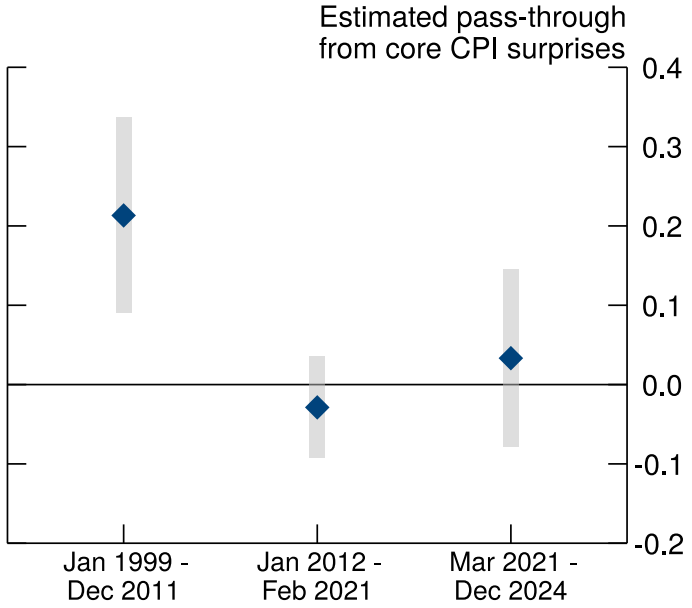
We first review some of the evidence that longer-term inflation expectations remained anchored in recent years before discussing the role anchoring played in containing the inflation surge and facilitating the decline in inflation.

Economists have often focused on the concept of anchoring as laid out by then Chairman Bernanke when he described anchored expectations as a situation in which longer-term inflation expectations are relatively insensitive to incoming data.[41] One line of research has put this concept into practice by measuring the degree of anchoring by the pass-through from surprises in data releases to financial market measures of longer-term inflation compensation.[42] Zero or low pass-through is consistent with well-anchored inflation expectations. Figure 10 shows results from Bundick and Smith (2024), which measures the pass-through from CPI inflation surprises to far-forward, longer-term inflation compensation over three distinct time periods. This analysis provides evidence that the 2012 inflation target helped to better anchor longer-term inflation expectations. Importantly, the improved degree of anchoring was largely maintained during the recent inflation surge.

[41] Bernanke, Ben S. 2007. "Inflation Expectations and Inflation Forecasting." Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts.

[42] For studies focusing on the U.S., see, for example, Gurkaynak, Levin, and Swanson (2010), Beechey, Johannsen, and Levin (2011), and Bundick and Smith (forthcoming).

Figure 10: Estimates of the sensitivity of longer-term inflation compensation to inflation data over different periods



Sources: Bundick and Smith (2024); updated to present by authors.

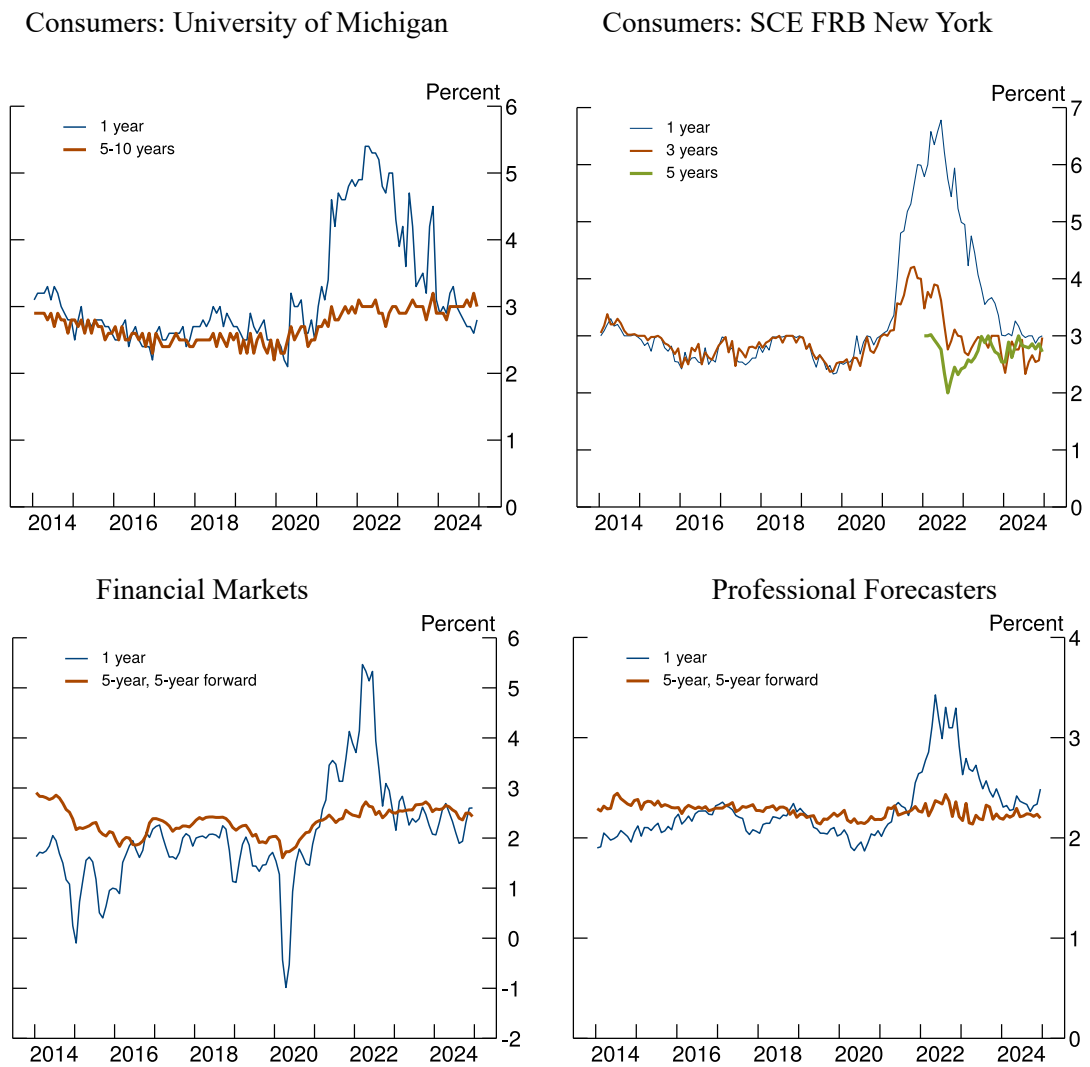
Note: The pass-through is measured from core CPI surprises to 1-year forward inflation compensation maturing in 10 years, derived from TIPS. A core inflation surprise is defined as the reported month-over-month percent change in core CPI less the Bloomberg median forecast ahead of the release. The March 2021 sample break is chosen because that is when 12-month PCE inflation first moved above 2 percent.

More generally, the degree of anchoring can be assessed by comparing the movements in shorter-term and longer-term measures of inflation expectations. Figure 11 shows measures of shorter-term and longer-term inflation expectations across consumers, financial markets, and professional forecasters. For all measures, longer-term inflation expectations increased very little in recent years despite large increases in shorter-term inflation expectations, consistent with generally well anchored inflation expectations.[43, 44]

[43] Lebow and Peneva (2024) use data from the University of Michigan Survey of Consumers to show that household perceptions of *current* inflation increased above their expectations for inflation 12 months ahead, suggesting that households thought inflation would come down over time.

[44] Beginning in early 2025, there was a large increase in the University of Michigan's measures of near-term and longer-term consumer inflation expectations. The increase in longer-term inflation expectations has not yet been corroborated in other survey or financial market measures and comes on the heels of a methodological change in the University of Michigan survey in late 2024.

Figure 11: Comparisons of shorter-term and longer-term inflation expectations across consumers, financial markets, and professional forecasters



Sources: Bureau of Economic Analysis via Haver Analytics; Survey of Consumers, University of Michigan; Federal Reserve Board; Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia; authors' calculations.

Across various measures, it appears that longer-term inflation expectations remained generally well anchored in recent years. This favorable outcome was likely supported by aggressive monetary policy actions. Recent research covering this period has shown that as inflation moved further above target and the FOMC began raising interest rates, perceptions of the private sector shifted to anticipate a strong Federal Reserve response to inflation.[45] The large

[45] See Bauer, Pflueger, and Sunderam (2024b), as well as Bundick, Smith, and Van der Meer (2024).

and rapid rate increases in the federal funds rate carried out in 2022 likely informed these perceptions and demonstrated the commitment of the Committee to the 2 percent inflation target.[46] While this evidence is not conclusive, it does seem to suggest that the conduct of monetary policy likely helped to prevent longer-term inflation expectations from becoming unanchored.

The stability of longer-term inflation expectations likely prevented larger and more lasting increases in inflation. Research has found that better-anchored inflation expectations can reduce the volatility and persistence of inflation.[47] Intuitively, if longer-term inflation expectations are unanchored, then otherwise transitory price pressures can be propagated through increases in inflation expectations. It is generally believed that this mechanism contributed to the Great Inflation of the 1970s and early 1980s. For example, some research looking at the 1970's and 80's has argued that the unanchoring of inflation expectations, together with short-lived supply shocks, lifted inflation expectations, allowing high inflation to persist and become embedded in the economy through price- and wage-setting behavior.[48] While the supply and demand shocks outlined in section 2 that led to the pandemic inflation surge were themselves large and persistent, the stability of longer-term inflation expectations likely helped to limit the extent and duration of elevated inflation.

Anchored longer-term inflation expectations may have also helped inflation to return toward 2 percent without a large increase in unemployment. Historically, disinflations have been costly in terms of unemployment because prolonged periods of economic slack were necessary to reduce longer-term inflation expectations. Given this experience, the post-2022 disinflation was unique because inflation declined without a large increase in unemployment. While supply improvements and nonlinearities in the Beveridge and Phillips curves likely played a key role in reducing inflation without necessitating economic slack, the anchoring of longer-term inflation expectations likely contributed to this outcome as well. This recent experience demonstrates the labor market benefits of maintaining well-anchored inflation expectations.

4. Lessons from the pandemic and inflation dynamics going forward

In this section, we review some lessons learned from the recent inflation surge and identify several indicators that could be helpful in identifying future inflationary pressures in a

[46] Bauer, Pflueger, and Sunderam (2024a) explicitly link the Committee's policy rate actions to perceptions about the perceived monetary policy rule, and the weights placed on inflation and unemployment. Some models of expectations formation explicitly link the central bank's reaction function (and the responsiveness to inflation) to the degree of anchoring (See, for example, Jorgensen and Lansing, 2021). In such models, an increase in the perceived response of the central bank to deviations of inflation from target would help to anchor inflation expectations. Knoteck, Mitchell, Pedemonte, and Shiroff (2024) provide some supportive evidence for this mechanism from a randomized control trial.

[47] See, for example, Orphanides and Williams (2004).

[48] See, for example, Ball and Mazumder (2011) for some evidence. Ireland (2007) also provides empirical estimates suggesting that short-lived supply shocks during this time were allowed to feed through to longer-term inflation expectations.

timely manner.[49] We also discuss how changes in the economy and our experience with inflation over the past five years could alter the assessment of inflation risks going forward.

4.1 Lessons learned from the pandemic inflation experience

The COVID-19 pandemic was an extraordinary event, both in its rarity and the magnitude of the economic shocks it triggered, which warrants against drawing too many conclusions from such a unique period. Nevertheless, we think there are two lessons that can be carried forward.

First, the pandemic period showed how the inflation environment can change abruptly. The extreme events driving the pandemic inflation surge revealed the potential for large shocks to generate severe supply and demand imbalances. The aggregate supply curve can appear flat but steepen quickly when large increases in demand encounter supply constraints, resulting in sharp increases in inflation. In this sense, the pandemic revealed information about inflation dynamics in an economy with severe supply and demand imbalances.

Given our better understanding that inflationary pressures may build suddenly and persist, timely indicators of supply–demand imbalances may be valuable to policymakers. These include vacancies, which vary by more than the unemployment rate when the unemployment rate is low, measures of strain in the market for inputs such as the GSCPI and the QPC, and measures of housing market rents, which are leading indicators of rent measures used in the PCE price index.[50]

Recent inflation dynamics were also amplified by nonlinearities. However, monitoring for nonlinearities in real-time is challenging, making it difficult to determine whether they are affecting the economy. The frequency of price adjustment could be a particularly useful real-time indicator.[51] Thresholds on the unemployment rate or the V-U ratio, such as that suggested in Benigno and Eggertsson (2024), may also be informative but should be used with caution because of limited data and challenges with identification of nonlinearities in a Phillips curve setting.

Second, the recent inflation surge reinforces the importance of keeping longer-term inflation expectations well anchored. Longer-term inflation expectations were thought to be generally well anchored heading into the inflation surge; however, their continued stability through such a severe stress scenario is noteworthy. While monetary policy likely played an important role in maintaining the anchor, our understanding of how inflation expectations are formed and influenced remains limited. Therefore, continued stability in the face of future large

[49] For lessons from the recent inflation surge regarding monetary policy strategies, see Chung et al. (2025).

[50] Foote et al. (2025) discuss various labor market indicators and how they can be used to assess the labor market's position relative to the Committee's full employment goal.

[51] Creating timely measures of the frequency of price change would require cooperation with the Bureau of Labor Statistics, as the necessary microdata are not publicly available.

disturbances is far from assured and measures of inflation expectations deserve continued monitoring.

4.2 Why inflation risks might be more balanced than was thought five years ago

At the time of the last framework review, several factors appeared to impart downside risks to the inflation outlook, including the experience with below-target inflation and risks posed by the ELB.[52] In this final section, we discuss why future risks to inflation might now be more balanced.

Inflation outcomes over the past five years revealed the potential for large shocks to the economy to trigger significant inflation surges. While it remains unclear if the underlying distribution of inflationary shocks has changed since 2019, our understanding of the potential risks has evolved. In particular, there is a greater appreciation for the fragility of supply chains and the role of capacity constraints in amplifying inflation. These supply effects are likely stronger and more persistent when they coincide with monetary and fiscal stimulus. Moreover, synchronized disruptions across geographies or shortages of essential inputs can pose a particular upside risk to inflation, given the limited substitutability of affected suppliers.

The behavior of short-term inflation expectations after the pandemic may also impart upside inflation risks. While longer-term inflation expectations remained well anchored, some research suggests that the recent period of high inflation may have contributed to inflation becoming more salient in the minds of the public.[53] Financial markets are more sensitive to inflation news and recent studies suggest that the public might be more attentive to inflation than they were before the pandemic.[54] Shorter-term inflation expectations may be more sensitive to inflationary news, and particularly news of higher prices, after a period of high inflation, although it is uncertain to what extent such changes will persist.[55] The evidence of feedback between actual and expected inflation discussed in section 3 suggests that any sustained changes

[52] See Gourio, Johannsen, and Lopez-Salido (2025).

[53] Studies by Pfauti (2024) and Weber et al. (2025) show that consumers' and firms' attention to inflation rose significantly during the recent inflation surge. Moreover, Gennaioli et al. (2024) show that inflationary shocks can prompt retrieval of past high-inflation experiences, which would then lead to a strong response of inflation expectations to inflationary news. The authors document that this mechanism was quantitatively important to explain the fast surge in short-term inflation expectations across all consumers in the U.S. Moreover, Stantcheva (2024) documents that in two waves of surveys in December 2023 and January 2024, consumers ranked inflation to be the top economic and social issue. Finally, Hilscher, Raviv, and Reis (2024) estimate the perceived probability of a future inflation disaster, that is, of future inflation exceeding 4 percent, to have stabilized after the inflation surge but at higher than pre-pandemic levels.

[54] Kroner (2025) shows that financial markets attention to inflation news remained significantly higher than before the pandemic during the disinflation period. Korenok and Munro (2024) provide suggestive evidence that public's attention to inflation remains high after inflation returns to pre-surge values.

[55] For instance, D'Acunto et al. (2021) show that the short-term inflation expectations of US households are more responsive to positive grocery price changes than negative grocery price changes. Cavallo, Cruces, and Perez-Truglia (2017) found a similar result.

to the behavior of short-term inflation expectations could be important for inflation dynamics going forward.[56]

Finally, reduced effective lower bound (ELB) risk could also be important in determining the risk to inflation going forward. At the time of the previous framework review, heightened ELB risk stemming from the low level of interest rates raised concerns that longer-term inflation expectation could drift lower. By some accounts, the neutral level of the federal funds rate has increased since just before the pandemic.[57] All else equal, a higher neutral rate reduces the likelihood that the funds rate will be constrained by the lower bound in the future, reducing downside risks to inflation and longer-term inflation expectations.[58] While we acknowledge that the ELB may yet remain a concern relative to previous decades when the neutral level of the funds rate was thought to be much higher, on the margin ELB risk appears to have declined somewhat since 2019. Whether the potential reduction in ELB risk is large enough to have a meaningful effect on inflation outcomes depends on the magnitude of the increase in the neutral federal funds rate, which is highly uncertain. For instance, FOMC participants had a wide range of projections for the longer-run federal funds rate in the most recent March 2025 SEP, ranging from 2.5 to 3.9 percent.

Overall, the risks to inflation appear more balanced than was perceived at the time of the last framework review, although the uncertainty of such an assessment is great. These greater upside inflation risks stem from a plausibly higher perceived likelihood of abrupt increases in inflation due to large shocks, disruptions to critical inputs, and nonlinearities; the potentially greater salience of inflation in the wake of the pandemic and its effects on inflation expectations; and the possibility that the neutral rate has moved higher since immediately before the pandemic and therefore the downside risks to inflation and longer-term inflation expectations from the ELB have possibly diminished.

[56] On the other hand, research by Armantier et al. (2022) found that, compared with the pre-pandemic period, consumer short-term inflation expectations did not become more sensitive to inflation news during the inflation surge. There is also evidence from Bundick and Smith (2024) that financial market measures of longer-term inflation compensation from financial markets do not seem to have become more responsive to inflation shocks since the pandemic.

[57] For example, estimates of the longer-run federal funds rate from the Federal Reserve Bank of New York's Survey of Primary Dealers and Survey Market Participants have increased since 2019.

[58] See Mertens and Williams (2021).

References

Abberger, Klaus, Anne K. Funk, Michael Lamla, Sarah M. Lein, and Stefanie Siegrist (2024). "The Pass-Through of Inflation Expectations into Prices and Wages: Evidence from an RCT Survey," CEPR Discussion Paper, no. 19595, https://cepr.org/publications/dp19595.

Abdelrahman, Hamza, and Luiz Edgard Oliveira (2023). "The Rise and Fall of Pandemic Excess Savings," FRBSF Economic Letter 2023-11. San Francisco: Federal Reserve Bank of San Francisco, May 8, https://www.frbsf.org/research-and-insights/publications/economic-letter/2023/05/rise-and-fall-of-pandemic-excess-savings.

Abdelrahman, Hamza, Luiz Edgard Oliveira, and Adam Hale Shapiro (2024). "The Rise and Fall of Pandemic Excess Wealth," FRBSF Economic Letter 2024-06. San Francisco: Federal Reserve Bank of San Francisco, February 26, https://www.frbsf.org/research-and-insights/publications/economic-letter/2024/02/rise-and-fall-pandemic-excess-wealth.

Adams, Brian, Lara Loewenstein, Hugh Montag, and Randal Verbrugge (2024). "Disentangling Rent Index Differences: Data, Methods, and Scope," *American Economic Review: Insights,* vol. 6 (June), pp. 230–45, https://doi.org/10.1257/aeri.20220685.

Aladangady, Aditya, David Cho, Laura Feiveson, and Eugenio Pinto (2022). "Excess Savings during the COVID-19 Pandemic," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, October 21, https://doi.org/10.17016/2380-7172.3223.

Amiti, Mary, Sebastian Heise, Fatih Karahan, and Ayşegül Şahin, (2024). "Inflation Strikes Back: The Role of Import Competition and the Labor Market," *NBER Macroeconomics Annual,* vol. 38, pp. 71–131, https://doi.org/10.1086/729195.

Armantier, Olivier, Gisem Koşar, Jason Somerville, Giorgio Topa, Wilbert van der Klaauw, and John C. Williams (2022). "The Curious Case of the Rise in Deflation Expectations," Federal Reserve Bank of New York Staff Reports 1037. New York: Federal Reserve Bank of New York, October (revised November 2023), https://www.newyorkfed.org/research/staff_reports/sr1037.

Ball, Laurence, and Sandeep Mazumder (2011). "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity,* Spring, pp. 337–405, https://www.brookings.edu/wp-content/uploads/2016/07/2011a_bpea_ball.pdf.

Balleer, Almut, and Marvin Noeller (2023). "Monetary policy in the presence of supply constraints: Evidence from German firm-level data," Ruhr Economic Papers 1060, https://www.rwi-essen.de/en/publications/scientific/ruhr-economic-papers/detail/monetary-policy-in-the-presence-of-supply-6524.

Barlevy, Gadi, R. Jason Faberman, Bart Hobijn, and Ayşegül Şahin (2024). "The Shifting Reasons for Beveridge Curve Shifts," *Journal of Economic Perspectives,* vol. 38 (Spring), pp. 83–106, https://doi.org/10.1257/jep.38.2.83.

Barnichon, Regis, and Adam Hale Shapiro (2022). "What's the Best Measure of Economic Slack?" FRBSF Economic Letter 2022-04. San Franscisco: Federal Reserve Bank of San Francisco, February 22, https://www.frbsf.org/wp-content/uploads/sites/4/el2022-04.pdf.

——— (2024). "Phillips Meets Beveridge," *Journal of Monetary Economics,* vol. 148, Supplement (November), 103660, https://doi.org/10.1016/j.jmoneco.2024.103660.

Barnichon, Regis, Luiz Edgar Oliveira, and Adam Hale Shapiro (2021). "Is the American Rescue Plan Taking Us Back to the '60s?" FRBSF Economic Letter 2021-27. San Franscisco: Federal Reserve Bank of San Francisco, October 18, https://www.frbsf.org/wp-content/uploads/el2021-27.pdf.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam (2024a). "Changing Perceptions and Post-Pandemic Monetary Policy," paper presented at "Reassessing the Effectiveness and Transmission of Monetary Policy," a symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyo., August 22–24, https://www.kansascityfed.org/Jackson%20Hole/documents/10337/pflueger_jh.pdf.

——— (2024b). "Perceptions about Monetary Policy," *Quarterly Journal of Economics,* vol. 139 (November), pp. 2227–78, https://doi.org/10.1093/qje/qjae021.

Baumann, Ursel, Annalisa Ferrando, Dimitris Georgarakos, Yuriy Gorodnichenko, and Timo Reinelt (2024). "SAFE to Update Inflation Expectations? New Survey Evidence on Euro Area Firms," NBER Working Paper Series 32504. Cambridge, Mass.: National Bureau of Economic Research, May, https://www.nber.org/papers/w32504.

Beaudry, Paul, Chenyu Hou, and Franck Portier (2025). "On the Fragility of the Nonlinear Phillips Curve View of Recent Inflation," CEPR Discussion Paper, no. 19950, https://cepr.org/publications/dp19950.

Beechey, Meredith J., Benjamin K. Johannsen, and Andrew T. Levin (2011). "Are long-run inflation expectations anchored more firmly in the Euro area than in the United States?" *American Economic Journal: Macroeconomics*, vol. 3 (2), pp. 104-129, https://doi.org/10.1257/mac.3.2.104.

Bernanke, Ben S. (2007). "Inflation Expectations and Inflation Forecasting," speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Mass., July 10, https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm.

Bernanke, Ben, and Olivier Blanchard (2025). "What Caused the US Pandemic-Era Inflation?" *American Economic Journal: Macroeconomics,* vol. 17 (3), pp. 1-35, https://doi.org/10.1257/mac.20230195.

Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan (2024a). "The Inflation Accelerator," NBER Working Paper Series 32531, Cambridge, Mass.: National Bureau of Economic Research, May, https://doi.org/10.3386/w32531.

Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan (2024b). "Nonlinear Inflation Dynamics in Menu Cost Economies," NBER Working Paper Series 32094, Cambridge, Mass.: National Bureau of Economic Research, January, https://doi.org/10.3386/w32094.

Blinder, Alan S., Michael Ehrmann, Jakob de Haan, and David-Jan Jansen (2024). "Central Bank Communication with the General Public: Promise or False Hope?" *Journal of Economic Literature*, vol. 62 (2), pp. 425-457, https://doi.org/10.1257/jel.20231683.

Bocola, Luigi, Alessandro Dovis, Kasper Jørgensen, and Rishabh Kirpalani (2024). "Bond Market Views of the Fed," NBER Working Paper Series 32620. Cambridge, Mass.: National Bureau of Economic Research, June, https://www.nber.org/papers/w32620.

Boehm, Christoph E., and Nitya Pandalai-Nayar (2022). "Convex Supply Curves," *American Economic Review,* vol. 112 (December), pp. 3941–69, https://doi.org/10.1257/aer.20210811.

Braun, Robin, Aaron Flaaen, and Sinem Hacioglu Hoke (2024). "Supply vs. Demand Factors Influencing Prices of Manufactured Goods," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, February 23, https://doi.org/10.17016/2380-7172.3465.

Bundick, Brent, and A. Lee Smith (2024). "Despite High Inflation, Longer-Term Inflation Expectations Remain Well Anchored," FRBKC Economic Bulletin. Kansas City: Federal Reserve Bank of Kansas City, May 31, https://www.kansascityfed.org/research/economic-bulletin/despite-high-inflation-longer-term-inflation-expectations-remain-well-anchored.

——— (forthcoming). "Did the Federal Reserve Break the Phillips Curve? Theory & Evidence of Anchoring Inflation Expectations," *Review of Economics and Statistics.* https://doi.org/10.1162/rest_a_01357.

Bundick, Brent, A. Lee Smith, and Luca Van der Meer (2024). "Maintaining the Anchor: An Evaluation of Inflation Targeting in the Face of COVID-19," Research Working Paper 24-15. Kansas City: Federal Reserve Bank of Kansas City, December, http://doi.org/10.18651/RWP2024-15.

Calvo, Guillermo A. (1983). "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, vol. 12 (3), pp. 383-398, https://doi.org/10.1016/0304-3932(83)90060-0.

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston (2023). "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, vol. 15 (January), pp. 1–47, https://doi.org/10.1257/mac.20200080.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara (2024). "Large Shocks Travel Fast," *American Economic Review: Insights,* vol. 6 (December), pp. 558–74, https://doi.org/10.1257/aeri.20230454.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia (2017). "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments," *American Economic Journal: Macroeconomics*, vol. 9 (July), pp. 1–35, https://doi.org/10.1257/mac.20150147.

Chung, Hess, Callum Jones, Antoine Lepetit, and Fernando Martin (2025). "Implications of Inflation Dynamics for Monetary Policy Strategies," Finance and Economics Discussion Series 2025-072. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.072.

Clemens, Jeffrey, Joshua D. Gottlieb, and Adam Hale Shapiro (2014). "How Much Do Medicare Cuts Reduce Inflation?" FRBSF Economic Letter 2014-28. San Francisco: Federal Reserve Bank of San Francisco, September 22, https://www.frbsf.org/research-and-insights/publications/economic-letter/2014/09/medicare-cuts-reduce-inflation-budget-control-act.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele (2020). "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics,* vol. 135 (February), pp. 165–219, https://doi.org/10.1093/qje/qjz029.

Comin, Diego A., Robert C. Johnson, and Callum J. Jones (2023). "Supply Chain Constraints and Inflation," NBER Working Paper Series 31179. Cambridge, Mass.: National Bureau of Economic Research, April (revised August 2024), https://doi.org/10.3386/w31179.

Considine, Sabrina, Brandon Miskanic, Nicolas Petrosky-Nadeau, and Deepika Baskar Prabhakar (2025). "What's Driving Labor Force Participation among Women?" FRBSF Economic Letter 2025-04. San Francisco: Federal Reserve Bank of San Francisco, February 10, https://www.frbsf.org/research-and-insights/publications/economic-letter/2025/02/whats-driving-labor-force-participation-among-women.

Curst, Erin, E., Kevin J. Lansing, and Nicolas Petrosky-Nadeau (2023). "Reducing inflation along a nonlinear Phillips Curve," FRBSF Economic Letter. San Francisco: Federal Reserve Bank of San Francisco, July 10, https://www.frbsf.org/research-and-insights/publications/economic-letter/2023/07/reducing-inflation-along-nonlinear-phillips-curve/.

Cuba-Borda, Pablo and Jean-Paul L'Huillier (2025). "Inflation is a Supply Phenomenon." Working Papers 137, Brandeis University, Department of Economics and International Business School, https://www.brandeis.edu/economics/RePEc/brd/doc/brandeis_wp137.pdf.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber (2021). "Exposure to Grocery Prices and Inflation Expectations," Journal of Political Economy, vol. 129 (May), pp. 1615–39, https://doi.org/10.1086/713192.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim (2023). "Pandemic-Era Inflation Drivers and Global Spillovers," NBER Working Paper Series 31887. Cambridge, Mass.: National Bureau of Economic Research, November (revised February 2025), https://doi.org/10.3386/w31887.

Dogra, Keshav, Sebastian Heise, Edward S. Knotek II, Brent Meyer, Robert W. Rich, Raphael S. Schoenle, Giorgio Topa, Wilbert van der Klaauw, and Wändi Bruine de Bruin (2023). "Estimates of Cost-Price Passthrough from Business Survey Data," Working Paper Series 23-14. Cleveland: Federal Reserve Bank of Cleveland, June, https://doi.org/10.26509/frbc-wp-202314.

Doser, Alexander, Ricardo Nunes, Nikhil Rao, and Viacheslav Sheremirov (2023). "Inflation Expectations and Nonlinearities in the Phillips Curve," *Journal of Applied Econometrics,* vol. 38 (June/July), pp. 453–71, https://doi.org/10.1002/jae.2963.

Dupor, Bill, and Marie Hogan (2025). "Was the Post-Lockdown Inflation Surge Mainly Supply Driven?" Working Paper 2025-007. St. Louis: Federal Reserve Bank of St. Louis, April, https://doi.org/10.20955/wp.2025.007.

Duncan, Roberto, Enrique Martínez García, and Luke Miller (2025). "Tempting FAIT: Flexible Average Inflation Targeting and the Post-COVID U.S. Inflation Surge," Working Paper 2511. Dallas: Federal Reserve Bank of Dallas, April, https://doi.org/10.24149/wp2511.

Duzhak, Evgeniya (2024). "Recent Spike in Immigration and Easing Labor Markets," FRBSF Economic Letter 2024-19. San Francisco: Federal Reserve Bank of San Francisco, July 15, https://www.frbsf.org/research-and-insights/publications/economic-letter/2024/07/recent-spike-in-immigration-and-easing-labor-markets.

Ferrante, Francesco, Sebastian Graves, and Matteo Iacoviello (2023). "The Inflationary Effects of Sectoral Reallocation," *Journal of Monetary Economics,* vol. 140, Supplement (November), pp. S64–S81, https://doi.org/10.1016/j.jmoneco.2023.03.003.

Figura, Andrew, and Chris Waller (2022). "What Does the Beveridge Curve Tell Us about the Likelihood of a Soft Landing?" FEDS Notes. Washington: Board of Governors of the Federal Reserve System, July 29, https://doi.org/10.17016/2380-7172.3190.

Foote, Christopher, Shigeru Fujita, Amanda Michaud, and Joshua Montes (2025). "Assessing Maximum Employment," Finance and Economics Discussion Series 2025-067. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.067.

Gallin, Joshua, Lara Loewenstein, Hugh Montag, and Randall Verbrugge (2024). "Sticky Continuing-Tenant Rents," paper prepared for the 2024 North American Meeting of the Urban Economics Association, held in Washington, D.C., September 21, https://www.bls.gov/pir/journal/mh01.pdf.

Gennaioli, Nicola, Marta Leva, Raphael Schoenle, and Andrei Shleifer (2024). "How Inflation Expectations De-Anchor: The Role of Selective Memory Cues," NBER Working Paper

Series 32633. Cambridge, Mass.: National Bureau of Economic Research, June, https://doi.org/10.3386/w32633.

Giannone, Domenico, and Giorgio Primiceri (2024). "The Drivers of Post-Pandemic Inflation," NBER Working Paper Series 32859. Cambridge, Mass.: National Bureau of Economic Research, August, https://doi.org/10.3386/w32859.

Glick, Reuven, and Sylvain Leduc (2022). "Will Workers Demand Cost-of-Living Adjustments?" FRBSF Economic Letter 2022-21. San Francisco: Federal Reserve Bank of San Francisco, August 8, https://www.frbsf.org/research-and-insights/publications/economic-letter/2022/08/will-workers-demand-cost-of-living-adjustments.

Glover, Andrew, José Mustre-del-Río, and Alice von Ende-Becker (2023). "How Much Have Record Corporate Profits Contributed to Recent Inflation?" FRBKC Economic Review. Kansas City: Federal Reserve Bank of Kansas City, January 12, https://doi.org/10.18651/ER/v108n1GloverMustredelRiovonEndeBecker.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, Ivan Werning (2022). "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review,* vol. 112 (May), pp. 1437–74, https://doi.org/10.1257/aer.20201063.

Gürkaynak, Refet S., Andrew Levin, and Eric Swanson (2010). "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., U.K., and Sweden," *Journal of the European Economic Association,* vol. 8 (December), pp. 1208–42, https://doi.org/10.1111/j.1542-4774.2010.tb00553.x.

Gourio, François, Benjamin K. Johannsen, and David López-Salido (2025). "The Origins, Structure, and Results of the Federal Reserve's 2019-2020 Review of Its Monetary Policy Framework," Finance and Economics Discussion Series 2025-065. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.065.

Hajdini, Ina (2023). "Mis-specified Forecasts and Myopia in an Estimated New Keynesian Model," Working Paper Series 22-03R. Cleveland: Federal Reserve Bank of Cleveland, March, https://doi.org/10.26509/frbc-wp-202203r.

Hajdini, Ina, Edward S. Knotek II, John Leer, Mathieu Pedemonte, Robert W. Rich, and Raphael S. Schoenle (2022). "Low Passthrough from Inflation Expectations to Income Growth Expectations: Why People Dislike Inflation," Working Paper Series 22-21R. Cleveland: Federal Reserve Bank of Cleveland, June, https://doi.org/10.26509/frbc-wp-202221.

——— (2024). "Indirect Consumer Inflation Expectations: Theory and Evidence," *Journal of Monetary Economics,* vol. 145, Supplement (July), 103568, https://doi.org/10.1016/j.jmoneco.2024.103568.

Harding, Martín, Jesper Lindé, and Mathias Trabandt (2023). "Understanding Post-COVID Inflation Dynamics," *Journal of Monetary Economics,* vol. 140, Supplement (November), pp. S101–S118, https://doi.org/10.1016/j.jmoneco.2023.05.012.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken (2012). "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps," *Review of Financial Studies,* vol. 25 (May), pp. 1588–1629, https://doi.org/10.1093/rfs/hhs003.

Hilscher, Jens, Alon Raviv, and Ricardo Reis (2024). "How Likely Is an Inflation Disaster?" Centre for Macroeconomics Discussion Paper Series CFM-DP2024-37. London: London School of Economics, September, https://www.lse.ac.uk/CFM/assets/pdf/CFM-Discussion-Papers-2024/CFMDP2024-37-Paper.pdf.

Ireland, Peter N. (2007). "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit and Banking,* vol. 39 (December), pp. 1851–82, https://doi.org/10.1111/j.1538-4616.2007.00091.x.

Jordà, Òscar, Celeste Liu, Fernanda Nechio, and Fabián Rivera-Reyes (2022). "Why Is U.S. Inflation Higher than in Other Countries?" FRBSF Economic Letter 2022-07. San Franscisco: Federal Reserve Bank of San Franscisco, March 28, https://www.frbsf.org/wp-content/uploads/el2022-07.pdf.

Klitgaard, Thomas, and Matthew Higgins (2023). "Spending Down Pandemic Savings Is an 'Only-in-the-U.S.' Phenomenon," Federal Reserve Bank of New York, *Liberty Street Economics* (blog), October 11, https://libertystreeteconomics.newyorkfed.org/2023/10/spending-down-pandemic-savings-is-an-only-in-the-u-s-phenomenon.

Korenok, Oleg, and David Munro (2024). "The Rockets and Feathers of Inflation Attention," http://dx.doi.org/10.2139/ssrn.4850841.

Kroner, T. Niklas (2025). "How Markets Process Macro News: The Importance of Investor Attention," Finance and Economics Discussion Series 2025-022. Washington: Board of Governors of the Federal Reserve System, March, https://doi.org/10.17016/FEDS.2025.022.

Lansing, Kevin J., Luiz Edgard Oliveira, and Adam Hale Shapiro (2022). "Will Rising Rents Push Up Future Inflation?" FRBSF Economic Letter 2022-03. San Francisco: Federal Reserve Bank of San Francisco, February 14, https://www.frbsf.org/research-and-insights/publications/economic-letter/2022/02/will-rising-rents-push-up-future-inflation.

Lansing, Kevin J. (2025). "Demand versus Supply: Which Is More Important for Inflation?" Federal Reserve Bank of San Francisco Working Paper 2025-08. https://doi.org/10.24148/wp2025-08.

Lebow, David, and Ekaterina Peneva (2024). "Inflation Perceptions during the Covid Pandemic and Recovery," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, January 19, https://doi.org/10.17016/2380-7172.3439.

Leduc, Sylvain, Huiyu Li, and Zheng Liu (2024). "Are Markups Driving the Ups and Downs of Inflation?" FRBSF Economic Letter 2024-12. San Francisco: Federal Reserve Bank of San Francisco, May 13, https://www.frbsf.org/research-and-insights/publications/economic-letter/2024/05/are-markups-driving-ups-and-downs-of-inflation.

Lipińska, Anna, Enrique Martínez Garcia, and Felipe Schwartzman (2025). "Pandemic and War Inflation: Lessons from the International Experience," Finance and Economics Discussion Series 2025-071. Washington: Board of Governors of the Federal Reserve System, August, https://doi.org/10.17016/FEDS.2025.071.

Loewenstein, Lara, Jason Meyer, and Randal J. Verbrugge (2024). "New-Tenant Rent Passthrough and the Future of Rent Inflation," Economic Commentary 2024-17. Cleveland: Federal Reserve Bank of Cleveland, October, https://doi.org/10.26509/frbc-ec-202417.

Luo, Shaowen, and Daniel Villar (2023). "Propagation of Shocks in an Input-Output Economy: Evidence from Disaggregated Prices," *Journal of Monetary Economics,* vol. 137 (July), pp. 26–46, https://doi.org/10.1016/j.jmoneco.2023.05.004.

Mertens, Thomas M., and John C. Williams (2021). "What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices," *American Economic Review,* vol. 111 (August), pp. 2473–2505, https://doi.org/10.1257/aer.20181461.

Meyer, Brent H., and Xuguang Simon Sheng (2025). "Unit Cost Expectations: Firms' Perspectives on Inflation," *European Economic Review,* vol. 174 (May), 104955, https://doi.org/10.1016/j.euroecorev.2025.104955.

Mondragon, John A., and Johannes Wieland (2022). "Housing Demand and Remote Work," NBER Working Paper Series 30041. Cambridge, Mass.: National Bureau of Economic Research, May, https://doi.org/10.3386/w30041.

Montag, Hugh, and Daniel Villar (forthcoming). "Post-Pandemic Drivers of Price Setting."

Najjar, Rami, and Adam Hale Shapiro (2025). "Does Monetary Policy Tightening Reduce Inflation?" FRBSF Economic Letter 2025-03. San Francisco: Federal Reserve Bank of San Francisco, February 3, https://www.frbsf.org/research-and-insights/publications/economic-letter/2025/02/does-monetary-policy-tightening-reduce-inflation.

Orphanides, Athanasios, and John C. Williams (2004). "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in Ben S. Bernanke and Michael Woodford, eds., *The Inflation-Targeting Debate.* Chicago: University of Chicago Press, pp. 201–46, https://www.nber.org/system/files/chapters/c9559/c9559.pdf.

Peneva, Ekaterina V., and Jeremy B. Rudd (2017). "The Passthrough of Labor Costs to Price Inflation," *Journal of Money, Credit and Banking,* vol. 49 (December), pp. 1777–1802, https://doi.org/10.1111/jmcb.12449.

Pfauti, Oliver (2024). "The Inflation Attention Threshold and Inflation Surgers," EMPCT Working Paper Series, no. 2024-02.

Pilossoph, Laura, and Jane Ryngaert (2024). "Job Search, Wages, and Inflation," NBER Working Paper Series 33042. Cambridge, Mass.: National Bureau of Economic Research, October, https://doi.org/10.3386/w33042.

Romer, Christina D., and David H. Romer (2023). "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years," *American Economic Review,* vol. 113 (June), pp. 1395–1423, https://doi.org/10.1257/aer.113.6.1395.

Shapiro, Adam Hale (forthcoming). "Decomposing Supply- and Demand-Driven Inflation," *Journal of Money, Credit and Banking*.

Smith, Simon C., Allan Timmermann, and Jonathan H. Wright (2024). "Nonlinear Phillips Curves," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, September 4, https://doi.org/10.17016/2380-7172.3596.

Stantcheva, Stefanie (2024). "Why Do We Dislike Inflation?" *Brookings Papers on Economic Activity,* Spring, pp. 1–45, https://www.brookings.edu/wp-content/uploads/2024/03/1_Stantcheva_unembargoed.pdf.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce (2025). "Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings," *Econometrica,* vol. 93 (January), pp. 229–64, https://doi.org/10.3982/ECTA22764.

Journal of Economic Perspectives—Volume 11, Number 1—Winter 1997—Pages 51–72

What We Know and Do Not Know About the Natural Rate of Unemployment

Olivier Blanchard and Lawrence F. Katz

Almost 30 years ago, Friedman (1968) and Phelps (1968) developed the concept of the "natural rate of unemployment." In what must be one of the longest sentences he ever wrote, Milton Friedman explained: "The natural rate of unemployment is the level which would be ground out by the Walrasian system of general equilibrium equations, provided that there is imbedded in them the actual structural characteristics of the labor and commodity markets, including market imperfections, stochastic variability in demands and supplies, the cost of gathering information about job vacancies and labor availabilities, the costs of mobility, and so on." Over the past three decades a large amount of research has attempted to formalize Friedman's long sentence and to identify, both theoretically and empirically, the determinants of the natural rate. It is this body of work we assess in this paper. We reach two main conclusions.

The first is that there has been considerable theoretical progress over the past 30 years. A framework has emerged, organized around two central ideas. The first is that the labor market is a market with a high level of traffic, with large flows of workers who have either lost their jobs or are looking for better ones. This by itself implies that there must be some "frictional unemployment." The second is that the nature of relations between firms and workers leads to wage setting that often differs substantially from competitive wage setting. This in turn is likely to result in a level of unemployment that differs from the efficient frictional level. We present

■ *Olivier Blanchard is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Lawrence F. Katz is Professor of Economics, Harvard University, Cambridge, Massachusetts. Both authors are Research Associates, National Bureau of Economic Research, Cambridge, Massachusetts. Their e-mail addresses are blanchar@mit.edu and katz@nber.harvard.edu, respectively.*

this framework below and then take it through its paces, showing how it helps to think about the complex relation between technological progress and the natural rate.

The second conclusion, however, is that empirical knowledge sadly lags behind. Economists are a long way from having a good quantitative understanding of the determinants of the natural rate, either across time or across countries. We review the standard equations used in empirical macroeconomic models to characterize the natural rate. We then take up two issues. Following the work of Blanchflower and Oswald (1995), the first is known as the issue of the "wage curve" versus the "Phillips curve": it is about the form of the dynamic relation between wages and unemployment. Far from an esoteric issue, it has major implications for which variables do or do not affect the natural rate. Yet the empirical evidence is far from settled. The second issue is the rise in European unemployment since the 1970s. The magnitude of the increase, the diversity of labor market institutions and of policies should have helped identify culprits. Yet while many suspects have been identified, none has been convicted.

We believe that progress can come from a better integration of the empirical work pursued by labor economists with some of the work developed by macroeconomists, and we end up with a plea for such joint work.

Thinking about the Natural Rate

Flows

The labor market in advanced economies is in a constant state of flux: some firms are expanding employment, others are reducing it. Firms are constantly changing their mix of positions, laying off some workers and hiring new ones. Workers quit in search of better opportunities, forcing firms to find replacements.[1] In the United States, flows of workers to or from employment are each equal to roughly 3 percent of employment per month, with about half of these flows coming from or going to unemployment, and the other half coming from or going to "out of the labor force."[2]

Such a volume of traffic clearly implies some frictional unemployment. But there is no reason in general to expect that the natural rate will be equal to the efficient frictional level. Even leaving aside standard distortions such as unemployment benefits or taxes on labor income, wages are unlikely to be set in such a way as to generate the efficient level of unemployment.

[1] Much progress has been made in the recent past in our knowledge of the distribution of changes in employment across firms over time, thanks in large part to the work of Davis, Haltiwanger and Schuh (1996).

[2] Most of our knowledge of flows of workers between labor market states in the United States comes from the Current Population Survey (CPS). The numbers given in the text are the CPS numbers adjusted for measurement error, which otherwise generates spuriously high estimated flows (Blanchard and Diamond, 1990).

Wage Setting

Three main approaches have been followed to think about wage determination.

The first has focused on the fact that, by the very nature of the labor market, both workers and firms typically have some bargaining power. The bargaining power of workers stems from the fact that they cannot be costlessly and instantaneously replaced. The bargaining power of firms arises because most workers cannot costlessly and instantaneously locate an equivalent job. Taking this as a starting point, the *matching approach*, following in particular Diamond (1982) and Pissarides (1990), has formalized wage determination as the result of bargaining (typically bilateral Nash bargaining) between each worker and each firm. This approach has delivered a number of important insights.[3]

In a depressed labor market, knowing that finding another job is likely to be difficult, workers will be willing to settle for a lower wage. The bargained wage will therefore be low, close to the workers' reservation wage (that is, the wage that makes a worker indifferent to being employed or unemployed). Conversely, in tight labor markets, the bargained wage will be much higher than the reservation wage. Thus, even if the reservation wage is constant, the wage will vary with labor market conditions.

The right measure of the state of the labor market is the exit rate from unemployment, defined as the number of hires divided by the number unemployed, rather than the unemployment rate itself. What matters to the unemployed is not how many of them there are, but how many of them there are in relation to the number of hires by firms.

Factors that affect the pace of reallocation, the matching process of workers and jobs, or the nature of bargaining will all affect wage determination. For example, at a given unemployment rate, increases in "churning"—increases in the flows of separations and hires by firms—will lead to an increase in the exit rate from unemployment and thus lead to an increase in the wage. To take another example, restrictions on firing, which partially insulate incumbent workers from wage pressure from the unemployed, will likely lead to a higher wage and decrease the role of labor market conditions in wage determination.

The second approach has focused instead on the complexity of transactions between firms and workers and their implications for wage setting. The starting point of the *efficiency wage approach* is that wages affect productivity. It follows that firms may want to pay workers more than their reservation wage in order to get and to keep better workers, to economize on turnover costs, or to motivate greater effort and cooperation from their workforce (Katz, 1986). Efficiency wage models imply that factors such as the organization of firms (and thus the nature of internal labor markets) and the nature of jobs (for example, the extent to which firms must trust workers to do the job right) potentially affect wage setting and, through it, affect the natural rate of unemployment.

[3] For recent models and applications along these lines, see for example Mortensen and Pissarides (1994) or Blanchard and Diamond (1990).

Specific theories range from those with a strong neoclassical flavor, such as Shapiro and Stiglitz (1984), to those that emphasize fuzzier (but not necessarily less relevant) notions, such as fairness (Akerlof and Yellen, 1990). At the more neoclassical end, the "shirking" model of Shapiro and Stiglitz leads to a wage relation in which the tighter the labor market—that is, the higher the exit rate from unemployment—the higher is the wage the firms have to pay to prevent shirking, an implication similar to that of matching models. At the other end, models based on fairness suggest that the reservation wage may depend on factors such as the level and the rate of growth of wages in the past, if workers have come to consider such wage increases as fair. Perhaps a better word than "reservation" wage in that context is "aspiration" wage, a word that sometimes appears in the literature (Oswald, 1986).

The third approach, known as the *competitive approach*, has emphasized that some unemployment is simply due to poor labor market opportunities. Especially at the bottom end of the skill distribution, workers have little or no bargaining power because they can be replaced easily. Efficiency wage considerations may also play much less of a role as many jobs are routine and differences in worker quality make less of a difference. In this case, how many unskilled workers want to work may simply depend on their reservation wages relative to their labor market opportunities. If labor opportunities are poor, unskilled workers will simply prefer not to work, or at least not to work all the time.

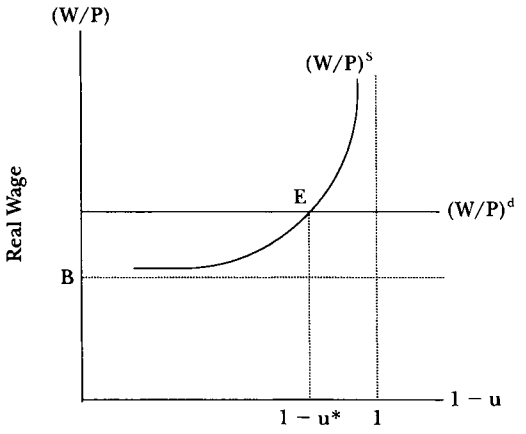
This approach suggests that the right measure of labor market conditions, at least for those in the low-wage labor market, may not be the unemployment rate, but rather the nonemployment rate. People facing poor opportunities relative to their reservation wages are likely to spend some of their time in and out of employment. During periods in which they are out of work, they may characterize themselves in household surveys as unemployed or out of the labor force, depending on the nature of the unemployment benefits system (Card and Riddell, 1993), but the distinction may be of little economic significance.

A Graphical Representation

The theories of wage determination we have just sketched differ in many ways. Nonetheless, they can be conveniently represented in a graph like Figure 1, which has the real wage on the vertical axis, where W is the nominal wage and P the price level, and one minus the unemployment rate on the horizontal axis.

Let us first consider what we call the "supply wage relation" in Figure 1, shown by the upward sloping curve $(W/P)^s$. All the models discussed so far have a basic implication: the tighter the labor market, the higher the real wage, given the reservation wage. This relation should be tailored to fit each theory. For example, as we saw earlier, the labor market variable suggested by matching models is the exit rate from unemployment rather than the unemployment rate itself. Under the competitive approach, the correct variable is nonemployment rather than unem-

Figure 1
The Determination of the Natural Rate



ployment. We shall treat the unemployment rate as a stand-in for the relevant labor market variable and ignore these differences here.

Consider now the "demand wage relation," represented by the horizontal line $(W/P)^d$ in Figure 1, which shows the real wage consistent with the employment decisions of firms. If we focus for simplicity on the medium run—that is, the run over which firms can adjust all factors of production including capital—we can think of the real wage paid by firms as independent of the level of employment (or equivalently, the unemployment rate). The simplest interpretation of this relation is as a long-run labor demand curve, giving the real wage consistent with other input prices and the condition that firms make zero pure profit. In this case, the "demand wage" that firms are willing to pay will be determined by the level of productivity, by the characteristics of the production function, by other input prices (including the interest rate), as well as by any factor that affects the wage firms can pay, such as payroll taxes. In imperfectly competitive goods markets, the behavior of the markup of price over marginal cost will also matter. The higher the markup, the higher the price of goods given the wage, the lower the real wage paid by firms.

The natural rate of unemployment is the rate at which the supply wage and the demand wage are equal. In Figure 1, the equilibrium is given by point E, with associated unemployment rate, u^*.[4] The natural rate of unemployment depends on

[4] For readers who prefer to see these relations in algebraic form, the supply wage function can be written as $W/P = B\,g(u, X_s)$; $g_u > 0$, where W is the nominal wage, P the price level, so that W/P is the real wage, B is the (real) reservation wage, u is the unemployment rate, and X_s includes all the factors that affect wage determination. The demand wage relation can be written as $W/P = A\,f(X_d)$, where A measures the level of total factor productivity (under the assumption of Harrod neutral technological progress), and X_d includes all the other factors that affect the real wage set by firms as described in the text. Combining these equations, the natural rate of unemployment is such that $g(u^*, X_s) = (A/B)f(X_d)$.

the level of productivity in relation to the reservation wage, as well as all the factors just mentioned that affect either the supply or the demand wage.[5]

The Natural Rate in a Changing World

Discussing all the determinants of the natural rate would take us too far. To illustrate how the framework above can shed light on a number of current debates, we focus on one theme, the effects of technological progress on the natural rate.

"Pure" Productivity Growth and the Natural Rate

Start with a simple fact. Unemployment rates viewed over the very long run—for example, over the past century—appear to be untrended in most nations (Layard, Nickell and Jackman, 1991), despite tremendous increases in productivity. Thus, any model should satisfy the condition that there is no long-run effect of the level of productivity on the natural rate of unemployment.

What restrictions does this impose on the framework we developed earlier? Return to Figure 1. An increase in productivity leads to a proportional increase in the demand wage and thus an upward shift in the demand wage relation $(W/P)^d$. When will such an increase in the wage leave the natural rate unaffected? Only if the reservation wage increases in line with the real wage, leading to an upward shift in the supply wage relation $(W/P)^s$, which offsets the effect of the shift in the demand wage relation on the unemployment rate. If this is the case, an increase in the level of productivity will lead to a proportional increase in the real wage, but will not affect the natural rate.

Is this condition likely to be satisfied? The answer is yes. To see why requires a closer look at the determinants of the reservation wage. The reservation wage depends first on what the unemployed do with their time, what we often call the utility of leisure but may also include a number of nonmarket activities, including working in the grey and black economy. A plausible benchmark is that increases in productivity affect market and nonmarket activities in roughly the same way. If this is the case, this component of the reservation wage will move in line with the real wage.

The reservation wage depends on the various types of income support the unemployed can expect to receive if unemployed, from unemployment benefits to other social insurance programs, to intrafamily insurance. Again, a plausible benchmark is that all these increase roughly in line with wages and thus with productivity. In the United States, the ratio of average weekly unemployment benefits to average weekly wages (for workers covered by unemployment insurance) has been quite

[5] Two formal treatments of the determination of the natural rate—and two important references—are given in the books by Layard, Nickell and Jackman (1991) and by Phelps (1994). Shorter treatments can be found in the articles by Woodford (1994) and by Bean (1994).

stable over the last 50 years (Advisory Council on Unemployment Compensation, 1994). Other social insurance components are not automatically linked to wages, but, at least for the last 50 years, they have also roughly increased in line with wages.

The reservation wage depends on nonlabor income and access to credit. How desperate the unemployed are to work, and thus how low a wage they are willing to accept, depends not only on the availability of income support, but also on whether they can draw on their assets or borrow against the future. Here again, a plausible benchmark is that an increase in productivity leads to equal increases in labor and nonlabor income, thus again leaving the ratio of the wage to the reservation wage roughly unaffected.

We have relied so far on the narrow interpretation of the reservation wage, namely as the wage that made workers indifferent to working or not working. As we saw earlier, reservation wages may embody workers' aspirations as well as hard realities. Under that interpretation, all that is required for changes in productivity not to affect the natural rate in the long run is that aspirations eventually adjust to reality, that workers in Portugal, for example, do not hold out for German wages. Whatever exact content we give to the notion of aspiration wages, this appears to be a mild requirement, also likely to be satisfied.

To summarize, the conditions that deliver long-run neutrality of the natural rate to changes in productivity need not hold as a matter of logic, but they all appear plausible. The discussion makes clear that long-run neutrality does not preclude dynamic effects of changes in productivity, or—more realistically, in a world with positive productivity growth—of changes in the rate of productivity growth, such as we have seen since the mid-1970s. How long and how large these dynamic effects may be depends on the exact nature of the reservation wage. For example, if reservation wages reflect primarily unemployment benefits, they are likely to adjust quickly to any change in the level or in the rate of growth of wages, and the effects of changes in productivity growth on the natural rate will be short-lived. But if reservation wages reflect primarily aspirations, the adjustment may be more drawn out. Suppose, for example, that the rate of productivity growth decreases. If, at a given unemployment rate, workers keep asking for wage increases corresponding to the previous higher rate of productivity growth, lower productivity growth will lead to a higher natural rate until aspirations have adjusted to the new realities.

Technological Progress, Structural Change and the Natural Rate

We have considered so far a very rarefied form of technological progress, one that affected productivity but did not affect the organization of production in any other way. Technological progress has many more dimensions, and these are also relevant to the determination of the natural rate.

Technological progress comes with structural change. Increases in the pace of technological progress are likely to come with a higher pace of reallocation across jobs. To the extent that this leads to larger flows of workers in the labor market, these will lead to a higher unemployment rate. These effects have been recently explored using matching models, for example, by Mortensen and Pissarides (1995).

Interestingly, while the empirical evidence remains largely unexplored, it does not suggest so far a strong relation between flows in the labor market and structural change. Perhaps the most striking piece of evidence is from eastern Europe, where high structural change—measured, say, by the Lilien (1982) measure, the standard deviation of employment changes across sectors—has been associated with low flows in the labor market and a nearly stagnant unemployment pool.[6]

Technological progress may be biased toward a particular group of workers, and this may in turn affect the natural rate. It is simplest here to think in terms of the competitive model and to think of two groups of workers: the skilled and the unskilled. The unskilled are paid a wage very close to their reservation wage. The skilled are paid a wage much higher than their reservation wage. Shifts in the demand for unskilled workers imply movements along a very flat labor supply; shifts in the demand for skilled workers imply movements along a very steep labor supply. Thus, a change in technology that raises the demand for skilled workers but reduces it for unskilled workers will not have symmetric effects; instead, the rise in unemployment for the unskilled will be substantial, while the decline in unemployment for the skilled will be small. Overall, the natural rate of unemployment will increase. Juhn, Murphy and Topel (1991) have shown that this may help explain much of the rise in nonemployment for prime-age American males since the early 1970s.

Labor Market Rigidities, Sclerosis and the Natural Rate

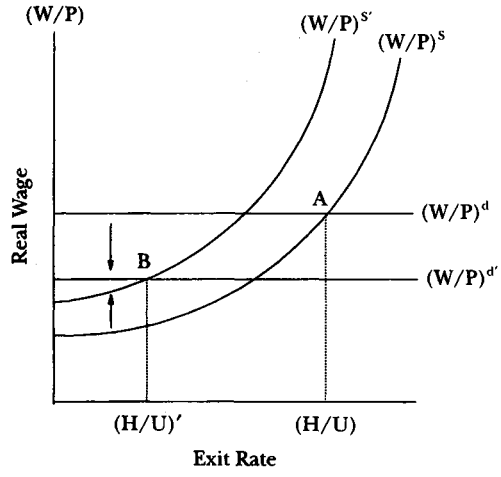
Legal and administrative restrictions on layoffs are an example of the "labor market rigidities" that are often held responsible for high European unemployment. For example, the OECD Jobs Report (1994) argues that, in an environment in which structural change and adaptation of firms is increasingly important, countries with high labor market rigidities will do poorly and have higher unemployment.

Matching models provide a useful framework for assessing this argument. They suggest that firing restrictions have three separate effects. First, firing restrictions lead firms to keep some workers they would otherwise have laid off. The result is fewer separations, and thus fewer hirings (the two must be equal in equilibrium) and smaller flows in the labor market. Second, such restrictions force firms to keep at times workers they do not need, thus decreasing average productivity and the demand wage. Third, raising the cost of firing strengthens the hand of the currently employed workers in bargaining, leading them to obtain a higher wage for given labor market conditions (for a given exit rate from unemployment).

The implications of these effects for the natural rate are shown in Figure 2. It will be more convenient here to think of the exit rate from unemployment—that is, the number of hires, H, divided by unemployment, U—as the relevant labor market variable for wage determination (as suggested by matching models). Thus the horizontal axis in Figure 2 plots the exit rate from unemployment rather than

[6] See, for example, Blanchard, Commander and Coricelli (1995) and Boeri (1996).

Figure 2
Firing Restrictions and the Natural Rate



one minus the unemployment rate as in Figure 1. The curve corresponding to the supply wage relation is still upward sloping: an increase in the exit rate from unemployment leads to an increase in the supply wage. The locus corresponding to the demand wage relation is still horizontal.

Because of lower average productivity, the demand wage decreases: the line $(W/P)^d$ shifts downward to $(W/P)^{d'}$. Because workers are stronger in bargaining, the supply wage relation shifts upward, from $(W/P)^s$ to $(W/P)^{s'}$. Thus, the equilibrium exit rate decreases from (H/U) to $(H/U)'$; equivalently, unemployment duration (the inverse of the exit rate in steady state) increases. This does not however imply that the natural unemployment rate increases. This is because the flow of traffic in the labor market, and thus H, is also smaller. The declines in (H/U) and H have an ambiguous effect on U and thus on the rate of unemployment. In words, higher firing restrictions lead to a sclerotic labor market, in which average productivity is lower, the flows of workers are lower and the duration of unemployment is higher; but, contrary to the claims of the OECD report, they do not necessarily lead to a higher rate of unemployment.[7]

Putting the threads together, what does our analysis suggest about the effects of technological progress on the natural rate? In short, that they are ambiguous. Take an increase in the rate of technological progress. If reservation wages reflect in part aspirations, and aspirations lag behind, the first effect of faster technological progress will be to decrease the natural rate for some time. But if faster technological

[7] Models that explore various implications of firing costs for employment and unemployment include Bentolila and Bertola (1990) and Hopenhayn and Rogerson (1993).

progress leads to faster churning, this will work the other way: other things equal, larger traffic in the labor market will lead to higher unemployment. And so will composition effects, if technological progress is biased against lower-skilled workers, as seems to be the case at this point. These effects do not add up to an unambiguous or simple answer. But the analysis, which builds on nearly all aspects of the framework we presented earlier (flows and matching, the nature of bargaining, competitive factors at the lower end of the skill distribution, efficiency wages and the role of aspirations) represents, we believe, progress.

However, this is the end of the good news. Suggesting a list of potential determinants of the natural rate is not the same as identifying their respective roles. As we now argue, empirical knowledge lags very much behind.

The Macroeconometric Evidence

Let us start by reviewing the standard characterization of the joint behavior of wage inflation, price inflation and unemployment in U.S. macroeconometric models. Macroeconometric models determine the natural rate through two equations, a "price equation" (which corresponds conceptually to what we called the demand wage relation) and a "wage equation" (which corresponds to the supply wage relation). In their simplest form, the equations can be written as

$$\Delta p_t = a_p + \Delta w_t + \epsilon_{pt}$$

$$\Delta w_t = a_w + \Delta p_{t-1} - \beta u_t + \epsilon_{wt},$$

where p is the logarithm of the price index, w is the logarithm of the nominal wage, so that Δp and Δw are price and wage inflation respectively, u is the unemployment rate, a_p and a_w are constants, and ϵ_p and ϵ_w are error terms.

Substituting the term for wage inflation from the second equation into the first equation, these equations yield the (expectations-augmented) "Phillips curve" reduced-form relation between the change in inflation and the unemployment rate

$$\Delta p_t = a + \Delta p_{t-1} - \beta u_t + \epsilon_t,$$

where $a = a_p + a_w$, and $\epsilon_t = \epsilon_{wt} + \epsilon_{pt}$. The natural rate of unemployment is typically interpreted as the rate of unemployment consistent with constant (nonaccelerating) inflation in this Phillips curve equation.

The first three columns of Table 1 report the results of estimating these relations using annual U.S. data for 1970–1995, with the consumer price index as a measure of the price level and compensation per hour in the business sector as a measure of the wage level. Even in their simplest form, these three equations fit the data quite well. Obviously, they can be improved upon. The actual wage equations of large macroeconometric models include richer lag structures, as well as other variables, such as trend productivity growth and wage/price control dummies. Actual price equations typically specify price inflation not only as a function of wage

Table 1
Price Inflation, Wage Inflation and Unemployment in the United States, 1970–1995

	Δp	Δw	Δp	Δw
Constant	−0.77	7.20	6.64	6.75
	(0.28)	(1.40)	(1.65)	(1.31)
Δp_{-1}		1.00[a]	1.00[a]	1.00[a]
Δw	1.00[a]			
$(w_{-1} - p_{-1}) - x_{-1}$				0.11
				(0.05)
Unemployment Rate (u)		−0.99	−1.02	−0.95
		(0.21)	(0.24)	(0.19)
Durbin-Watson Statistic	1.26	1.18	1.48	1.38
$\bar{R}^2$	0.76	0.67	0.71	0.72

[a] Indicates that the coefficient is constrained to equal 1.
Notes: All regressions include 26 annual observations for 1970 to 1995. $p = 100 \times \log$ (Consumer Price Index), using the CPI-U; $w = 100 \times \log$ (Compensation per hour in the business sector); u is the civilian unemployment rate measured in percentage points. $x = \log$ (cyclically adjusted labor productivity) = predicted value from a regression of log (business sector output per hour) on a linear time trend and its square for the 1969 to 1994 period.
Source: Citibase.

inflation, but also of productivity growth and the rate of change of other input prices. A state-of-the-art specification of the reduced-form Phillips curve is given, for example, by Gordon in his paper for this issue. It includes longer lags of past inflation, a slowly moving natural rate, price control dummies, and the rate of change of food and energy prices. But the basic structure of these equations remains close to those of the bare-bones specifications presented in the first three columns of Table 1, so that we can use these simpler equations as a basis for discussion.

We take up two issues in the next two sections. The first one is that of the consistency of the standard macroeconometric wage equation linking wage inflation to the unemployment rate with the demand wage relation suggested by theory. The second is why these empirical macroeconometric relations that have worked reasonably well for the United States have failed so miserably for Europe over the last 20 years. This failure motivates us to summarize what we think has been learned (or not learned) from the search for causes of the rise of European unemployment.

Wages and Unemployment

If we interpret lagged inflation as a proxy for expected inflation in the wage equation above, we can think of the empirical wage equation as giving the *rate of change* of the (expected) real wage as a function of the unemployment rate. Is such

a relation consistent with the supply wage relation implied in our framework and illustrated by the $(W/P)^s$ locus in Figure 1?

The answer would appear to be no. Recall that the supply wage relation implies, conditional on the reservation wage, a relation between the *level* of the real wage and the unemployment rate. But the reservation wage is not observable. Thus, what turns out to be crucial here is the exact nature of the reservation wage and the implied empirical relation between wages and unemployment when one cannot explicitly condition on the reservation wage.

Suppose, to take an extreme case, that the reservation wage is really an aspiration wage, and that this aspiration wage in turn is based on the real wage paid last year. This will generate exactly the relation observed in the data (ignoring the difference between expected and actual price inflation): the rate of change of the wage will depend on the unemployment rate. But this is indeed an extreme case. In general, as we argued earlier, reservation wages will reflect in part factors that grow with productivity. Then the relation will be one in which the real wage is equal, or at least returns over time, to a level determined by both the level of productivity and the unemployment rate.

This suggests that the wage equation found in macroeconomic models may be misspecified and that a more appropriate specification is

$$\Delta w_t = a_w + \Delta p_{t-1} - \lambda(w_{t-1} - p_{t-1} - x_{t-1}) - \beta u_t + \epsilon_{wt},$$

where x is the log of the productivity level. If λ is positive, this equation implies that while it may take some time, the real wage eventually adjusts to a level determined by productivity and the unemployment rate.[8] The standard macroeconometric wage equation can be seen as imposing the possibly incorrect restriction that λ is equal to zero.

To see whether this is the case, we estimate this augmented empirical wage equation and report the results in the last column of Table 1. The variable x is constructed as the fitted value from a first-stage regression of the logarithm of labor productivity on a quadratic time trend. The point estimate of λ is actually *negative* (-0.11 with a standard error of 0.05), suggesting the absence of any long-term level relation between the real wage and the unemployment rate.[9]

But this does not quite settle the issue. One may question whether we can trust the estimated value of λ in column 4. Issues of specification are many and have been discussed at length in the empirical literature. They range from the treatment of expected inflation, to the measurement of the underlying productivity growth trend, to the noninclusion of other possibly highly autocorrelated wage determinants, and so on. In an important contribution, Blanchflower and Oswald (1994)

[8] This error correction specification was first proposed and estimated in this context by Sargan (1964).
[9] This result is consistent with estimation by others. Grubb (1986) estimates a similar specification for each OECD country. He also finds an estimated negative value of λ for the United States, but a positive value for most other countries. Recent estimates by Turner, Richardson and Rauffet (1996) report a small but positive value of λ for the United States (0.08).

have suggested one way of avoiding some, if not all, of these issues, by looking at the behavior of wages and unemployment rates across regions (or states) within a country over time. Their argument is simple and compelling. If many of these difficult-to-measure variables, such as expectations of inflation or the underlying productivity growth trend, are common across regions within a nation, then they can be dealt with by using time effects in a panel data regression model.

Let us be more precise here. Consider the n U.S. states, indexed by i, and assume that for each one, the relation between wages and unemployment has the following form:

$$\Delta w_{it} = a_{wi} + \Delta p_{t-1} - \lambda(w_{i,t-1} - p_{t-1} - x_{t-1}) - \beta u_{it} + \epsilon_{wit}.$$

This equation is a straightforward extension to the state level of the augmented wage equation just given above and estimated on aggregate data in column 4 of Table 1. The rate of change in the nominal wage in state i is a function of lagged inflation (assumed common to all states), the lagged log real wage adjusted for productivity (where productivity is assumed to be common to all states) and the state-specific unemployment rate.[10] This equation can be estimated on panel data by replacing all the terms that vary through time but are common to all states by time effects, and thus running the change in the log nominal wage in each state on the lagged log nominal wage, the state unemployment rate, time and state dummies. Note also that by writing $\Delta w_{it} = w_{it} - w_{it-1}$ and moving w_{it-1} to the righthand side, this equation can be rewritten in the following form, which makes the results in Table 2 below easier to interpret:

$$w_{it} = a_{wi} + (1 - \lambda) w_{i,t-1} - \beta u_{it} + d_t + \epsilon_{wit},$$

where the d_t are time fixed effects. This equation says that, other things equal, the wage w_{it} depends on the lagged wage w_{it-1} with coefficient $1 - \lambda$.

Estimating this equation using yearly data for states of the United States, and for regions in the United Kingdom, Blanchflower and Oswald (1994) estimate λ to be close to one (equivalently, $1 - \lambda$ to be close to zero). They conclude that once region and time effects are included, there is "little sign of autoregression in wage equations" (p. 284) so that the wage adjusts to the unemployment rate rapidly with little in the way of interesting dynamics.

If Blanchflower and Oswald are right, the problem is then to reconcile micro and macro empirical results. In the process of exploring their results for this paper, we have concluded however, as have Card and Hyslop (1996) and Bell (1996a), that their results concerning the lack of wage dynamics appear to be nonrobust, at least for the United States. The details of our analysis are given in Blanchard and Katz (1996). But the basic source of the difference between their results and ours

[10] As noted by Bell (1996a), the assumption that productivity trends are common to all states appears empirically questionable. So is the assumption of common inflation rates. In other words, estimation with regional data still has to confront the issue of state specific unobservables. We leave this issue aside here.

Table 2
The Relation Between Wages and Unemployment, U.S. States, 1980–1991

	Dependent variable = Log State Wage (w_{it})					
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
	Weekly UI	*Hourly Manuf.*	*Hourly ORG CPS*	*Hourly ORG CPS*	*Hourly March CPS*	*Annual March CPS*
Lagged log wage (w_{it-1})	0.978	0.953	0.911	0.920	0.633	0.258
	(0.010)	(0.020)	(0.017)	(0.017)	(0.030)	(0.039)
Log unemployment rate (u_{it})	−0.034	−0.019	−0.041	−0.039	−0.057	−0.062
	(0.002)	(0.004)	(0.003)	(0.003)	(0.008)	(0.017)
Adjusted wage[a]	No	No	Yes	Yes	Yes	No
State average worker controls[b]	No	No	No	No	No	Yes
Standard error ($\hat{\sigma}$)	0.008	0.011	0.011	0.011	0.028	0.042

Notes: All regressions except column 2 contain 612 observations (51 states × 12 years) for 1980 to 1991; column 2 contains 561 observations for 1980 to 1990. All regressions include state and year fixed effects. The estimates are obtained by weighted OLS using as weights the relative number of workers in each state in each year. Standard errors are in parentheses. i indexes state; t indexes time.

[a] The dependent variable (w_{it}) and the lagged dependent variable (w_{it-1}) are log stage average wages adjusted for characteristics of workers in the state using year-specific OLS log earnings regressions (using CPS sampling weights) as a wage prediction model; w_{it} is the mean wage residual (coefficient on state dummy variable) for state s in year t from such a cross-section log hourly earnings regression. Individual-level controls include schooling dummies, a quartic in experience, a female dummy, interactions of the quartic in experience with the female dummy, marital status, marital status interacted with female, a nonwhite dummy, and a full-time dummy in column 3. The adjustment equation for 4 also includes 10 one-digit industry dummies. The adjustment equation for 5 is described in Card and Hyslop (1996).

[b] Includes average worker characteristics by state and year from the Merged Outgoing Rotation Groups of the Current Population Survey (CPS) analogous to individual controls in wage adjustment equation used for column 4.

Sources of Wage Measures: Column 1, log (average weekly earnings in employment covered by unemployment insurance), U.S. Department of Labor (1996); 2, log (average hourly earnings of production workers in manufacturing), BLS; 3 & 4, adjusted log hourly earnings, Merged Outgoing Rotation Group Files of CPS; 5, adjusted log hourly earnings from March CPS, Card and Hyslop (1996); 6, average log annual earnings, March CPS, Card and Hyslop (1996). Annual state unemployment rates are from BLS, LABSTAT.

is, we believe, their use of annual earnings from the March Current Population Survey (CPS) as their basic wage measure. This annual earnings measure deviates from a theoretically appropriate wage rate measure since it mixes up actual changes in wages with changes in hours of work. The March CPS samples are also relatively small for measuring annual state average wages. Systematic measurement error from the use of an inappropriate wage measure and sampling error in the lagged wage measure both bias down the coefficient of the lagged log wage in a regression of the current log wage on lagged log wage and other covariates (equivalently, they lead to an upward biased estimate of λ).

Both problems can be corrected by using more appropriate measures of the

wage and larger data sets. A sampling of our estimates from regressions of the log wage on the lagged log wage, the (log) unemployment rate, state and year dummies for U.S. states over the 1980–1991 period is presented in Table 2.[11] The estimates in the first four columns of Table 2 use administrative data on average weekly earnings of covered workers from the U.S. unemployment insurance system, data on average hourly earnings of manufacturing production workers from a large-scale establishment survey and data on hourly earnings from the Merged Annual Outgoing Rotation Groups of the CPS. These regressions consistently yield estimates of the coefficient on the lagged log wage close to 1 (0.91 to 0.97) implying a value of λ close to 0 (0.03 to 0.09). As shown in column 4, such large coefficients on the lagged log wage remain even when adjusting the wage measure for detailed individual worker characteristics.[12] The last two columns show the effects of using the March CPS (column 5) and annual earnings as the measure of wages (column 6)—the sample and the earnings measure used by Blanchflower and Oswald. The use of the March CPS sample reduces the estimated coefficient on the lagged wage to 0.63. The use of annual earnings further reduces the coefficient to 0.26, an estimate close to that obtained by Blanchflower and Oswald.

To summarize, estimates from state regressions suggest that there is a long-run relation between wages and unemployment, but that, in contrast to the results of Blanchflower and Oswald, the adjustment to this level is slow. Why have we spent a section on the dynamics of the wage equation? Why does it matter whether λ is close to zero or instead close to one? For two reasons.

The first is that the dynamics of the wage equation determine the dynamic effects of variables, such as oil price shocks and payroll taxes, which are included in the demand wage and supply wage relations discussed earlier, on the natural rate of unemployment. The second is that the discussion shows the extent of our ignorance about wage behavior and, in turn, about the exact nature of the reservation wage. Does this dependence of wages on lagged wages reflect the effects of an individual's own lagged wage, or of the lagged wages of others in the relevant labor market, or of (unmeasured) variables correlated with lagged wages? We do not know. The most interesting piece of evidence here is Card's (1990) study of wage determination in Canadian union contracts. Card uses the interaction of unanticipated inflation with partial wage indexation to decompose real wage changes during the course of a contract into "intended" and "unintended" components. He finds that intended and unintended changes in real wages over the course of the previous contract have similar effects on wages in the next contract: a 10 percent intended or unintended increase in real wages in the previous contract is associated with a 4 percent higher wage in the new contract. This suggests that the actual own lagged wages, not just unobservables correlated with lagged wages, play a role in

[11] We use the log rather than the level of the unemployment rate to be consistent with Blanchflower and Oswald's specifications.

[12] Bell (1996b) finds similar results using weekly wage measures for Canada, Germany and the United Kingdom.

affecting current wages. Other pieces of evidence are relevant here, such as the evidence that workers who have lost high-wage jobs appear to hold out for high wages, even when jobs with such high wages are unlikely to be offered to them (Kruse, 1988; Summers, 1986). The issue strikes us as an important topic for research.

The Puzzle of European Unemployment

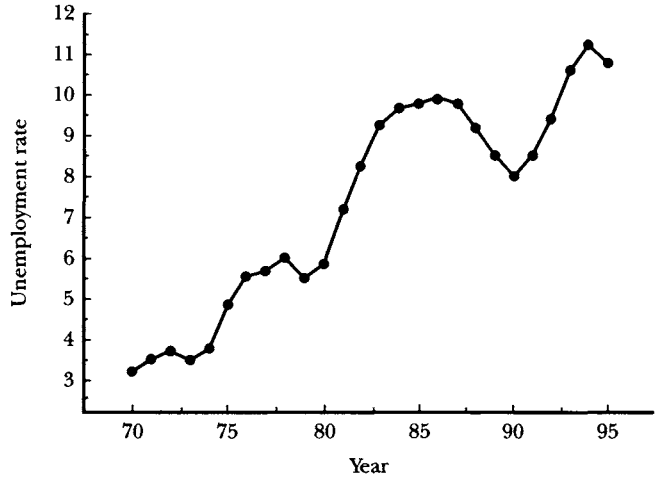
The evolution of the unemployment rate in OECD-Europe since 1970 is shown in Figure 3. The unemployment rate is defined as the total number of unemployed in all OECD-Europe divided by the total labor force. Since European inflation in the last five years has been roughly constant, there is a little doubt that most of the rise in actual unemployment corresponds to a large increase in the natural rate of unemployment. One would have thought that such an increase in the natural rate, together with the wealth of Europe's cross-sectional variation in institutions and outcomes, would allow economists to identify which variables affect or do not affect the natural rate. It is fair to say that this has not been the case.[13]

The initial line of research, started in the late 1970s, focused on the effects of the increase in the price of energy and the slowdown in total factor productivity growth (Bruno and Sachs, 1985). But while this line of argument can explain much of what happened in the late 1970s, it cannot easily explain why the natural rate kept increasing in the 1980s and has remained so high in the 1990s. The price of oil has now returned to the level of the early 1970s. And it is hard to believe that aspirations of workers have not adjusted by now to the reality of lower productivity growth.

Another popular explanation put forward in the mid-1980s was that high taxes were the source of high unemployment. The motivation was clear: tax rates had increased substantially since the 1960s, and tax rates were much higher in Europe than in the United States. Our framework indeed suggests that some taxes can have a permanent effect on the natural rate. For example, if part of the reservation wage reflects untaxed activities, then higher payroll or income taxes will decrease the ratio of the wage to the reservation wage, leading to an increase in the natural rate of unemployment. In terms of Figure 1, an increase in payroll taxes will lead to a downward shift in the horizontal demand wage schedule but will not generate as large a downward shift in the supply wage schedule, thereby generating an increase in u^*. The role of taxes was a main focus of a multicountry study organized by Richard Layard and Steve Nickell in

[13] For detailed and careful assessments of progress to date, read the surveys by Bean (1994) and by Elmeskov (1993). The technical parts of the OECD Jobs Study (1994, Parts I and II) also review theories and facts in great detail.

Figure 3
Unemployment in OECD-Europe



the mid-1980s. The effort was judged only partially successful at the time (Bean, Layard and Nickell, 1986) and has not withstood the test of time well. The cross-sectional evidence within Europe does not reveal much correlation between tax rates and unemployment rates, nor between changes in tax rates and changes in unemployment.[14] In the recent study by Jackman, Layard and Nickell (1996) summarizing their interpretation of the increase in European unemployment, taxes no longer appear in the equation used to explain time series and cross-country movements in unemployment.

Since the early 1990s, the focus has turned to three other explanations for high European unemployment: labor market rigidities, skill-biased technological progress and hysteresis.

As noted earlier, the major theme of the OECD (1994) Jobs Study is that labor market rigidities explain much of the rise in unemployment in Europe. The argument developed in the official report is that widening globalization has strengthened competition and increased the need for flexibility. Thus, while legal restrictions on firing and hiring have not become more stringent—indeed, they have typically become less stringent—the actual cost that such restrictions impose on firms has substantially increased. Despite the OECD endorsement, this approach faces conceptual and empirical problems. At the theoretical level, as we saw earlier, while rigidities can indeed increase cost and lead to labor market sclerosis, it is not

[14] A dissenting view is presented by Phelps (1994), who argues that there are both good theoretical reasons as well as substantial empirical evidence for the proposition that higher payroll taxes increase the natural rate.

clear that they lead to high unemployment. At the empirical level, the evidence on job and worker turnover collected so far does not suggest (somewhat surprisingly) much increase in the pace of reallocation and in labor market turnover. Finally, the cross-country evidence on the relation of unemployment to rigidities is less than fully supportive. For example, although Spain and Portugal are classified by the OECD as having the most stringent legal restrictions on layoffs in Europe, Spain's unemployment rate is equal to nearly four times that of Portugal (Blanchard and Jimeno, 1995).

The second strand of research has examined the role of skill-biased technological progress. The argument is that for the last 20 years or so, technological progress has been associated with an acceleration in the rate at which labor demand shifts in favor of skilled relative to unskilled workers. Compared to the United States, Europe has higher relative minimum wages, higher social benefits for the nonemployed and stronger pressures for wage compression. Thus, the argument goes, this skill-biased technological shift (typically viewed as arising from the computer revolution) has led to a larger increase in the unemployment of the unskilled and in overall unemployment in Europe than it has in the United States (Krugman, 1994). Rising unemployment and slow private sector employment growth in Europe and rising wage inequality and declining real earnings of the less skilled in the United States are viewed as flip sides of the same coin. Research here is just starting, and it is too early to tell how this approach will fare. What is clear is that the hypothesis works well when looking at two observations, the United States and Europe. It is less clear that it can explain the cross-country differences within Europe (Card, Kramarz and Lemieux, 1996; Freeman and Katz, 1995; Nickell and Bell, 1996).

The difficulty in finding a common explanation for what appears as a largely common increase in the natural rate across European countries has led some economists, including one of the present authors (Blanchard and Summers, 1986), to explore explanations based on "hysteresis"—the notion that the history of unemployment itself may have long-lasting effects on the natural rate.[15] From this perspective, various shocks to the European economy, including the oil shocks of the 1970s and the tight monetary policies of the 1980s, led to high actual unemployment. By the late 1980s, inflation was low, but sustained high unemployment for a decade had led in turn to an increase in the natural rate of unemployment.

There are a number of potential channels for hysteresis in unemployment. A long period of high actual unemployment leads to an increase in the proportion of long-term unemployed. If the long-term unemployed either lose skills or become less effective in their search, their effect on bargaining will decrease, leading to a

[15] Formally, the word "hysteresis" should be used only for the case where the natural rate is truly history dependent, so that the time series process for the unemployment rate contains a unit root. Using hysteresis to refer to long-lasting rather than permanent effects of the history of actual unemployment rates on the natural rate is a minor abuse of language.

higher natural rate of unemployment (Blanchard, 1991). Sociological factors may also increase the reservation wage of some of the unemployed: a long period of unemployment changes society's attitudes toward the unemployed. It becomes more socially acceptable to be unemployed and to use existing benefits to their utmost (Lindbeck, 1995). Family insurance may develop to help the unemployed, in particular the unemployed youth. Hysteresis may also come from the political response to unemployment. Higher prolonged unemployment creates pressure for government policies to offer more generous programs aimed at helping the unemployed. These changes decrease the pain, but they are also likely to increase the natural rate in the process.

The empirical case for hysteresis is far from tight, however. At the macroeconomic level, a recent paper by Ball (1996) presents evidence that those OECD countries that experienced larger or longer disinflations in the 1980s, especially those with both large disinflations and generous unemployment benefits, have also experienced larger increases in the natural rate of unemployment from 1980 to 1990. The most natural explanation is one based on hysteresis. But the evidence on the relative and absolute importance of the specific channels for hysteresis is weaker. For example, whether higher unemployment leads over time to more generous unemployment benefits appears to be empirically ambiguous (Di Tella and McCullough, 1995; Saint-Paul, 1995). While there is evidence of loss of skills, morale and search effectiveness by the long-term unemployed, it remains to be shown that these can generate the degree of hysteresis needed to explain the evolution of European unemployment since the mid-1980s.

Conclusions

We have argued that while substantial conceptual progress has been made in thinking about the natural rate of unemployment, empirical knowledge lags behind. Economists are a long way from having a good quantitative understanding of the determinants of the natural rate, either across time or across countries.

Burgeoning empirical research on rich micro data sets has improved our understanding of the importance in wage determination and the behavior of the unemployed of some of the mechanisms highlighted by alternative models of the natural rate. Research on industry and establishment wage differentials has shed some light on the relative importance of efficiency wage, bargaining and competitive factors for wage determination in different parts of the labor market (Krueger and Summers, 1988; Abowd and Lemieux, 1993). Much has been learned concerning the effects of the generosity and potential duration of unemployment benefits on the job search behavior of the unemployed and the duration of unemployment (Katz and Meyer, 1990; Meyer, 1990). New work is exploring the roles of unemployment insurance and intrafamily insurance in

providing support for the unemployed and potentially affecting the evolution of reservation wages (Gruber, 1997; Gruber and Cullen, 1996). The research on the wage curve by Blanchflower and Oswald (1994) and our own work on regional labor market dynamics (Blanchard and Katz, 1992) highlight the possibility of using state and regional data as testing grounds for macroeconomic models of the labor market.

Nevertheless, we feel that a divide has grown between macroeconomists and labor economists, at least on this side of the Atlantic. Too much theoretical work on the natural rate of unemployment by macroeconomists is divorced from microeconomic evidence, and too much microeconometric work seems in search of a broader theoretical framework for interpretation. This is unhealthy. We thus end with a plea for more joint efforts by macro and labor economists to better integrate theoretical and empirical work on wage determination and unemployment.

■ *We thank Marianne Bertrand for superb research assistance and David Card for providing us with earnings data by state from the March Current Population Surveys. We are grateful to Charlie Bean, David Blanchflower, David Card, Peter Diamond, Brad De Long, Richard Freeman, Juan Jimeno, Alan Krueger, Richard Layard, Greg Mankiw, Steve Nickell, Andrew Oswald, Edmund Phelps, Gilles Saint-Paul, Robert Solow and Timothy Taylor for comments.*

References

Abowd, John, and Thomas Lemieux, "The Effects of Product Market Competition on Collec(tive Bargaining Agreements: The Case of Foreign Competition in Canada," *Quarterly Journal of Economics,* November 1993, *108,* 983–1014.

Advisory Council on Unemployment Compensation (ACUC), *Report and Recommendations.* Washington, D.C.: ACUC, 1994.

Akerlof, George, and Janet Yellen, "The Fair Wage-Effort Hypothesis and Unemployment," *Quarterly Journal of Economics,* 1990, *105*:2, 255–84.

Ball, Laurence, "Disinflation and the NAIRU." NBER Working Paper No. 5520, 1996.

Bean, Charles, "European Unemployment: A Survey," *Journal of Economic Literature,* June 1994, *32,* 573–619.

Bean, Charles, Richard Layard, and Stephen Nickell, eds., "Unemployment," *Economica,* 1986, *53,* Supplement, S1–22.

Bell, Brian, "Wage Curve or Phillips Curve?," mimeo, Nuffield College, Oxford, 1996a.

Bell, Brian, "Wither the Wage Curve?," mimeo, Nuffield College, Oxford, 1996b.

Bentolila, Samuel, and Giuseppe Bertola, "Firing Costs and Labour Demand: How Bad is Eurosclerosis?," *Review of Economic Studies,* June 1990, *57,* 381–402.

Blanchard, Olivier, "Wage Bargaining and Unemployment Persistence," *Journal of Money, Credit and Banking,* August 1991, *23,* 277–92.

Blanchard, Olivier, and Peter Diamond, "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity,* 1990, *2,* 85–155.

Blanchard, Olivier, and Juan Jimeno, "Structural Unemployment: Spain versus Portugal," *American Economic Review,* May 1995, *85,* 212–18.

Blanchard, Olivier, and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity,* 1992, *1,* 1–76.

Blanchard, Olivier, and Lawrence Katz, "Wages and Unemployment: A Reassessment of the Wage Curve," mimeo, 1996.

Blanchard, Olivier, and Lawrence Summers, "Hysteresis and the European Unemployment Problem." In Fischer, Stanley, ed., *NBER Macroeconomics Annual.* Vol. 1, Cambridge: Massachusetts Institute of Technology Press, 1986, pp. 15–78.

Blanchard, Olivier, Simon Commander, and Fabrizio Coricelli, "Unemployment and Restructuring in Eastern Europe." In Commander, S., and F. Coricelli, eds., *Unemployment, Restructuring, and the Labor Market in Eastern Europe and Russia.* Washington, D.C.: World Bank, 1995, pp. 289–330.

Blanchflower, David, and Andrew Oswald, *The Wage Curve.* Cambridge, Mass. and London: Massachusetts Institute of Technology Press, 1994.

Blanchflower, David, and Andrew Oswald, "An Introduction to the Wage Curve," *Journal of Economic Perspectives,* Summer 1995, *9:*3, 153–67.

Boeri, Tito, "Labor Market Flows in the Midst of Structural Change," mimeo, OECD, 1996.

Bruno, Michael, and Jeffrey Sachs, *The Economics of Worldwide Stagflation.* Oxford: Basil Blackwell, 1985.

Card, David, "Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts," *American Economic Review,* September 1990, *80,* 669–88.

Card, David, and Dean Hyslop, "Does Inflation 'Grease the Wheels' of the Labor Market?" NBER Working Paper No. 5538, 1996.

Card, David, and W. Craig Riddell, "A Comparative Analysis of Unemployment in Canada and the United States." In Card, David, and Richard Freeman, eds., *Small Differences That Matter.* Chicago: University of Chicago Press, 1993, pp. 149–89.

Card, David, Francis Kramarz, and Thomas Lemieux, "Changes in the Relative Structure of Wages and Employment: A Comparison of the United States, Canada, and France." NBER Working Paper No. 5487, 1996.

Davis, Steven, John Haltiwanger, and Scott Schuh, *Job Creation and Job Destruction.* Cambridge, Mass.: Massachusetts Institute of Technology Press, 1996.

Diamond, Peter, "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies,* April 1982, *49,* 217–27.

Di Tella, R., and R. MacCulloch, "The Determination of Unemployment Benefits," mimeo, Oxford University, 1995.

Elmeskov, Jorgen, "High and Persistent Unemployment: Assessment of the Problem and its Causes." OECD Working Paper No. 132, Paris, 1993.

Freeman, Richard, and Lawrence Katz, eds., *Differences and Changes in Wage Structures.* Chicago: University of Chicago Press, 1995.

Friedman, Milton, "The Role of Monetary Policy," *American Economic Review,* April 1968, *58,* 1–21.

Grubb, David, "Topics in the OECD Phillips Curve," *Economic Journal,* March 1986, *96,* 55–79.

Gruber, J., "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review,* forthcoming 1997.

Gruber, J., and J. Berry Cullen, "Spousal Labor Supply as Insurance: Does Unemployment Insurance Crowd Out the Added Worker Effect?" NBER Working Paper No. 5608, 1996.

Hall, Robert, "A Theory of the Natural Rate of Unemployment and the Duration of Unemployment," *Journal of Monetary Economics,* April 1979, *5,* 153–69.

Hopenhayn, Hugo, and Richard Rogerson, "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy,* 1993, *101:*5, 915–38.

Jackman, Richard, Richard Layard, and Stephen Nickell, "Combatting Unemployment: Is Flexibility Enough?," mimeo LSE, 1996; presented at the OECD Conference on Interactions between Structural Reform, Macroeconomic Policies and Economic Performance, March 1996.

Juhn, Chinhui, Kevin M. Murphy, and Robert Topel, "Why Has the Natural Rate Increased Over Time?," *Brookings Papers on Economic Activity,* 1991, *2,* 75–142.

Katz, Lawrence, "Efficiency Wage Theories: A Partial Evaluation," *NBER Macroeconomics Annual,* 1986, *1,* 235–76.

Katz, Lawrence F., and Bruce D. Meyer, "The Impact of the Potential Duration of Unemployment Insurance Benefits on the Duration of Unemployment," *Journal of Public Economics,* February 1990, *41,* 45–72.

Krueger, Alan, and Lawrence Summers, "Efficiency Wages and the Inter-Industry Wage Structure," *Econometrica,* March 1988, *56,* 259–93.

Krugman, Paul, "Past and Prospective Causes of High Unemployment." In *Reducing Unemployment: Current Issues and Policy Options.* Kansas City, Mo.: Federal Reserve Bank of Kansas City, 1994, pp. 49–80.

Kruse, Douglas, "International Trade and the Labor Market Experience of Displaced Workers: Evidence from the Displaced Workers Survey," *Industrial and Labor Relations Review,* April 1988, *41,* 402–17.

Layard, Richard, Stephen Nickell, and Richard Jackman, *Unemployment: Macroeconomic Performance and the Labour Market.* Cambridge: Oxford University Press, 1991.

Lilien, David, "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy,* August 1982, *90,* 777–92.

Lindbeck, A., "Hazardous Welfare-State Dynamics," *American Economic Review,* 1995, *85:2,* 9–15.

Meyer, Bruce, "Unemployment Insurance and Unemployment Spells," *Econometrica,* July 1990, *58,* 757–82.

Mortensen, Dale, and Christopher Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies,* July 1994, *61:3,* 397–416.

Mortensen, Dale, and Christopher Pissarides, "Technological Progress, Job Creation and Job Destruction," mimeo, 1995.

Nickell, Stephen, and Brian Bell, "Changes in the Distribution of Wages and Unemployment in OECD Countries," *American Economic Review,* May 1996, *86,* 302–08.

OECD, *The OECD Jobs Study.* Paris: OECD, 1994.

Oswald, Andrew, "Is Wage Rigidity Caused by 'Layoffs by Seniority'?" In Beckerman, Wilfred, ed., *Wage Rigidity and Unemployment.* Baltimore, Md.: Johns Hopkins University Press, 1986, pp. 77–95.

Phelps, Edmund, "Money-Wage Dynamics and Labor-Market Equilibrium," *Journal of Political Economy,* July/August 1968, *76,* Part 2, 678–711.

Phelps, Edmund, *Structural Slumps: The Modern Equilibrium Theory of Unemployment, Interest, and Assets.* Cambridge, Mass: Harvard University Press, 1994.

Pissarides, Christopher, *Equilibrium Unemployment Theory.* Oxford: Basil Blackwell, 1990.

Saint-Paul, Gilles, "Some Political Aspects of Unemployment," *European Economic Review,* April 1995, *39,* 575–92.

Sargan, J., "Wages and Prices in the United Kingdom." In Hart et al., eds., *Econometric Analysis for National Economic Planning.* Butterworths: Colston Research Society, 1964, pp. 25–54.

Shapiro, Carl, and Joseph Stiglitz, "Equilibrium Unemployment as a Discipline Device," *American Economic Review,* June 1984, *74,* 433–44.

Summers, Lawrence, "Why is Unemployment so Very High Near Full Employment?," *Brookings Papers on Economic Activity,* 1986, *2,* 339–83.

Turner, Dave, Pete Richardson, and Sylvie Rauffet, "Modelling the Supply Side of the Seven Major OECD Economies." OECD Economics Department Working Paper No. 167, Paris, 1996.

U.S. Department of Labor, *UI Financial Data Handbook. ET Handbook 394.* Washington, D.C.: Employment and Training Administration, Unemployment Insurance Service, 1995.

Woodford, Michael, "Structural Slumps," *Journal of Economic Literature,* December 1994, *32:4,* 1784–815.

PROJECT MUSE®

Comments and Discussion

Ayşegül Şahin, Justin Wolfers



Brookings Papers on Economic Activity, Spring 2020, pp. 139-165 (Article)

Published by Johns Hopkins University Press
DOI: *https://doi.org/10.1353/eca.2020.0022*

For additional information about this article
https://muse.jhu.edu/article/787847

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

KATHARINE G. ABRAHAM
University of Maryland

JOHN C. HALTIWANGER
University of Maryland

LEA E. RENDELL
University of Maryland

How Tight Is the US Labor Market?

ABSTRACT We construct a generalized measure of labor market tightness based on the ratio of vacancies to effective searchers. Our generalized measure exhibits substantially less volatility than the standard measure defined as the ratio of vacancies to unemployment. Effective searchers include not only the unemployed but also those who are out of the labor force and the employed. These groups account for a substantial share of hires and their presence mutes the effects of the pronounced countercyclical movements in unemployment. The effective searcher measure also distinguishes different groups among the unemployed. During protracted contractions, the distribution of unemployment shifts toward the long-term unemployed, a group with lower relative search intensities, contributing to the smaller proportional increase in effective searchers as compared to the simple unemployment count. The Beveridge curve constructed using effective searchers is much more stable than the standard Beveridge curve. Further, the matching function for hires based on our generalized measure outperforms the matching function based on the ratio of vacancies to unemployment. Our approach thus reduces the unexplained residual variation required in the matching function to be consistent with real world data.

Conflict of Interest Disclosure: Katharine G. Abraham is professor of economics and survey methodology at the University of Maryland; John C. Haltiwanger is the Dudley and Louisa Dillard Professor of Economics and Distinguished University Professor of Economics at the University of Maryland; Lea E. Rendell is a graduate student in economics at the University of Maryland. Beyond these affiliations, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation. The views expressed in this paper are those of the authors and do not necessarily reflect those of the University of Maryland.

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

For a simple summary of labor market conditions, observers and analysts long have turned to the unemployment rate. Unemployment exhibits clearly cyclical behavior, rising during downturns and falling during recoveries. By this metric, the labor market was tighter at the end of 2019 than it had been for half a century. Search and matching models of the labor market (Diamond 1982; Blanchard and Diamond 1989, 1992; Mortensen and Pissarides 1994; Pissarides 2000) imply that unemployment (or more generally job searchers) should be considered together with job openings in assessing labor market tightness. In these models, a higher ratio of job openings to unemployment makes it harder to recruit workers and easier to find jobs, thus indicating that the labor market is tighter. Official US statistics on job openings are available only since 2000, but the ratio of vacancies to unemployment at the end of 2019 was substantially higher than at any point since they began to be collected.

While both the unemployment rate and the ratio of vacancies to unemployment imply that the labor market was extraordinarily tight as of the end of 2019, there is reason to question whether these standard measures in fact capture the true degree of labor market tightness. This paper will focus on the measure at the heart of conventional search and matching models—the ratio of vacancies (V) to unemployment (U). In the search and matching framework, both the job-filling rate and the job-finding rate should vary systematically with V/U. In fact, however, these rates often deviate substantially from the rates implied by calibrated matching functions with V/U as the driving variable. This was especially true during and following the Great Recession, a finding often explained by an appeal to unobserved fluctuations in matching efficiency (Elsby, Michaels, and Ratner 2015). One of this paper's central goals is to explore whether and to what extent the use of broader measures of effective vacancies and effective searchers in the matching function can improve its fit.

One reason V/U might perform poorly in tracking job-filling and job-finding rates is that the number of unemployed people may be a poor proxy for the availability of workers to fill vacant jobs. For one thing, different groups among the unemployed may be more or less attached to the labor market and more or less likely to move into employment. In a seminal paper published in *Brookings Papers on Economic Activity*, Perry (1970) noted that women's rising labor force participation and the entry of the baby boom generation into the labor force could have raised measured unemployment independently of underlying labor market conditions. In recent years, researchers have argued that the higher-than-usual share of long-term unemployment among the unemployed following the Great

Recession implied effective unemployment lower than suggested by the unemployment rate (Krueger, Cramer, and Cho 2014). Further, as noted by Blanchard and Diamond (1989) and emphasized in a number of recent empirical studies, most new hires originate from out of the labor force or from another job (job-to-job flows) rather than from unemployment. The standard model ignores these job seekers altogether. Another potentially important factor is that the search intensity of potential workers in any given category may vary over time (Krueger and Mueller 2010, 2011; Davis 2011; Hall and Schulhofer-Wohl 2018).

Simply looking at job vacancies may be problematic as well. Existing evidence implies that the intensity with which firms recruit to fill their vacancies varies over time. Along with doing a better job of measuring the number of effective searchers, accounting for this variation also could help to account for observed variation in both job-filling and job-finding rates (Davis, Faberman, and Haltiwanger 2013; Gavazza, Mongey, and Violante 2018; Mongey and Violante 2020).

Building on the search and matching literature, including several prior studies that have made use of augmented measures of either effective searchers or effective vacancies, we propose a generalized measure of labor market tightness that addresses some important limitations of the standard measures. We begin our analysis by broadening the concept of effective job searchers to account not only for job candidates drawn from among the unemployed but also from among those currently out of the labor force or already working, as well as for the heterogeneity within each of these three groups of searchers. In defining the groups used for our analysis, we disaggregate the population relatively finely by labor market status in order to capture as much of the heterogeneity in job search behavior as possible. We construct a measure of effective searchers by taking a weighted sum across twenty-two different groups within the population age sixteen and older, with the weights based on the relative base period job-finding rates for each of the different groups as a proxy for relative job search intensities.

This generalized measure of effective searchers, expressed as a share of the population, exhibits much less volatility than the unemployment rate constructed on the same basis. One reason is that searchers who are out of the labor force or employed mute the effects of the pronounced countercyclical movements in unemployment on the broader generalized measure. Second, composition effects matter. In particular, during a deep and extended contraction such as the Great Recession, the composition of unemployment shifts toward the long-term unemployed, a group with

lower relative job-finding rates. This further dampens the increase in the number of effective searchers in the broader measure during economic downturns.

One of the puzzles of recent labor market history has been the pronounced and persistent outward shift in the Beveridge curve relating job vacancies and unemployment following the Great Recession. We find that over the period from 1994 through 2019 the Beveridge curve constructed using vacancies and effective searchers is much more stable than the curve constructed using vacancies and unemployment. Further, using our generalized measure of effective searchers rather than unemployment in the measure of labor market tightness reduces the unexplained residual variation in the job-filling rate (hires per vacancy) and in the job-finding rate of unemployed workers (hires from unemployment per unemployed person). Over the 1994–2019 period, the residual unexplained variation in the job-filling rate calibrated using our generalized measure of effective searchers is only about half as large as that based on the standard measure. For the job-finding rate among the unemployed, the residual unexplained variation of the calibrated series using our generalized effective searchers measure is about a third lower than that based on the standard measure.

Our baseline analysis is conservative in that it neglects some of the important factors that a generalized measure of labor market tightness ideally would take into account. First, the relative intensity of search on the part of job seekers in any given group may change over time. Second, the intensity with which employers recruit to fill their vacant jobs also may change over time. We extend our baseline analysis with an exploratory investigation of time-varying relative search and recruiting intensities. Given the measurement challenges associated with this exercise, we view the results as suggestive rather than conclusive. Even so, our efforts further reduce the unexplained residual variation in job-filling rates. The evidence regarding the job-finding rates of the unemployed is mixed. Our accounting for time-varying recruiting intensity reduces the unexplained residual variation in job-finding rates, but our accounting for time-varying search intensity actually increases it. Better measures of both search intensity and recruiting intensity could well produce better results.

I. A Broader Perspective on Labor Market Tightness

A matching function that relates hires to unemployment and vacancies is the cornerstone of modern macro models of the labor market. Especially in recent years, however, it has been necessary to posit substantial

fluctuations in matching efficiency to account for observed fluctuations in job finding and job filling. Fluctuations in matching efficiency similarly have been offered as an explanation for shifts in the Beveridge curve that describes the inverse relationship between vacancies and unemployment (Elsby, Michaels, and Ratner 2015).

An alternative approach that we build on here is to consider whether using more suitable arguments in the matching function can eliminate the apparent instability of the standard model. Rather than equating effective searchers with the unemployed, this broader perspective on labor market tightness recognizes that the people available to fill open jobs also include those out of the labor force or in another job. Constructing an aggregate measure of effective searchers also requires a way to measure the search intensity of those in the effective searcher pool. Relatedly, a broader perspective should recognize that labor market tightness depends not only on the number of posted job openings but also on how hard employers are trying to fill those jobs.

I.A. The Pool of Effective Searchers

A lengthy literature has examined how changes in the composition of the unemployed may affect the interpretation of the official unemployment rate. One strand of the literature, launched by the seminal work of Perry (1970) and further developed by Shimer (2001), Aaronson and others (2015), and Barnichon and Mesters (2018), among others, focuses on the demographic composition of the unemployed. Another strand focuses on the relative numbers of long-term and short-term unemployed (Kaitz 1970; Krueger, Cramer, and Cho 2014). Whether because of lower search intensity, loss of human capital, or employer unwillingness to hire them (Abraham and others 2019), the long-term unemployed have lower job-finding rates and may contribute proportionately less than the short-term unemployed to the pool of effective searchers. How a person entered unemployment also may be important. As an example, the job-finding pattern among those laid off from a job differs considerably from the pattern for other groups among the unemployed (Katz 1986; Katz and Meyer 1990; Fujita and Moscarini 2017).

A comprehensive measure of effective job searchers also needs to account for people who are outside of the labor force. Those in this group have a much lower job-finding rate than the average unemployed person, but there are many more of them. In a typical month, the number of people who enter employment directly from out of the labor force is much larger than the number entering directly from unemployment

(Hornstein, Kudlyak, and Lange 2014). People who are out of the labor force but say they want a job are much more likely to enter employment than the rest of the out-of-the-labor-force population (Blanchard and Diamond 1989; Jones and Riddell 1999; Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017; Hall and Schulhofer-Wohl 2018).

The employed are a final group of searchers (Sedlacek 2016; Hall and Schulhofer-Wohl 2018). In the canonical search-and-matching model, vacancies include the job openings created by departing employees. Symmetrically, the measurement of effective searchers should take into account on-the-job searchers who may fill those jobs. Available survey data suggest that on-the-job search is prevalent (Black 1980; Blau and Robins 1990; Faberman and others 2017) and administrative data show that a large share of hires are people moving from one job to another (Haltiwanger, Hyatt, and McEntarfer 2018; Haltiwanger and others 2018).

I.B. Job Search Intensity

In addition to properly identifying the effective searcher population, a full accounting of effective search activity also needs to incorporate search intensity. Approaches used in the literature to do this include directly measuring search activities, making use of information on the gap between individuals' desired and actual hours, and the approach that we adopt—inferring relative search intensity from relative job-finding rates.

In an early effort to measure job search intensity directly, Shimer (2004) uses information from the Current Population Survey (CPS) on the number of different search methods reported by the unemployed as a search intensity proxy. Several studies, including DeLoach and Kurt (2013), Gomme and Lkhagvasuren (2015), and Mukoyama, Patterson, and Şahin (2018), measure search intensity among the unemployed using data from the American Time Use Survey (ATUS) on time devoted to job search. ATUS data are available only beginning in 2003, but Mukoyama, Patterson, and Şahin (2018) combine them with CPS data on job search methods to construct a longer search intensity series. They model the relationship of search time to the search methods reported by the unemployed, then use that estimated relationship to construct a longer search intensity series for the unemployed. Studies using ATUS data reach conflicting conclusions about whether search intensity among the unemployed is procyclical or countercyclical. Ahn and Shao (2017) use ATUS data to study the cyclicality of job search among the employed. Because the ATUS does not ask what respondents are doing while they are at work, ATUS measures of job

search among the employed seem especially likely to miss at least some job search activity.

Faberman and others (2020) use data on the gap between desired and actual hours to proxy for job search intensity. They show that, in 2013–2015 data from the Survey of Consumer Expectations (SCE), this gap is correlated with a direct measure of search intensity. They use the SCE data to calculate the average difference between desired and actual hours for each of thirty-nine groups defined based on labor force status and demographic characteristics. Treating the gaps for defined groups as constant over time, the authors use CPS data to construct a time series of aggregate slack, which they define as the total gap between desired and actual hours divided by total desired hours.

Finally, job-finding rates have been used to proxy for search intensity. The simplest version of this approach uses group-specific job-finding rates in a base period to weight the number of people in each group to produce an aggregate measure of effective searchers. A notable example of this approach is the Federal Reserve Bank of Richmond's Hornstein-Kudlyak-Lange Non-Employment Index (NEI), which uses long-run average job-finding rates to aggregate its nine groups of effective searchers among the unemployed and those out of the labor force (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017).

Several studies, including Veracierto (2011), Hornstein and Kudlyak (2016), and Sedlacek (2016), have sought to infer the variation in within-group search intensities from changes in relative job-finding rates. Among the studies adopting this general approach, Hall and Schulhofer-Wohl (2018) offer the most comprehensive characterization of the job searcher pool, considering sixteen groups of job seekers—thirteen groups among the unemployed and two among those out of the labor force plus the employed.

The basic strategy in all of these studies is, in effect, to infer what is happening to group-specific search intensities based on how having more or fewer people in any given group affects the number of matches. If adding people to a group makes a larger than expected contribution to the number of matches realized when the labor market is tight, for example, procyclicality in search intensity is a plausible explanation.[1] A limitation

1. Alternatively, the cross-group differences that are the basis for the suggested inference about job search intensity could be attributable to differences in the pattern of the shocks experienced by different groups of searchers. This is a less parsimonious explanation, and it is not entirely apparent what the source of such shocks might be, though it cannot be ruled out.

is that cyclical variation in search intensity that is common across groups cannot be distinguished empirically from the elasticity of matching with respect to the (properly measured) ratio of vacancies to searchers in the standard matching function or common changes in matching efficiency. A modeler can hope to quantify changes in aggregate search intensity that result from changes in the relative sizes of groups with relatively procyclical or relatively countercyclical job-finding rates. The potentially more important changes in search intensity that are common across groups, however, cannot be quantified using this approach.

I.C. Time-Varying Employer Recruiting Intensity

A final factor missing from the standard search-and-matching model is employer recruiting intensity. Empirical implementations of the standard model use data on the number of job openings. The intensity with which employers recruit to fill their openings can vary considerably, however, depending both on the company's own circumstances and on aggregate labor market conditions.

The most literal interpretation of recruiting intensity is the time and effort devoted to advertising the firm's job openings, processing applications, and so on, but other aspects of firms' recruiting behavior may be even more important. As an example, in a tighter labor market, employers may choose to consider job candidates with criminal records (Casselman 2018; Smialek 2019), lower the levels of education and experience they require, offer better working conditions, or raise wages. All of these can be viewed as changes in recruiting intensity, in the sense that employers are trying harder to fill their vacant jobs.

Modestino, Shoag, and Balance (forthcoming) show that, controlling for occupation, the shares of online job advertisements stating a requirement for a college degree or four or more years of experience rose during the Great Recession. The changes were larger in states and occupations that experienced a larger increase in the supply of workers. Using establishment-level Job Openings and Labor Turnover Survey (JOLTS) data, Davis, Faberman, and Haltiwanger (2013) show that, holding aggregate conditions constant, employers with a larger gross hiring rate fill their jobs at a faster pace. They interpret this finding through the lens of recruiting intensity—that is, they infer that recruiting intensity is positively associated with the gross hiring rate. Their findings suggest that the decline in hiring rates during the Great Recession should have led to a corresponding decline in recruiting intensity. Cross-sectional heterogeneity in recruiting intensity also may imply changes over time in overall recruiting intensity attributable

to composition effects (Davis, Faberman, and Haltiwanger 2012b). Building on the results of Davis, Faberman, and Haltiwanger (2013), Gavazza, Mongey, and Violante (2018) develop an equilibrium model of the recruiting intensity response to negative shocks. Their model incorporates both the gross hiring rate channel and a second channel in which, in a weaker labor market, firms of all types exert less effort to fill their jobs. Based on a calibration exercise, they conclude that the latter effect is more important than the former. Mongey and Violante (2020) argue similarly that the effect of labor market slackness on firms' chosen recruiting intensity is a key driver of the residual variation in match rates during the Great Recession.

II. An Organizing Framework

In the standard specification of the canonical search-and-matching model (Diamond 1982; Blanchard and Diamond 1989, 1992; Mortensen and Pissarides 1994; Pissarides 2000), employers create job openings they would like to fill (V), and unemployed individuals (U) search among these job openings for employment. The process of matching unemployed workers to vacant jobs is represented by a production function, often assumed to be Cobb-Douglas in form, with vacancies and unemployment as the inputs and matches (hires) as the output:

$$H_t = m(V_t, U_t) = \mu_t V_t^{1-\alpha} U_t^{\alpha}, \tag{1}$$

where H is hires, V is the number of job openings, U is the number of unemployed people, t is the time period, μ_t is a time-varying match efficiency parameter, and $\alpha(1 - \alpha)$ is the elasticity of the matching function with respect to unemployment (vacancies). In this framework, labor market tightness (θ_t) is expressed as:

$$\theta_t = \frac{V_t}{U_t}. \tag{2}$$

This relationship may be viewed through the lens of the job-finding rate, expressed as hires per unemployed worker:

$$\frac{H_t}{U_t} = \mu_t \left(\frac{V_t}{U_t}\right)^{1-\alpha} = \mu_t (\theta_t)^{1-\alpha}. \tag{3}$$

All else the same, when the labor market is tighter (when θ_t is larger), an unemployed individual is more likely to find a job. An alternative but equivalent approach is to view this relationship through the lens of the job-filling rate, expressed as hires per vacant job:

$$\frac{H_t}{V_t} = \mu_t \left(\frac{U_t}{V_t} \right)^{\alpha} = \mu_t \left(\frac{1}{\theta_t} \right)^{\alpha}. \tag{4}$$

All else the same, when the labor market is tighter, an employer is less likely to be able to recruit an unemployed person to fill a vacant job.

Since the matching function as written in equation (1) has constant returns to scale, it can be expressed as a relationship among the hiring rate h, vacancy rate v, and unemployment rate u:

$$h_t = m(v_t, u_t) = \mu_t v_t^{1-\alpha} u_t^{\alpha}, \tag{5}$$

where $h = H/\text{E}$, $v = V/\text{E}$, and $u = U/\text{E}$ and E is employment. An additional constraint is that, in steady state, the number of separations (inflows to unemployment) must equal the number of hires (outflows from unemployment). This steady state relationship can be expressed:

$$\delta_t = h_t = m(v_t, u_t) = \mu_t v_t^{1-\alpha} u_t^{\alpha}, \tag{6}$$

where δ is the separation rate (in this case separations from employment into unemployment expressed as a fraction of employment) and the other terms are as previously defined.[2] The downward sloping relationship between the unemployment rate and the vacancy rate implied by equation (6) commonly is termed the Beveridge curve.

In this framework, shifts in either δ or μ will shift the position of the Beveridge curve. Improvement in the matching function (an increase in μ), for example, shifts the Beveridge curve inward (lowering unemployment for given vacancies), while deterioration in the matching function shifts the Beveridge curve outward (raising unemployment for given vacancies). Shifts in μ also will affect the job-filling rate and the job-finding rate.

2. Nothing fundamental changes if this expression is modified to allow for steady state growth at rate g in desired employment, in which case the left-hand side becomes $\delta + g$. Although the standard Beveridge curve specification expresses vacancies and unemployment relative to employment, as in equation (6), for comparison with the generalized Beveridge curve examined later, our empirical implementation works with vacancies and unemployment relative to population.

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

In the model as just sketched out, $\theta_t = \frac{V_t}{U_t}$, but unemployment and vacancies are imperfect proxies for the measures of effective searchers and effective vacancies that we believe should be the objects of interest. Empirically, the standard framework requires substantial variation in μ_t to account for observed shifts in the Beveridge curve, as well as to explain the considerable residual variation between actual job-filling and job-finding rates and those implied by the model. One reason for this may be that, by ignoring heterogeneity among the unemployed and job search among those who are out of the labor force or employed, the standard specification misrepresents the stock of potential job candidates. If effective job seekers of each of the different types moved together with unemployment over time, it would not be important to account for them separately (Broersma and van Ours 1999; Sedlacek 2016), but this cannot be the case as increases in any one group imply decreases in others. In addition, the standard tightness measure does not account either for temporal variation in search intensity or for temporal variation in employer recruiting intensity.

We can elaborate the simple model to account for these complexities. Building on the standard hiring function, we can write:

$$H_t = m\left(\sum_j \rho_t^{v_j} V_{jt}, \sum_i \rho_t^{s_i} S_{it}\right) = \mu_t \left(\sum_j \rho_t^{v_j} V_{jt}\right)^{1-\alpha} \left(\sum_i \rho_t^{s_i} S_{it}\right)^{\alpha}, \tag{7}$$

where V_j represents the number of job openings from firms of type j, S_i represents the number of job searchers of type i, $\rho_t^{v_j}$ represents the intensity of employer recruiting effort for firms of type j at time t, and $\rho_t^{s_i}$ represents the intensity of job search on the part of searchers of type i at time t.[3] Note that the elasticity of the matching function specified in equation (7) as well as the unexplained residual variation in the actual number of matches as compared to the number implied by the model are likely to differ from those associated with the standard model specified in equation (1).

In this expanded framework, labor market tightness can be written as:

$$\tilde{\theta}_t = \frac{\sum_j \rho_t^{v_j} V_{jt}}{\sum_i \rho_t^{s_i} S_{it}}. \tag{8}$$

3. This formulation builds on specifications of generalized matching functions in Davis (2011), Davis, Faberman, and Haltiwanger (2013), and Abraham (2015).

We refer to the numerator of this expression as effective vacancies (EV_t) and the denominator as effective searchers (ES_t). The steady state equilibrium of hires equal to separations is now:

$$(9) \qquad \delta_t = h_t = m\left(\sum_j \rho_t^{v_j} v_{jt}, \sum_i \rho_t^{s_i} s_{it}\right) = \mu_t \left(\sum_j \rho_t^{v_j} v_{jt}\right)^{1-\alpha} \left(\sum_i \rho_t^{s_i} s_{it}\right)^{\alpha},$$

where separations are now all separations from employment and variables are rates expressed as fractions of the population. Over the course of a business cycle, absent changes in matching efficiency or other factors that shift the position of the generalized Beveridge curve, effective searchers and effective vacancies will move inversely as implied by equation (9).[4]

In this generalized setting, the job-filling rate is given by:

$$(10) \qquad \frac{H_t}{V_t} = \mu_t \left(\frac{1}{\tilde{\theta}_t}\right)^{\alpha} \rho_t^v, \text{ where } \rho_t^v = \frac{\sum_j \rho_t^{v_j} V_{jt}}{V_t}.$$

For the generalized model, the ratio of hires to unemployment is no longer a job-finding rate since not all hires come from among the unemployed. To characterize the job-finding rate for any subgroup or collection of subgroups, we build on the transformation of the matching function used by Hall and Schulhofer-Wohl (2018). Following Hall and Schulhofer-Wohl (2018), job-finding rates differ across groups only because of differences in their effective search intensity, meaning that they are the same on a search intensity–adjusted basis. The common search intensity–adjusted job-finding rate varies with the tightness of the labor market:

$$(11) \qquad \frac{H_t}{ES_t} = \frac{H_{it}}{ES_{it}} = \frac{H_{it}}{\rho_t^{s_i} S_{it}} = \tilde{f}_t = \tilde{f}_{it} = A_t T_t^{\eta},$$

where $T_t = V_t/H_t$ is average vacancy duration and A_t are any common time effects on job-finding rates not captured by vacancy duration. Equation (11)

4. Something we have not considered explicitly is the possibility of mismatch between vacant jobs and effective job seekers. Although commonly cited by business leaders and policy officials as an important contributor to unemployment, especially during periods when the labor market is weak (Abraham 2015), available evidence suggests that mismatch plays at most a modest role in explaining aggregate unemployment fluctuations (Şahin and others 2014; Crump and others 2019). In our framework, we will think of mismatch as captured by μ.

is a transformation of (7). To see this, define $A_t = \mu_t^{1+\eta} (\rho_t^v)^{\eta}$. Then with appropriate substitution we have:

$$H_t = \mu_t \left(\sum\nolimits_j \rho_t^{v_j} V_{jt} \right)^{\eta/(1+\eta)} \left(\sum\nolimits_i \rho_t^{S_i} S_{it} \right)^{1/(1+\eta)}, \tag{12}$$

where $\alpha = 1/(1 + \eta)$. Returning to the job-finding rate for group i, we can write:

$$\frac{H_{it}}{S_{it}} = f_{it} = \rho_t^{S_i} A_t T_t^{\eta}, \tag{13}$$

Equation (13) can be used to quantify the implied job-finding rate for the unemployed. Using the subgroups among the unemployed defined in the general model, we have:

$$\frac{H_{ut}}{U_t} = \frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta} (\rho_t^v)^{\eta} (V_t/H_t)^{\eta} \sum\nolimits_{i \in u} (\rho_t^{S_i} S_{it}/S_{ut}). \tag{14}$$

The standard model with a single group and ρ_t^v equal to 1.0 is just a special case of equation (14) given by:

$$\frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta} (V_t/H_t)^{\eta}. \tag{15}$$

Equations (7)–(15) lay out the aspirational general model that measures both effective vacancies and effective searchers, allowing for cross-sectional and time series variation in search intensity for searchers and recruiting intensity for vacancies. Our primary goal is to evaluate the extent to which this generalization overcomes some of the known limitations of the standard framework. We report empirical analyses that compare and contrast calibrations of the standard and general models, emphasizing the residual unexplained variation in the matching function associated with each model manifest in the Beveridge curve, job-filling rate and job-finding rate.

In our empirical analysis, we focus first on a simpler version of equations (7)–(15) that allows for cross-sectional variation in relative search intensities across groups but not for variation in those relative search intensities over time. That is, we begin with a specification in which $\rho_t^{S_i} = \gamma_i$ for all t and $\rho_t^{v_j} = 1$ for all j and t. Our goal in this portion of the analysis is to quantify how much allowing for a broader group of effective searchers

can improve the performance of the matching function. Then we explore specifications that permit relative search and recruiting intensities to vary over time.

III. Creating a Measure of Labor Market Tightness Based on Effective Searchers

Our measures of effective searchers build on several earlier papers, including the research underlying the Federal Reserve Bank of Richmond's Non-Employment Index (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017) and most especially the work of Hall and Schulhofer-Wohl (2018). Similar to these other papers, we use CPS micro data to track flows across labor market states and from job to job. Our analysis makes use of job-finding rates for each of twenty-two groups—thirteen groups among the unemployed (as in Hall and Schulhofer-Wohl 2018); three among those who are out of the labor force but say they want a job and four among others who are out of the labor force (as in the Richmond Fed NEI); and two among the employed.[5] Table 1 shows the full list of twenty-two groups. Following Hall and Schulhofer-Wohl (2018), we adjust each group's job-finding rate to hold demographics constant at the group's 2005–2007 values. In addition, following Fujita, Moscarini, and Postel-Vinay (2019), we make an adjustment for a change in the procedures used to collect CPS data that otherwise would lead to an understatement in job-changing rates among employed individuals. Details of the construction of the demographically adjusted job-finding rates are provided in online appendix B. As already described, we interpret the cross-group variation in (demographically adjusted) job-finding rates as variation in search intensity.

Table 1 shows, for 2006 and 2010, the population shares of each of our twenty-two groups, together with each group's demographically adjusted job-finding rate. Those who were recently temporarily laid off have the highest rate and those not in the labor force (retired) the lowest rate. Involuntary part-time workers have a job-finding rate that is twice that of other employed people, though still relatively low compared to other identified groups. Not surprisingly, average job-finding rates fell between

5. The featured Richmond Fed index distinguishes only two groups among the unemployed (short-term and long-term) and does not include employed searchers, though there is a second version that allows for search among those working part-time for economic reasons. In their work, Hall and Schulhofer-Wohl (2018) distinguish only two groups among those out of the labor force (want a job and do not want a job) and treat the employed as a single group.

Table 1. Estimated Relative Job-Finding Rates

	Share	*JFR*	*Rel. JFR*	*Rel. JFR (raw)*
2006				
Unemployed: Recently left job	0.16	39.46	0.71	0.61
Unemployed: Recently permanently laid off	0.21	32.80	0.59	0.52
Unemployed: Recently temporarily laid off	0.23	55.22	1.00	1.00
Unemployed: Temporary job recently ended	0.12	38.63	0.70	0.60
Unemployed: Recently newly entered	0.11	21.26	0.38	0.33
Unemployed: Recently reentered	0.34	29.89	0.54	0.46
Unemployed: Left job months ago	0.15	27.86	0.50	0.43
Unemployed: Permanently laid off months ago	0.36	21.19	0.38	0.33
Unemployed: Temporarily laid off months ago	0.16	44.28	0.80	0.69
Unemployed: Temporary job ended months ago	0.13	26.03	0.47	0.40
Unemployed: Newly entered months ago	0.12	14.75	0.27	0.23
Unemployed: Reentered months ago	0.45	23.44	0.42	0.37
Unemployed: Long-term unemployed	0.43	17.41	0.32	0.27
Want Job: Discouraged	0.15	14.74	0.27	0.23
Want Job: Looked last 12 months	0.43	14.24	0.26	0.22
Want Job: Other	1.24	15.26	0.28	0.24
Not in Labor Force: In school	4.34	9.41	0.17	0.15
Not in Labor Force: Retired	15.51	1.56	0.03	0.02
Not in Labor Force: Disabled	4.67	1.96	0.04	0.03
Not in Labor Force: Other	7.44	8.87	0.16	0.14
Employed: Involuntary part-time	1.79	5.12	0.09	0.08
Employed: Not involuntary part-time	61.44	2.22	0.04	0.03
2010				
Unemployed: Recently left job	0.09	27.81	0.54	0.48
Unemployed: Recently permanently laid off	0.29	23.12	0.45	0.38
Unemployed: Recently temporarily laid off	0.28	51.80	1.00	1.00
Unemployed: Temporary job recently ended	0.13	32.88	0.63	0.56
Unemployed: Recently newly entered	0.12	12.65	0.24	0.22
Unemployed: Recently reentered	0.27	21.30	0.41	0.37
Unemployed: Left job months ago	0.16	19.29	0.37	0.32
Unemployed: Permanently laid off months ago	0.90	14.41	0.28	0.24
Unemployed: Temporarily laid off months ago	0.26	36.15	0.70	0.60
Unemployed: Temporary job ended months ago	0.24	20.06	0.39	0.33
Unemployed: Newly entered months ago	0.24	9.41	0.18	0.16
Unemployed: Reentered months ago	0.57	16.45	0.32	0.28
Unemployed: Long-term unemployed	2.14	10.92	0.21	0.18
Want Job: Discouraged	0.47	11.33	0.22	0.19
Want Job: Looked last 12 months	0.52	9.76	0.19	0.17
Want Job: Other	1.27	12.30	0.24	0.21
Not in Labor Force: In school	5.07	6.28	0.12	0.11
Not in Labor Force: Retired	15.56	1.41	0.03	0.02
Not in Labor Force: Disabled	5.17	1.42	0.03	0.02
Not in Labor Force: Other	7.26	6.76	0.13	0.12
Employed: Involuntary part-time	3.73	3.63	0.07	0.06
Employed: Not involuntary part-time	55.27	1.77	0.03	0.03

Source: Authors' calculations using CPS.

Notes: Job-finding rates (JFR) estimated using CPS survey data linking households month to month. Relative job-finding rate (Rel. JFR) calculated by dividing all job-finding rates by job-finding rate for unemployed recently laid off. Recently unemployed groups refer to those unemployed 0–4 weeks. "Unemployed months ago" refers to those unemployed 5–26 weeks. "Long-term unemployed" refers to those unemployed 27 weeks or more.

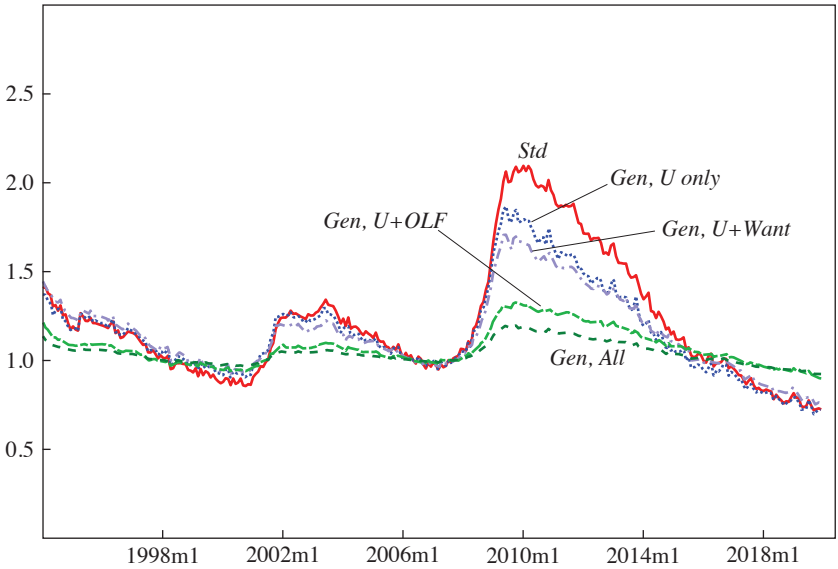
2006 and 2010 as the economy worsened following the onset of the Great Recession. Our primary interest, however, lies with the relative job-finding rates across the different groups. These are much more stable—the correlation in relative job-finding rates between 2006 and 2010 is 0.98. We construct our measure of effective searchers by weighting each of the twenty-two groups by its relative 2006 job-finding rate, shown in the third column of the top panel of table 1. All of the relative job-finding rates are defined with reference to the demographically adjusted 2006 job-finding rate of those who were recently temporarily laid off.

Table 1 also reports relative job-finding rates for 2006 and 2010 calculated directly from the CPS micro data not controlling for changing demographics, shown in the column denoted "Rel. JFR (raw)." Although our baseline analysis uses the demographically adjusted 2006 relative job-finding rate estimates, the results are very similar if we instead use the raw 2006 relative job-finding rates. Later in the paper, we present results based on demographically adjusted relative job-finding rates that are allowed to vary over time.

Figure 1 displays the standard measure of searchers (the simple unemployment count) along with several alternative measures that move toward our fully generalized effective searcher measure. The first of these considers only the unemployed as effective searchers but allows for changes in the composition of the unemployed across the thirteen groups we have specified. The second alternative measure incorporates the three groups among those out of the labor force who say they want a job, and the third adds the four groups among others who are out of the labor force. The fourth and final effective searcher measure is the fully generalized measure that adds the two groups of employed people. Each of the generalized measures weights the different groups it includes in accord with their relative 2006 job-finding rates.

The alternative series shown in figure 1 are highly correlated but distinctly different in their volatility. Allowing for heterogeneity among the unemployed yields a measure that is less cyclically volatile than the standard unemployment measure. Including, in turn, those who are out of the labor force but want a job, others who are out of the labor force, and finally the employed yields progressively less volatile measures. The first column of table 2 reports the standard deviations over the 1994–2019 period of the normalized series plotted in figure 1. Whereas unemployment as a share of the population, normalized to equal 1 in 2006, has a standard deviation of 0.34 over the 26-year period, the standard deviations of the alternative measures, also expressed as a share of the population and

Figure 1. Standard versus Generalized Measures of Searchers (Constant Relative Search Intensities)



Source: Authors' calculations using Current Population Survey (CPS).

Notes: All measures ratios to population age 16+ normalized to 1 in 2006. Std = unemployed; Gen, U only = generalized measure, unemployed only; Gen, U + Want = generalized measure, unemployed plus want a job; Gen, U + OLF = generalized measure, excludes employed; Gen, All = generalized measure, all twenty-two groups.

Table 2. Cyclical Volatility of Alternative Searcher Measures: Constant Relative Job Search Intensities

	Standard deviation, 1994–2019	*March 2001*	*June 2009*	*December 2019*
Standard	0.34	0.94	2.06	0.72
General, unemployed only	0.26	1.00	1.86	0.72
General, unemployed + want job	0.22	0.99	1.71	0.76
General, unemployed + OLF	0.10	0.97	1.32	0.90
General, all	0.06	0.99	1.19	0.92
U6	0.35	0.89	2.00	0.81
Richmond Fed NEI	0.11	0.96	1.34	0.89

Source: Authors' calculations using CPS and Richmond Fed NEI.

Notes: All measures normalized to 1 in 2006. Standard = unemployed; General, unemployed only = effective searchers with unemployed only; General, unemployed + want job = effective searchers with unemployed and want a job; General, unemployed + OLF = effective searchers excluding employed; General, all = all effective searchers; U6 = U6 measure of labor underutilization.

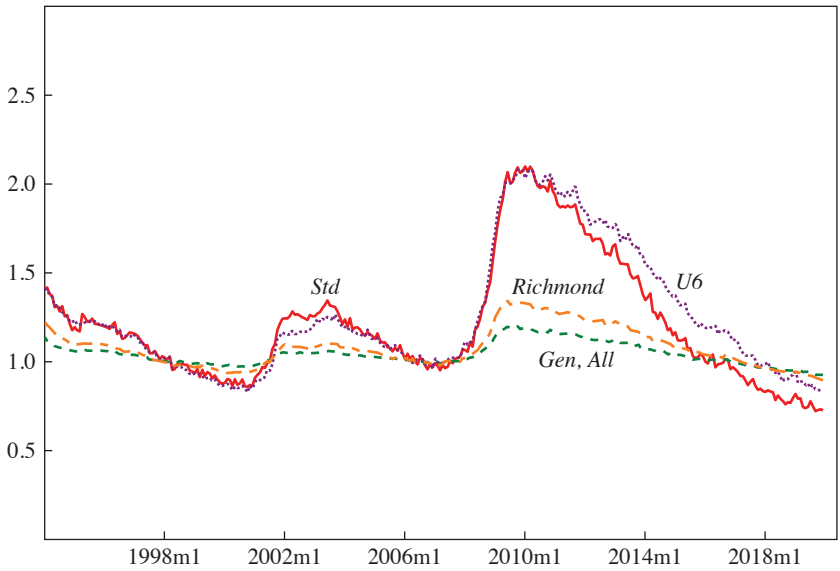
normalized to equal 1 in 2006, are progressively lower. The standard deviation of our fully generalized measure, calculated on a basis comparable to that of the unemployment measure, is just 0.06.

Another way to look at the alternative measures is to ask how the current level of effective searchers as a share of the population compares to the level at points of time in the past. The end of the prolonged expansion that lasted through most of the 1990s and into the early 2000s offers one interesting point of comparison. The standard measure of effective searchers (unemployment) relative to the population recorded in December 2019 was about 23 percent below its March 2001 value. In contrast, the December 2019 fully generalized measure of effective searchers relative to the population was only about 6 percent below its March 2001 value. Another interesting point of comparison is with the values for June 2009, the trough of the Great Recession. Whereas the standard measure of effective searchers fell by about 65 percent between June 2009 and December 2019, the fully generalized measure fell by just 23 percent.

One important reason for the muted cyclicality of the baseline generalized measure with fixed relative job search intensities as compared to the standard measure is that the generalized measure counts more people as effective searchers. The largest proportional fluctuations in group size over a typical cycle are the fluctuations among the unemployed. In the standard measure, any proportional increase in the number of unemployed people is de facto a proportional increase in the number of effective searchers. In the baseline version of the generalized measure, in contrast, the unemployed are only a fraction of all effective searchers, and increases in unemployment thus mechanically have a smaller proportional effect on the aggregate number of effective searchers. Composition also matters. When economic conditions are weak, the share of the unemployed who are long-term unemployed rises, and the long-term unemployed have lower job-finding rates than other unemployed people. In addition, in the baseline version of the generalized measure, even though the employed have lower relative search intensities than the unemployed, their search intensities are positive. This means that, during a downturn, reductions in the number of employed effective searchers partially offset increases in the number of effective searchers associated with rising unemployment.

Figure 2 compares our generalized measure of effective searchers with two alternative measures. The first is an index based on the Bureau of Labor Statistics U6 measure of slackness. The U6 measure counts the marginally attached and involuntary part-timers along with the unemployed but like

Figure 2. Standard, U6, Richmond Fed, and Generalized Effective Searcher Measures (Constant Relative Search Intensities)



Sources: Authors' calculations using CPS and Richmond Fed NEI.

Notes: All measures ratios to population age 16+ normalized to 1 in 2006. Std = unemployed; Gen, All = generalized measure, all twenty-two groups; U6 = U6 measure of labor underutilization; Richmond = Richmond Fed NEI.

the official unemployment rate weights all of them equally. The headline Richmond Fed NEI (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017) incorporates both the unemployed and those out of the labor force but distinguishes fewer groups among the unemployed than our generalized measure (two versus thirteen) and does not incorporate employed searchers. Similar to our measure, the different groups used in the Richmond Fed index are weighted based on persistent differences in their average relative job-finding rates. For consistency, we normalize both the U6 index and the Richmond Fed index by the population age sixteen and older.

As can be seen in figure 2, the U6 index has about the same volatility as the standard measure. The cyclical variation in the Richmond Fed index is proportionally less than that of the standard measure but greater than that of our generalized index of effective searchers. Table 2 reports summary statistics for both the U6 and the Richmond Fed indexes. Consistent with the visual impression conveyed by figure 2, over the 1994–2019

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

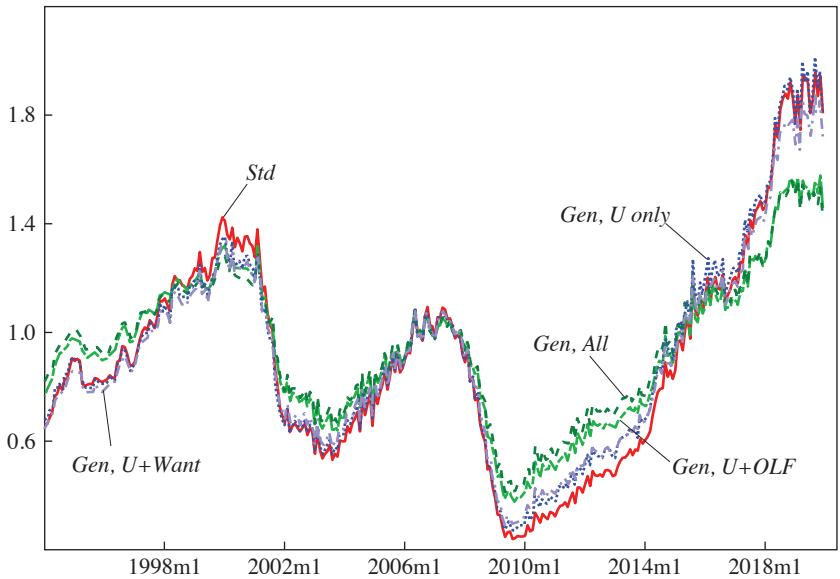
period as a whole, the volatility of the U6 index is very similar to that of the unemployment index. The volatility of the Richmond Fed index lies between that of the standard measure and our generalized measure.

As already mentioned, in constructing our generalized measures of effective searchers, we have used estimated relative job-finding rates that hold demographics constant over time. We also have constructed similar measures using the simple average 2006 job-finding rates for the twenty-two groups. Comparing these measures, shown in online appendix figure A.1, to those in figure 1 suggests that controlling for demographics is relatively unimportant. As a further sensitivity analysis, we also have asked how much difference it makes that we have broken the unemployed into thirteen different groups, as opposed to distinguishing just between the short-term and the long-term unemployed. Measures constructed using the latter approach are shown in online appendix figure A.2. In addition, we have constructed an effective searcher measure that breaks the population into just five groups—short-term unemployed, long-term unemployed, out of the labor force and want a job, out of the labor force and do not want a job, and employed. The comparison between this measure and our fully generalized measure is shown in online appendix figure A.3. Over our time period, effective searcher measures constructed using these alternative approaches look broadly similar to measures constructed using our baseline approach. We note, however, that the similar behavior of these alternative searcher measures over this period does not necessarily imply they always will behave so similarly. There are marked differences in base period job-finding rates across the groups that are pooled together in the alternative measures, and, in a different period, distinguishing among them could make more of a difference. Later in the paper, we consider the performance of each of the alternative job searcher measures in explaining job-filling and job-finding rates.

We now are ready to compare labor tightness measured using effective searchers as opposed to unemployment. The numerator for all of the tightness measures shown in figure 3 is vacancies. This is the published JOLTS series from 2001:M1 to 2019:M12 and the Davis, Faberman, and Haltiwanger (2012a) backcast series for 1994:M1–2000:M12.[6] The denominators of the alternative generalized measures incorporate successively

6. We also have produced a labor market tightness measure using the Barnichon (2010) series for the pre-JOLTS period. We find results that are quite similar whether we use the Davis, Faberman, and Haltiwanger (2012a) or the Barnichon (2010) series.

Figure 3. Standard versus Generalized Measures of Labor Market Tightness (Constant Relative Search/Intensities, Recruiting Intensity = 1)



Sources: Authors' calculations using CPS, Job Openings and Labor Turnover Survey (JOLTS) for 2000:M12–2019:M12 and backcast vacancies from Davis, Faberman, and Haltiwanger (2012a) for 1994:M1–2000:M11 (DFH (2012a) vacancies hereafter).

Notes: All measures normalized to 1 in 2006. Std = V/U; Gen, U only = V/ES (unemployed only); Gen, U + Want = V/ES (unemployed plus want a job); Gen, U + OLF = V/ES (excludes employed); Gen, All = V/ES (all twenty- two groups).

more encompassing pools of effective searchers, in each case with the subgroups we have defined weighted in accord with their relative search intensities. The first generalized labor market tightness measure shown in figure 3 considers only the unemployed as effective searchers, while allowing for changes in unemployment composition. The second incorporates people who are out of the labor force but say they want a job, and the third adds the remainder of those out of the labor force. The final, fully generalized tightness measure also treats the employed as effective searchers. Once again, for ease of comparison, all of the measures in figure 3 have been normalized to equal 1 on average in 2006.

The more inclusive generalized tightness measures displayed in figure 3 are markedly less cyclical than the standard tightness measure—in particular, they fell much less steeply during the Great Recession and subsequently have risen less. The December 2019 values of the generalized measures

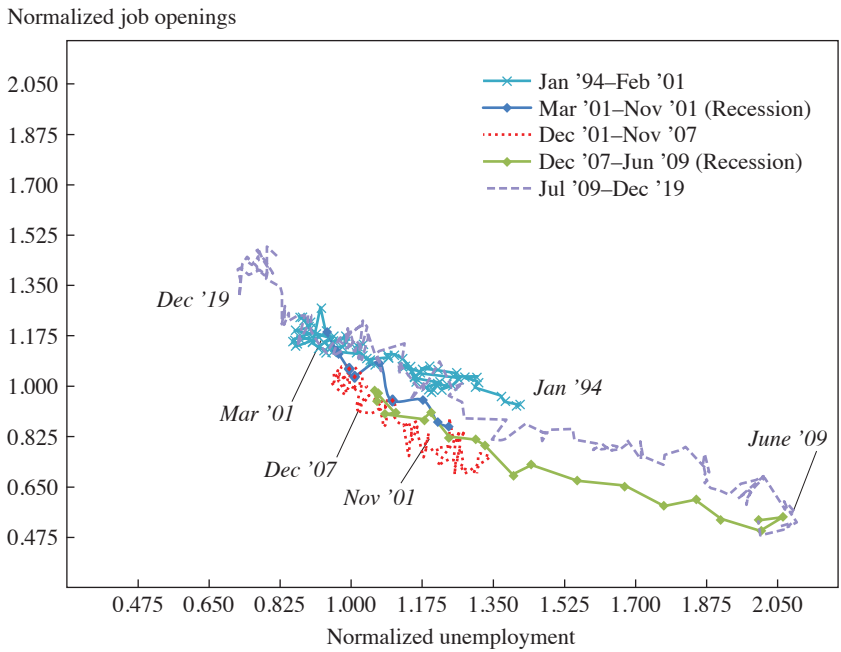
incorporating only unemployed searchers or adding just the want a job group are not very different from the value of the standard measure, but the other measures are substantially lower, implying that the labor market was not as tight at that point as implied by the standard measure. To put this into context, the December 2019 value of the standard labor market tightness measure is about 42 percent higher than in March 2001. In contrast, the generalized measure using effective searchers is only about 17 percent higher. In short, our generalized measure using effective searchers in place of unemployment suggests a significantly different evolution of labor market tightness than the standard measure.

For comparison, online appendix figure A.4 plots labor market tightness measures constructed using the U6 index and the Richmond Fed index along with the standard and our generalized measure. The value of the tightness measure based on the U6 index is about 27 percent higher in December 2019 than it was in March 2001, and the value of the tightness measure based on the Richmond Fed index is about 18 percent higher.

IV. Beveridge Curve

A closely related but distinct way to look at the properties of the effective versus standard measures of searchers and job openings is through the lens of the Beveridge curve. Figure 4 displays the standard Beveridge curve using monthly data on vacancies and unemployment from 1994:M1 to 2019:M12. For this purpose, we use normalized unemployment and vacancy series defined relative to their 2006 average values. Plotting these series against one another makes clear their inverse relationship. In addition to the familiar downward sloping relationship between vacancies and unemployment, the figure also shows the substantial outward shift in that relationship during the long, slow recovery from the Great Recession. To illustrate, consider the period two to four years after the trough of the Great Recession (from June 2011 through May 2013), as compared to the comparable period following the trough of the 2001 recession (from November 2003 through October 2005). Job openings are only slightly lower over this portion of the recovery from the Great Recession than during the corresponding period following the 2001 recession (by about 8 percent) but unemployment is much higher (by about 48 percent). The marked increase in unemployment compared to that associated in the past with a similar level of vacancies led many to speculate that, following the Great Recession, there had been a decline in matching efficiency in the labor market.

Figure 4. Standard Beveridge Curve

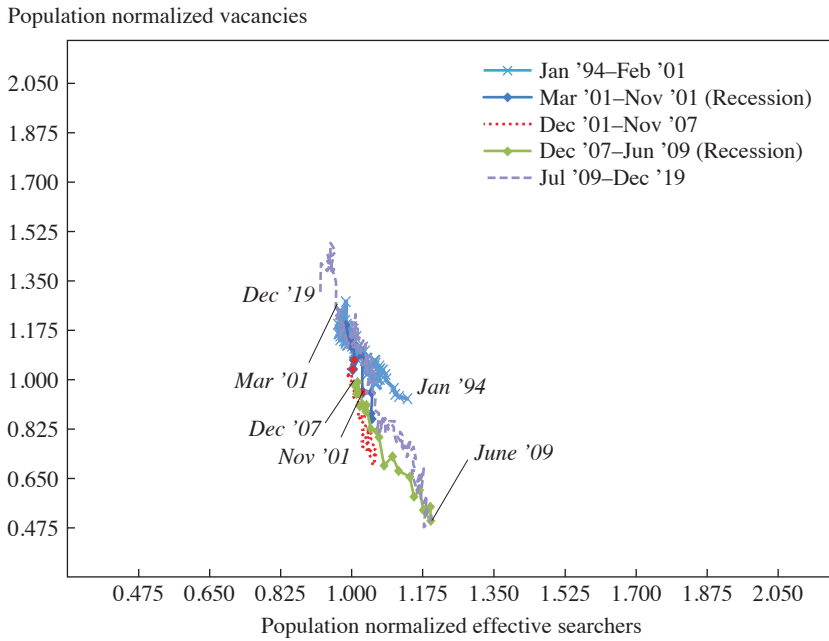


Source: Authors' calculations using job openings (vacancies) from JOLTS for 2000:M12–2019:M12 and DFH (2012a) vacancies for 1994:M1–2000:M11; unemployed from CPS.

Note: Both series as rates relative to population age 16+ normalized to 1 in 2006.

Figure 5 depicts the generalized Beveridge curve using effective searchers (using our baseline constant relative search intensities measure) with series normalized so the values plotted are all relative to their 2006 averages. The generalized Beveridge curve shown in figure 5 is much steeper than the standard version shown in figure 4. As noted previously when discussing figures 1 and 2, the proportional variation in effective searchers over time is much smaller than the proportional variation in unemployment. This translates into a normalized Beveridge curve that spans a much shorter distance along the horizontal axis than does the standard Beveridge curve and also is much more stable than the standard curve during the period following the Great Recession. Consider again the period two to four years after the trough of the Great Recession as compared to the period following the 2001 recession. As before, vacancies are slightly lower from two to four years after the Great Recession (by about

Figure 5. Beveridge Curve using Vacancies and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)



Source: Authors' calculations using job openings (vacancies) from JOLTS for 2000:M12–2019:M12 and DFH (2012a) vacancies for 1994:M1–2000:M11; effective searchers from CPS.

Notes: Both series as rates relative to population age 16+ normalized to 1 in 2006. Effective searchers as described in text.

8 percent), but whereas unemployment was *much* higher over the same period (about 48 percent), effective searchers were only slightly higher (about 10 percent).

One way to summarize the general shape of the Beveridge curve using different effective searcher measures is to fit descriptive regressions of vacancies on the various measures. Online appendix table B.3 shows that, consistent with figures 4 and 5, the standard Beveridge curve has a slope that is well below 1.0 in absolute value, whereas the Beveridge curve based on our fully generalized effective searcher measure is much steeper. The table also shows how moving in steps from the standard measure to our fully generalized measure leads to a Beveridge curve that is increasingly steep, reflecting the progressively lower volatility of the more encompassing measures. The slope of the Beveridge curve using the U6 measure is

similar to that for the standard unemployment measure; the slope using the Richmond Fed index lies between that for the standard and the fully generalized measures.

V. Implications of Redefining Effective Searchers for the Matching Function

We have argued that the measure of labor market tightness using our generalized effective searcher series should be preferred conceptually to the standard measure, but we would like to have evidence that it actually does a better job of explaining the temporal variation in job-filling and job-finding rates. To evaluate the performance of the alternative measures in the matching function, we return to equation (1), the standard specification, and equation (7), the generalized specification, and ask how well each performs in tracking actual job-filling and job-finding rates. The targets we seek to match are the job-filling rate based on JOLTS data and the job-finding rate among the unemployed based on CPS gross flows data. We begin with our baseline model that defines $\rho_t^{V_i} = 1$ and $\rho_t^{S_i} = \gamma_i$. The next section of the paper will consider time-varying recruiting and relative search intensities.

V.A. Actual versus Model-Based Patterns in the Job-Filling Rate

Both the standard and the generalized matching function have implications for the evolution of the job-filling rate (H/V) as illustrated in equations (4) and (10). Because the left-hand sides of equations (4) and (10) are the same and are based on readily available data, we can compare the residual variation in the calibrated job-filling rates obtained using the standard and the generalized matching functions. For the present case, the specification of equation (10) is given by:

$$\frac{H_t}{V_t} = \mu_t \left(\frac{\sum_i \gamma_i S_{it}}{V_t} \right)^{\alpha}. \tag{10′}$$

In addition to our generalized measure with twenty-two different groups of effective searchers, we also consider versions of the generalized labor market tightness measure based on the U6 index and the Richmond Fed index.

In addition to the vacancy and unemployment (effective searcher) measures appearing in equations (4) and (10′), the calibrated job-filling rate associated with each of these tightness measures also depends on the

elasticity of the matching function. To give each set of matching function arguments the best possible chance to fit the data well, we have estimated a separate α for each using a simplified version of the method proposed by Hall and Schulhofer-Wohl (2018). This method relates the job-finding rates for specific groups to vacancy duration from the JOLTS. As described more fully in online appendix B, this yields an estimate of the elasticity η of the job-finding rate with respect to vacancy duration for each measure of effective searchers. Given this estimate, we compute $\alpha = 1/(1 + \eta)$. For the standard model, this gives us an estimate for η of 1.04 (standard error 0.05) implying a value for α of 0.49. For the model with our effective searcher specification, the estimate is $\eta = 0.75$ (0.04) implying $\alpha = 0.57$. For the U6 measure, we obtain $\alpha = 0.48$ ($\eta = 1.1$ [0.05]) and for the Richmond Fed measure we obtain $\alpha = 0.60$ ($\eta = 0.67$ [0.05]). All of these estimates are reasonably similar and well within the middle of the range of estimates in the matching function literature (Petrongolo and Pissarides 2001). Our main results are broadly unchanged if we apply a common value of the matching function elasticity within the range of the separate estimates to calibrate the job-filling rates using the different tightness measures.[7]

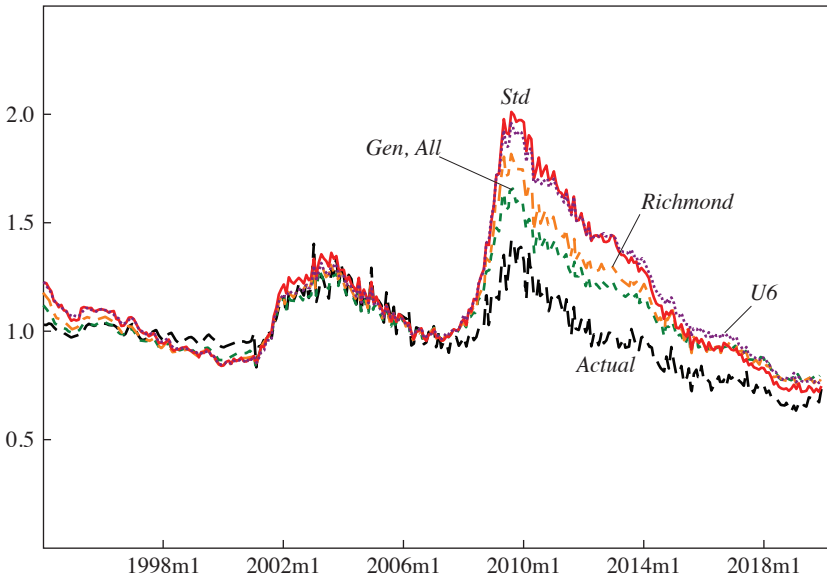
Figure 6 presents the actual and calibrated job-filling rates from equations (4) and (10′) using the standard, generalized, U6, and Richmond Fed measures. Once again, all series have been normalized to average 1 in 2006.[8] The calibrated job-filling rate based on equation (10′) and our generalized tightness measure tracks the actual job-filling rate much more closely than the calibrated rate based on equation (4) and the standard tightness measure. The U6 measure performs no better than the standard measure. The Richmond Fed measure performs substantially better than the standard measure, but not as well as our generalized measure.

To quantify the improvements in performance, panel A of table 3 reports the root-mean-square error (RMSE) of the calibrated job-filling rates as compared to the actual rates based on the different tightness measures. The generalized measure produces an RMSE of 0.13, a little more than half

7. In their evaluations of job-filling and job-finding rates using alternative measures of searchers and vacancies, Davis (2011), Davis, Faberman, and Haltiwanger (2013), and Mongey and Violante (2020) use $\alpha = 0.5$.

8. The model-specific normalization for the calibrated job-filling rates implies that we are permitting the mean matching efficiency for each model to be different. More specifically, it is permitted to vary in such a way that the calibrated job-filling rate in all models is equal to the actual job-filling rate on average in 2006.

Figure 6. Actual versus Calibrated Job-Filling Rates with Standard and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)



Source: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and Richmond Fed NEI.
Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups); U6 = calibrated using V/U6 index; Richmond = calibrated using V/Richmond Fed NEI.

as large as the RMSE of 0.25 produced using the standard measure. In other words, using the generalized measure with effective searchers in the calibration substantially reduces the unexplained residual variation in the job-filling rate.[9]

For comparison purposes, we also show the RMSEs in the calibrated job-filling rates for the U6 and Richmond Fed indexes. The RMSE for the U6 index is identical to that for the standard measure; the RMSE for the Richmond Fed index is intermediate between those for the standard and the generalized model. Table 3 also reports summary statistics for the alternative calibrated job-filling rates compared to the actual. The calibrated

9. Note that the calibration exercise we carry out is quite different in nature from regressing the actual job-filling rate on the calibrated values. See online appendix B for further discussion.

Table 3. Relative Performance for Job-Filling and Job-Finding Rates using Standard versus Effective Searchers: Constant Relative Search Intensities, Recruiting Intensity = 1

	Std dev	*Corr w/ actual*	*RMSE*	*Ratio RMSE to standard*
A. Job-filling rate				
Actual	0.17	1.00	NA	NA
Standard	0.29	0.75	0.25	1.00
General, all	0.19	0.81	0.13	0.54
U6	0.27	0.70	0.25	1.00
Richmond Fed NEI	0.23	0.78	0.18	0.71
B. Job-filling rate using $\alpha = 0.57$ for all measures				
Standard	0.35	0.74	0.31	1.00
Gen, U only	0.31	0.79	0.25	0.81
Gen, U + want	0.28	0.78	0.23	0.75
Gen, U + OLF	0.21	0.78	0.16	0.53
Gen, all	0.19	0.81	0.13	0.43
C. Job-finding rate for the unemployed				
Actual	0.16	1.00	NA	NA
Standard ($\alpha = 0.49$)	0.20	0.31	0.26	1.00
Standard ($\alpha = 0.57$)	0.14	0.32	0.22	0.85
General, all	0.15	0.62	0.17	0.67

Source: Authors' calculations using CPS; JOLTS; Davis, Faberman, and Haltiwanger (2012a) vacancies; Richmond Fed NEI; and Bureau of Labor Statistics Gross Flows.

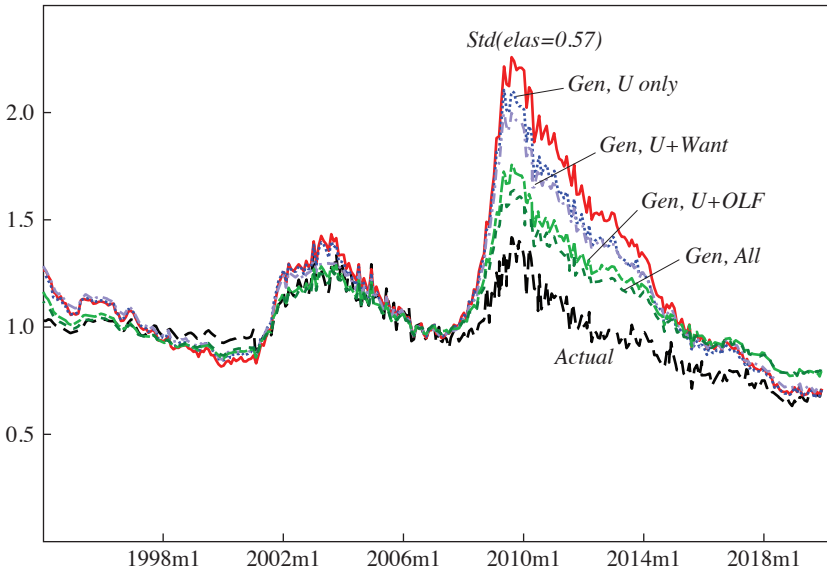
Notes: Panel A shows statistics for calibrated job-filling rates using matching function elasticities specific to each. Panel B shows statistics using same matching function elasticity for all measures. Panel C shows statistics for calibrated job-finding rate for the unemployed.

job-filling rate using effective searchers has a higher correlation and a standard deviation closer to the actual job-filling rate than the alternative calibrated series.

To help with understanding the factors underlying the improvement in performance of the generalized versus the standard tightness measure in tracking the job-filling rate, figure 7 presents calibrated rates based on a series of measures that incorporate the differences between the two in stages. For these figures and the associated analysis, we use the same matching function elasticity of $\alpha = 0.57$ for all of the counterfactual measures and for the standard measure. Panel B of table 3 reports the RMSEs in the calibrated job-filling rate using each of the different versions of the generalized tightness measure shown in figure 7. On its own, allowing for heterogeneity among the unemployed reduces the RMSE of the residual unexplained variation by about 20 percent. We gain an additional 5 percent by including in the pool of effective searchers those out of the labor force who want a job, an additional 20 percent by including the remaining people who are out of the labor force, and an additional 10 percent by including

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

Figure 7. Actual versus Calibrated Job-Filling Rates with Standard and Effective Searchers, Using $\alpha = 0.57$ for All Measures (Constant Relative Search Intensities, Recruiting Intensity = 1)



Source: Authors' calculations using CPS, JOLTS, and DFH (2012a) vacancies.

Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, U only = calibrated using V/ES (unemployed only); Gen, U + Want = calibrated using V/ES (unemployed plus want a job); Gen, U + OLF = calibrated using V/ES (excludes employed); Gen, All = calibrated V/ES (all twenty-two groups).

the employed. Appropriate caution should be used in interpreting these figures, as the numbers we have reported are not an exact decomposition, but they do provide useful guidance with respect to which features of our generalized measure account for its better performance.

We also have explored a number of sensitivity checks that we summarize briefly here; details are shown in online appendix A. First, as shown in online appendix figure A.5, we replicate the analysis of job-filling rates using simple averages of the direct (raw) relative job-finding rates from table 1 rather than the relative job-finding rates that abstract from demographics. We also replicate the findings with a generalized measure that breaks the unemployed into just two groups, the short-term and the long-term unemployed, rather than thirteen more disaggregated groups (online appendix figure A.6) and with another version that disaggregates the population into only five groups (short-term unemployed, long-term unemployed, out of the labor force and want a job, out of the labor force

and do not want a job, and employed; online appendix figure A.7). In both cases, the results are broadly similar to those we have just reported. The unexplained residual variation in the job-filling rate is the same using the raw job-finding rates and slightly higher when either limiting the disaggregation of the unemployed to just two groups or using the five-group disaggregation described above (RMSE = 0.14 for both as compared to RMSE = 0.13 for the fully disaggregated specification).

In addition, we have replicated the job-filling rate analysis using the Barnichon (2010) vacancy estimates based on help-wanted advertising for the 1994:M1–2000:M12 period in place of the series based on the methodology described by Davis, Faberman, and Haltiwanger (2012a). Again, the results are broadly similar (see online appendix figure A.8), though over the 1994:M1–2000:M12 period for which we must use projected vacancies, the generalized measure using Davis, Faberman, and Haltiwanger's (2012a) methodology performs better in tracking the job-filling rate than the Barnichon (2010) series. Over that period, the RMSE using the Davis, Faberman, and Haltiwanger series is 0.04 while the RMSE using the Barnichon series is 0.09.

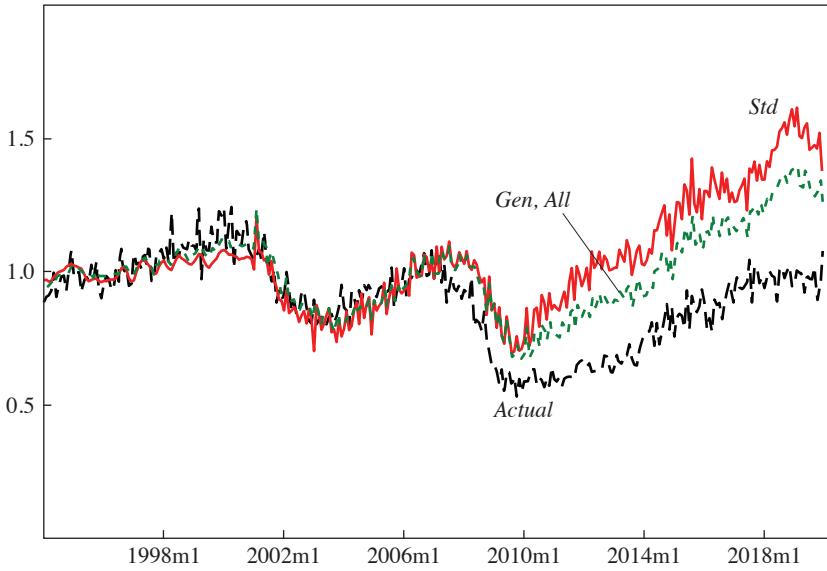
V.B. Actual versus Model-Based Patterns in the Job-Finding Rate

We now turn to investigating the performance of the generalized versus standard matching function for tracking job-finding rates. As noted in section II, overall hires per unemployed person is not a meaningful outcome measure for the generalized model with effective searchers. Instead, we use equations (14) and (15) to calibrate the job-finding rate of the unemployed. In the case we are currently considering, the specification of equation (14) is given by:

$$(14') \qquad \frac{H_{ut}}{U_t} = \frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta} \left(V_t / H_t\right)^{\eta} \left(\sum\nolimits_{i \in u} \gamma_i S_{it} / S_{ut}\right).$$

Like the standard model of equation (15), the right-hand side of equation (14′) includes vacancy duration—vacancies relative to hires—but with the difference that this is now hires from all sources, not just hires from unemployment. This reflects the fact that, in the generalized model, the job-finding rate per effective searcher is assumed to be the same across all effective searcher groups. The generalized model also has an extra term that reflects the ratio of effective searchers among the unemployed to the number unemployed, with a higher ratio implying a larger number of hires per unemployed person.

Figure 8. Actual versus Calibrated Job-Finding Rates for the Unemployed with Standard and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)



Sources: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and BLS Gross Flows.
Notes: All rates normalized to 1 in 2006. Actual = job-finding rate for unemployed from published BLS Gross Flows; Std = calibrated using V/U with $\alpha = 0.49$ ($\eta = 1.04$); Gen, All = calibrated using V/ES (all twenty-two groups).

Figure 8 shows the actual and calibrated job-finding rates for the unemployed for the generalized and standard matching functions based on using equations (14′) and (15), respectively. As with our analysis of job-filling rates, we normalize both the actual and the calibrated series to be equal to 1 in 2006.[10] The actual job-finding rate is highly procyclical and falls especially sharply in the Great Recession. Both the generalized and the standard matching function track the job-finding rate among the unemployed reasonably well during the period prior to the Great Recession, though the generalized matching function performs somewhat better over that period. The generalized and the standard matching function track the sharp decline in the job-finding rate among the unemployed during the Great Recession about equally well, but the standard model implies a faster

10. As before, this implies that we are permitting the mean matching efficiency to differ across the standard and generalized models in such a way that both yield values in 2006 with mean equal to 1.

recovery whereas the calibrated generalized model lies closer to what actually happened.

Panel C of table 3 quantifies the improvement in performance in the calibrated job-finding rate for the unemployed from using the generalized rather than the standard matching function. The RMSE of the generalized calibration of the job-finding rate from unemployment is about a third lower than the RMSE for the standard calibration. As with the calibrated job-filling rate, the calibrated job-finding rate using effective searchers has a higher correlation and a standard deviation closer to the actual job-finding rate than the standard calibrated series.

Summary statistics reported in table 3 also shed light on the difference that the choice of α can make in the results obtained. Using $\alpha = 0.57$ (the elasticity for the generalized matching function) reduces the RMSE of the standard calibration by about 15 percent compared to its RMSE when using $\alpha = 0.49$. On the other hand, comparing the results in panels A and B of table 3, increasing α from 0.49 to 0.57 worsens the relative performance of the standard model for calibrating the job-filling rate. The implication of these conflicting effects is that the relatively poor performance of the standard model cannot be rescued with an alternative estimate of the matching function elasticity.

As already noted in comparing figure 1 with online appendix figures A.1, A.2, and A.3, the time series behavior of effective searchers is quite similar whether we use raw or demographically adjusted base period job-finding rates to weight the different groups of searchers; two groups or thirteen groups among the unemployed; or five broader groups of searchers rather than twenty-two groups as in our baseline analysis. In online appendix figures A.9, A.10, and A.11, we show results for the job-finding rate from unemployment for all three of these variants. The RMSE for the measure based on raw base period job-finding rates is nearly identical to that for our baseline specification. Both the RMSE for the generalized model with only two unemployment groups and the RMSE when we use just five groups rather than twenty-two groups are about 70 percent of that for the standard model, as compared to 67 percent for the fully general model.

VI. Exploring Time Variation in Recruiting Intensity and Search Intensity

The generalized model we outlined in section II includes both cross-sectional and time series variation in both recruiting intensities and search intensities, but thus far we have considered only the cross-sectional variation.

Research on alternative approaches to measuring the variation in recruiting intensity and relative search intensities is still in its early stages. In this section, we present results from an exploratory analysis of the variation in these intensities over our 1994:M1–2019:M12 sample period. For recruiting intensity, we use the Davis, Faberman, and Haltiwanger (2013) methodology to construct an aggregate recruiting intensity index. Using cross-sectional micro data, Davis, Faberman, and Haltiwanger (2013) estimate the elasticity between the job-filling rate and the gross hiring rate at the establishment level as $\phi = 0.82$. Applying this micro-based elasticity, they proxy aggregate recruiting intensity as $\rho_t^v = \left(\frac{H_t}{E_t}\right)^{\phi}$. We apply that proxy here.[11] It also is possible that, in addition to depending on the gross hiring rate, recruiting intensity may change over time due to composition effects (Davis, Faberman, and Haltiwanger 2012b) or to the endogenous responses of firms to overall labor market conditions, independent of their gross hiring rate (Gavazza, Mongey, and Violante 2018; Mongey and Violante 2020). We leave these possibilities for future exploration.

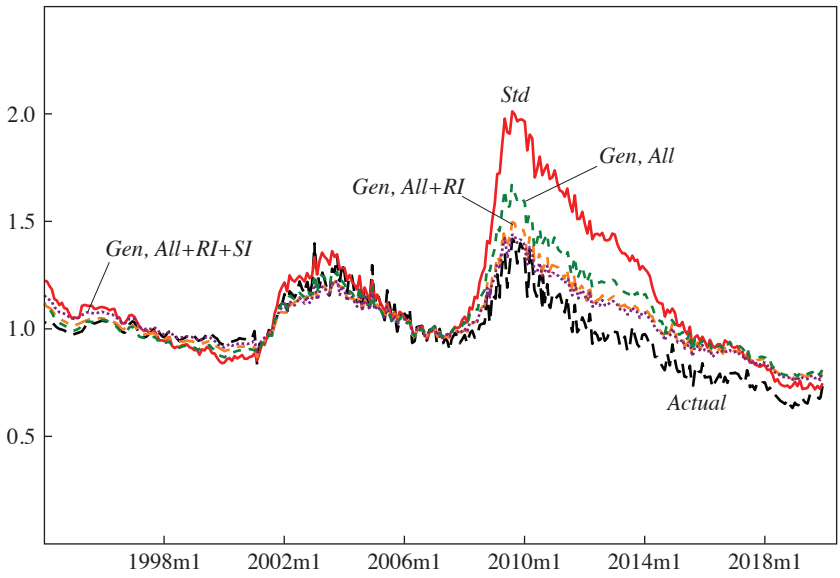
For time variation in relative job search intensities, we return to the specifications of Hall and Schulhofer-Wohl (2018). Their estimation allows not only for differences in average search intensities (what we have denoted as γ_i) but also for differences across job searcher types in the elasticity of job finding with respect to vacancy duration. We build on this insight to construct time-varying relative search intensities for our twenty-two groups, exploiting group-specific differences in the time variation in relative job-finding rates to infer what is happening to relative search intensities.[12] Details are reported in online appendix B. We do not use the common variation in job-finding rates in reweighting the different job searcher groups as we cannot distinguish changes in job-finding rates inherent in the matching function from those due to changes in search intensity.

Because the approach we have taken to constructing our time-varying recruiting intensity and relative job search intensity measures makes use of information related to the number of hires over time, some might be concerned that we have somehow built in the better performance of our model for explaining the job-filling and job-finding rates. Given the nature of our exercise, however, we do not believe this to be the case. We are

11. Online appendix figure A.12 shows how this recruiting intensity measure has moved over time.

12. See online appendix figure A.13 for the effective searcher measure using time-varying relative search intensities.

Figure 9. Actual versus Calibrated Job-Filling Rates with Standard and Generalized Tightness Measures including Time-Varying Recruiting (RI) and Search Intensity (SI)



Source: Authors' calculations using CPS, JOLTS, and DFH (2012a) vacancies.

Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and constant recruiting intensity; Gen, All + RI = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and time-varying recruiting intensity; Gen, All + RI + SI = calibrated using V/ES (all twenty-two groups) with time-varying search intensities and recruiting intensity.

evaluating alternative functional forms for the matching function by comparing the calibrated job-filling and job-finding rates based on each to the actual data.[13] There is no inherent reason that incorporating time-varying recruiting and search intensity into these calibrated rates in fact will reduce the unexplained residual variation in the outcome of interest. Indeed, as will become clear, we obtain mixed results when we extend the general model in these directions.

Figure 9 depicts the actual and calibrated job-filling rates from the standard model, the general model with fixed relative job search intensities, the extension of that specification to include time-varying recruiting intensity, and finally a specification that includes both time-varying recruiting

13. Online appendix B includes the functional forms of job-filling and job-finding rates in this extended version of the general model.

Table 4. Relative Performance for Job-Filling and Job-Finding Rates using Standard versus General Model: Time-Varying Recruiting Intensity (RI) and Relative Job Search Intensity (SI)

	Std dev	*Corr w/ actual*	*RMSE*	*Ratio RMSE to standard*
A. Job-filling rate				
Actual	0.17	1.00	NA	NA
Standard	0.29	0.75	0.25	1.00
General, constant SI, RI = 1	0.19	0.81	0.13	0.54
General, time-varying RI	0.15	0.88	0.09	0.38
General, time-varying SI	0.17	0.84	0.12	0.49
General, time-varying RI and SI	0.14	0.91	0.09	0.34
B. Job-finding rate for the unemployed				
Actual	0.16	1.00	NA	NA
Standard	0.20	0.31	0.26	1.00
General, constant SI, RI = 1	0.15	0.62	0.17	0.67
General, time-varying RI	0.18	0.77	0.14	0.54
General, time-varying SI	0.29	0.68	0.22	0.85
General, time-varying RI and SI	0.31	0.75	0.22	0.83

Source: Authors' calculations using CPS; JOLTS; Davis, Faberman, and Haltiwanger (2012a) vacancies; and Bureau of Labor Statistics Gross Flows.

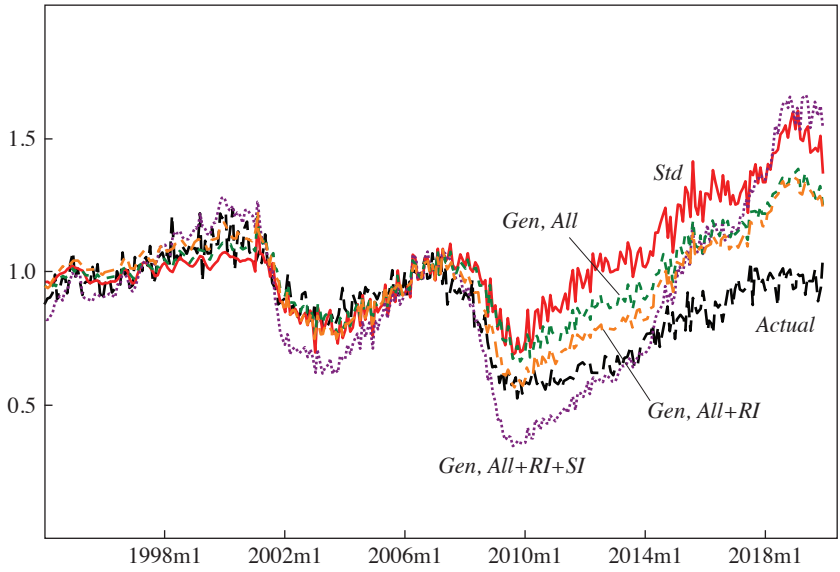
Notes: Panel A shows statistics for calibrated job-filling rates using matching function elasticities specific to each. Panel B shows statistics for calibrated job-finding rate for the unemployed.

intensity and time-varying relative search intensities. As can be seen in table 4, adding time-varying recruiting intensity helps to reduce the residual unexplained variation in the job-filling rate, lowering the RMSE in the generalized model's projected job-filling rate from 0.13 to 0.09. Allowing in addition for time-varying relative job search intensity has a much smaller effect that is not apparent when the numbers are rounded to two digits to the right of the decimal point, though the fourth column of the table shows that the RMSE declines relative to that in the standard model.

Figure 10 depicts the analogous exercise for the job-finding rate for the unemployed. Just adding time-varying recruiting intensity yields a modest improvement in the RMSE for the residual unexplained variation relative to the baseline case, lowering it from 0.17 to 0.14, but incorporating the time series variation in relative job search intensities leads to a worsening in the model's performance for tracking the job-finding rate among the unemployed, raising the RMSE to 0.22.

We regard the results shown in figures 9 and 10 as more suggestive than conclusive. The proxies we use for time variation in recruiting intensity and job search intensities are relatively crude and indirect. Developing better measures of these intensities is an important direction for future research.

Figure 10. Actual versus Calibrated Job-Finding Rates for the Unemployed with Standard and Generalized Tightness Measures including Time-Varying Recruiting (RI) and Search Intensity (SI)



Source: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and BLS Gross Flows.

Notes: All rates normalized to 1 in 2006. Actual = job-finding rate for unemployed from published BLS Gross Flows; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and constant recruiting intensity; Gen, All + RI = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and time-varying recruiting intensity; Gen, All + RI + SI = calibrated using V/ES (all twenty-two groups) with time-varying search intensities and recruiting intensity.

VII. Conclusions and Next Steps

The generalized measure of labor market tightness we have constructed based on the ratio of vacancies to effective searchers suggests that the US labor market was considerably less tight at the end of 2019 than implied by the standard ratio of vacancies to unemployment. The differing behavior of the two measures reflects the fact that the standard tightness measure does not account for important variation in search behavior on the part of workers. Job searchers include not only the unemployed but also those who are out of the labor force and the employed. In downturns, a more general index of effective searchers rises proportionally less than unemployment. The fact that the unemployed are only about 30 percent of all effective searchers contributes to this result, as any percentage increase in unemployment has a proportionally smaller effect on the overall number of effective

searchers. Another contributing factor is that, during a protracted contraction such as the Great Recession, the distribution of unemployment shifts toward the long-term unemployed, meaning that effective searchers rise less than the simple unemployment count.

The central question motivating our analysis is whether substituting a generalized measure of effective searchers for the standard unemployment measure reduces the need to appeal to fluctuations in matching efficiency to explain what is happening to employment flows. We observe that the Beveridge curve constructed using effective searchers is much more stable than the standard Beveridge curve. Further, the matching function for hires with our generalized measure of labor market tightness as its argument outperforms the matching function based on the ratio of vacancies to unemployment. Specifically, the calibrated job-filling rate (hires per vacancy) using the generalized measure tracks the actual job-filling rate much more closely than the job-filling rate calibrated using the standard measure of labor market tightness. The calibrated job-finding rate among the unemployed (hires from unemployment per unemployed person) based on the generalized measure also comes closer to tracking the actual series than the calibrated rate based on the standard measure. These findings imply that our approach reduces the unexplained residual variation required in the matching function to be consistent with the real-world data.

Our baseline effective searcher measure is constructed using data for twenty-two separate population groups. We also have examined a number of alternatives that are less inclusive or based on more aggregated population groups. Taken together, the results make clear that the key to better matching the actual time series behavior of the job-filling and job-finding rates is having a broad-based measure of effective searchers that also distinguishes among core groups. Measures that include the employed and, especially, those out of the labor force do better than those limited to the unemployed, but the very detailed disaggregation of the broader groups we have adopted does little better than the simpler five-category breakout we examined as a sensitivity check. Given the large differences in base-period job-finding rates across the detailed categories combined in the five-category version of the generalized measure, it might seem surprising that the further breakouts we apply do not add more to the performance of the job-filling and job-finding models. Over our period, the changes in composition occurring within the five more aggregated groups are not large enough to make much difference, but this might not always be the case. As there is little cost to using more disaggregated groups to construct the generalized tightness measure, it seems preferable to us to do so.

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

Our baseline measure of labor market tightness undoubtedly could be improved and built upon. In a suggestive analysis, we find that incorporating proxies for time variation in relative job search intensities and also taking into account variation in recruiting intensity across employers further reduces the unexplained variation in job-filling rates. We find mixed evidence on using these proxies for explaining the fluctuations in job-finding rates among the unemployed. There clearly is more to be done to develop time series measures of search and recruiting intensity.

One topic that we have deliberately avoided but that is of critically important interest is whether and how the generalized labor market tightness approach could improve our understanding of wage and price pressures. It would be interesting to explore the estimation of Phillips curve–type relationships using a generalized measure of labor market tightness rather than the unemployment rate gap as the central explanatory variable. Even if it is true that labor market tightness is a better predictor of wage and price changes than the unemployment rate, however, there are other sources of instability in the Phillips curve relationship that seem likely to pose problems for the estimation of such relationships. Still, given that estimating and interpreting Phillips curves is an active area of research and highly relevant for policymakers, there would be value in exploring the role of generalized labor market tightness measures in this context.

ACKNOWLEDGMENTS We are grateful to Daniel Aaronson, Steven Davis, Jason Faberman, Robert Hall, Chris Nekarda, Ayşegül Şahin, Sam Schulhofer-Wohl, Jim Stock, Robert Valletta, and Justin Wolfers for helpful conversations and comments on an earlier draft; to colleagues at the University of Maryland for comments at an early stage of this research that helped to sharpen our thinking about these issues; and to participants in both the Federal Reserve System's June 2019 Conference on Monetary Policy Strategy, Tools and Communication Practices and the March 2020 Brookings Panel on Economic Activity meeting for their comments.

References

Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan. 2015. "Changing Labor Force Composition and the Natural Rate of Unemployment," Chicago Fed Letter 338.

Abraham, Katharine G. 2015. "Is Skill Mismatch Impeding U.S. Economic Recovery?" *ILR Review* 68, no. 2: 291–313.

Abraham, Katharine G., John C. Haltiwanger, Kristin Sandusky, and James R. Spletzer. 2019. "The Consequences of Long-Term Unemployment: Evidence from Linked Survey and Administrative Data." *ILR Review* 72, no. 2: 266–99.

Ahn, Hie Joo, and Ling Shao. 2017. "Precautionary On-the-Job Search over the Business Cycle." Finance and Economics Discussion Series 2017–025. Washington: Board of Governors of the Federal Reserve System.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109, no. 3: 175–78.

Barnichon, Regis, and Geert Mesters. 2018. "On the Demographic Adjustment of Unemployment." *Review of Economics and Statistics* 100, no. 2: 219–31.

Black, Matthew. 1980. "Pecuniary Implications of On-the-Job Search and Quit Activity." *Review of Economics and Statistics* 62, no. 2: 222–29.

Blanchard, Olivier, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, no. 1, 1–76.

Blanchard, Olivier, and Peter A. Diamond. 1992. "The Flow Approach to Labor Markets." *American Economic Review* 82, no. 2: 354–59.

Blau, David M., and Philip K. Robins. 1990. "Job Search Outcomes for the Employed and Unemployed." *Journal of Political Economy* 98, no. 3: 637–55.

Broersma, Lourens, and Jan C. van Ours. 1999. "Job Searchers, Job Matching and the Elasticity of Matching." *Labour Economics* 6:77–93.

Casselman, Ben. 2018. "As Labor Pool Shrinks, Prison Time Is Less of a Hurdle." *New York Times*, January 13.

Crump, Richard K., Stefano Eusepi, Marc Gionnoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring U*." *Brookings Papers on Economic Activity*, Spring, 143–214.

Davis, Steven J. 2011. Comment on "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data," by Alan Krueger and Andreas Mueller. *Brookings Papers on Economic Activity*, Spring, 58–70.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012a. "Labor Market Flows in the Cross Section and Over Time." *Journal of Monetary Economics* 59, no. 1: 1–18.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012b. "Recruiting Intensity during and after the Great Recession." *American Economic Review* 102, no. 3: 584–88.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring," *Quarterly Journal of Economics*, 128, no. 2: 581–622.

DeLoach, Stephen B., and Mark Kurt. 2013. "Discouraging Workers: Estimating the Impacts of Macroeconomic Shocks on the Search Intensity of the Unemployed." *Journal of Labor Research* 34:433–54.

Diamond, Peter A. 1982. "Wage Determination and Efficiency in Search Equilibrium." *Review of Economic Studies* 49, no. 2: 217–27.

Elsby, Michael, Ryan Michaels, and David Ratner. 2015. "The Beveridge Curve: A Survey." *Journal of Economic Literature* 53, no. 3: 571–630.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2017. "Job Search Behavior among the Employed and the Non-Employed." Working Paper 23731. Cambridge, Mass.: National Bureau of Economic Research.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2020. "The Shadow Margins of Labor Market Slack." Working Paper 26852. Cambridge, Mass.: National Bureau of Economic Research.

Fujita, Shigeru, and Guiseppe Moscarini. 2017. "Recall and Unemployment." *American Economic Review* 107, no. 12: 3875–916.

Fujita, Shigeru, Guiseppe Moscarini, and Fabien Postel-Vinay. 2019. "Measuring Employer-to-Employer Reallocation." Working Paper.

Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante. 2018. "Aggregate Recruiting Intensity." *American Economic Review* 108, no. 8: 2088–127.

Gomme, Paul, and Damba Lkhagvasuren. 2015. "Worker Search Effort as an Amplification Mechanism." *Journal of Monetary Economics* 75:106–22.

Hall, Robert, and Sam Schulhofer-Wohl. 2018. "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers." *American Economic Journal: Macroeconomics* 10, no. 1: 1–32.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erica McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics* 10, no. 2: 52–85.

Haltiwanger, John C., Henry R. Hyatt, and Erica McEntarfer. 2018. "Who Moves Up the Job Ladder?" *Journal of Labor Economics* 36, no. S1: S301–36.

Hornstein, Andreas, and Marianna Kudlyak. 2016. "Estimating Matching Efficiency with Variable Search Effort." Working Paper 2016–24. San Francisco: Federal Reserve Bank of San Francisco.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange. 2014. "Measuring Resource Utilization in the Labor Market." *Federal Reserve Bank of Richmond Economic Quarterly* 100, no. 1: 1–21.

Jones, Stephen R. G., and W. Craig Riddell. 1999. "The Measurement of Unemployment: An Empirical Approach." *Econometrica* 67, no. 1: 147–62.

Kaitz, Hyman. 1970. "Analyzing the Length of Spells of Unemployment." *Monthly Labor Review* 93, no. 11: 11–20.

Katz, Lawrence F. 1986. "Layoffs, Recalls and the Duration of Unemployment." Working Paper 1825. Cambridge, Mass.: National Bureau of Economic Research.

Katz, Lawrence F., and Bruce D. Meyer. 1990. "The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment." *Journal of Public Economics* 41, no. 1: 45–72.

Krueger, Alan B., Judd Cramer, and David Cho. 2014. "Are the Long-Term Unemployed on the Margins of the Labor Market?" *Brookings Papers on Economic Activity*, Spring, 229–80.

Krueger, Alan B., and Andreas Mueller. 2010. "Job Search and Unemployment Insurance: New Evidence from Time Use Data." *Journal of Public Economics* 94:298–307.

Krueger, Alan B., and Andreas Mueller. 2011. "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data." *Brookings Papers on Economic Activity,* Spring, 1–57.

Kudlyak, Marianna. 2017. "Measuring Labor Utilization: The Non-Employment Index." Federal Reserve Bank of San Francisco Economic Letter 2017–08, March 17.

Modestino, Alicia Sasser, Danial Shoag, and Joshua Balance. Forthcoming. "Upskilling: Do Employers Demand Greater Skill when Workers Are Plentiful?" *Review of Economics and Statistics*.

Mongey, Simon, and Gianlucia Violante. 2020. "Macro Recruiting Intensity from Micro Data." Working Paper 26231. Cambridge, Mass.: National Bureau of Economic Research.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61, no. 3: 397–415.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin. 2018. "Job Search Behavior over the Business Cycle." *American Economic Journal: Macroeconomics* 10, no. 1: 190–215.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity*, no. 3: 411–48.

Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39, no. 2: 390–431.

Pissarides, Christopher. 2000. *Equilibrium Unemployment Theory*, 2nd ed. Cambridge, Mass.: MIT Press.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104, no. 11: 3529–64.

Sedlacek, Petr. 2016. "The Aggregate Matching Function and Job Search from Employment and Out of the Labor Force." *Review of Economic Dynamics* 21:16–28.

Shimer, Robert. 2001. "The Impact of Young Workers on the Aggregate Labor Market." *Quarterly Journal of Economics* 116, no. 3: 969–1008.

Shimer, Robert. 2004. "Search Intensity." Working paper. Chicago: University of Chicago.

Smialek, Jeanna. 2019. "The Economy Is Strong and Inflation Is Low. That's What Worries the Fed." *New York Times*, May 21.

Veracierto, Marcelo. 2011. "Worker Flows and Matching Efficiency." *Economic Perspectives* 35, no. 4: 147–69.

Comments and Discussion

COMMENT BY

AYŞEGÜL ŞAHIN Abraham, Haltiwanger, and Rendell developed a generalized measure of labor market tightness which takes into account intensive and extensive margins of search activity on both demand and supply sides of the labor market. Their measure captures the hiring process in the US economy better than the standard measure of labor market tightness. Their success is probably not surprising given that they build on two well-documented facts about the hiring process in the United States. The first is that the majority of jobs are filled by workers who are not unemployed. The second is that the number of vacancies is an imperfect proxy for firms' total recruiting effort since firms vary their recruiting intensity over the business cycle.

Assessment of labor market conditions is a fundamental issue in macroeconomics and a key input to implementation of monetary and fiscal policy. The authors' work is a valuable addition to the wealth of labor market indicators developed in the last decade to better evaluate labor market developments in light of ongoing secular trends. With the unemployment rate jumping from 3.5 percent in February 2020 to 14.7 percent in April 2020, the labor market will be our main focus of attention for years to come. Understanding how workers search for employment opportunities, how firms fill their open positions, and how search activity responds to macroeconomic conditions will help us in characterizing the adjustment path of the US labor market to the COVID-19 shock. This comment reviews and interprets the authors' findings and suggests new directions of research.

FRAMEWORK Abraham, Haltiwanger, and Rendell use the matching function specification—a key building block of the Diamond-Mortensen-Pissarides search and matching framework—to characterize the behavior

of hiring in the US economy. In its basic form, the matching function takes the Cobb-Douglas form and specifies hires, h_t, as a function of two inputs: vacancies posted by firms looking to hire, v_t, and unemployed workers looking for jobs, u_t:

$$h_t = \Phi_t (u_t)^{\alpha} (v_t)^{1-\alpha}, \tag{1}$$

where $\alpha \in (0,1)$ is the unemployment share and Φ_t is the aggregate matching efficiency parameter. The job-finding and the job-filling rates can be written as a function of labor market tightness ($\theta_t = v_t/u_t$) as $\Phi_t\ \theta_t^{1-\alpha}$ and $\Phi_t (1/\theta_t)^{\alpha}$, respectively. This specification implies that θ_t is the only determinant of the hiring process if Φ_t does not vary over time. However, as discussed by the authors, this specification ignores the fact that most hires originate from employment or nonparticipation. Moreover, it does not take into account the intensive margins of firm and worker search effort. A generalized tightness measure, $\tilde{\theta}_t$, which incorporates these factors can be defined as

$$\tilde{\theta}_t = \frac{\sum_j \rho_t^{V_j} V_{jt}}{\sum_i \rho_t^{S_i} S_{it}}, \tag{2}$$

where $\rho_t^{V_j}$ is the recruiting intensity of vacancy group in firm type j and $\rho_t^{S_i}$ is the search intensity of searcher group i. The interpretation of $\tilde{\theta}_t$ is straightforward: it is the ratio of the recruiting-intensity-weighted sum of vacancies to the search-effort-weighted sum of searchers in the economy. While searchers include all individuals regardless of their labor force status, not all groups contribute equally to the total search effort in the economy. The contribution of each group depends on their relative search intensity, $\rho_t^{S_i}$.

Abraham, Haltiwanger, and Rendell do not explore the variation in the composition of firms by age, size, or industry over the business cycle. Instead, they use the direct measure of aggregate recruiting intensity from Davis, Faberman, and Haltiwanger (2013). With the aggregate recruiting intensity ρ_t^V in hand, they compute the effective vacancies as $\rho_t^V V_t$. The implementation of the generalized measure of effective searchers is more involved since it requires using detailed micro data from the Current Population Survey (CPS). First, the authors identify twenty-two distinct labor force states following Hall and Schulhofer-Wohl (2018) and compute the job-finding rates of these groups using the matched CPS data. They assume that the relative job-finding rate of each group corresponds to its relative search intensity. This assumption allows them to compute an

aggregate measure of effective searchers by weighting each group *i* by its relative job-finding rate. Among these twenty-two detailed labor market groups, thirteen groups are unemployed; seven groups are those who are out of the labor force, and two are the employed. Table 1 of the paper shows that these groups vary considerably in their job-finding rates with recently temporarily laid-off workers having the highest job-finding rate.

ARE TWENTY-TWO LABOR MARKET STATES NECESSARY? Abraham, Haltiwanger, and Rendell differentiate between various detailed labor market states and end up using twenty-two distinct groups in their implementation. While this level of disaggregation allows them to capture even small variations in the relative job-finding rates of searchers, it restricts their analysis to the post-1994 period since the questions that allow them to differentiate between these twenty-two states are not available in the CPS before 1994. I show that focusing only on a few broad labor market states is sufficient to capture almost all the variation in the measure of effective searchers over time. This simplification is useful for two reasons. First, it makes it possible to compute the generalized tightness measure without using the CPS micro data. Second, it allows us to compute the generalized tightness measure starting in 1948.

Starting with five broad labor force states, the measure of effective searchers can be simplified as:

$$S_t = \underbrace{\rho^{UST} U_t^{ST} + \rho^{ULT} U_t^{LT}}_{\text{Unemployed}} + \underbrace{\rho^{W} W_t + \rho^{N} N_t}_{\text{Nonparticipants}} + \underbrace{\rho^{E} E_t}_{\text{Employed}}.$$

I focus on unemployment, nonparticipation, and employment and further distinguish between short-term (U^{ST}) and long-term unemployed (U^{LT}) among the unemployed and between those who want a job (W) and those who do not (N) among nonparticipants. I approximate the relative job-finding rates of these groups using the raw job-finding rates from the authors' table 1 and choose the following weights:

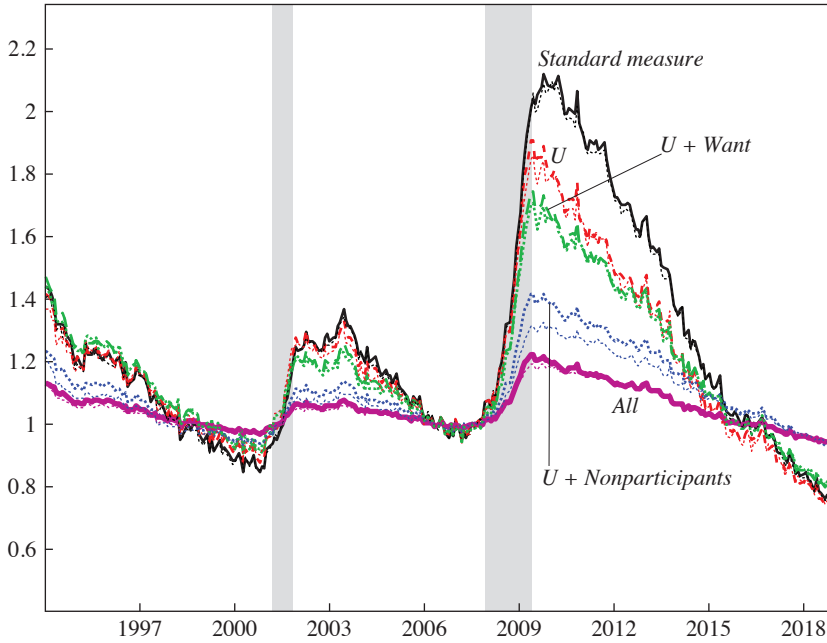
$$\rho^{UST} = 1, \rho^{ULT} = 0.48, \rho^{W} = 0.40, \rho^{N} = 0.09, \rho^{E} = 0.07.$$

I then compute the rate of effective searchers as

$$S_t = \frac{S_t}{P_t} = \frac{S_t}{U_t + N_t + E_t}$$

and normalize it to 1 in 2006. Figure 1 shows that the series computed using only five labor market states line up remarkably well with the authors'

Figure 1. Measures of Effective Searchers using Five Labor Force States



Source: Author's calculations based on the Current Population Survey (CPS).

Notes: AHR measures indicated by thin dotted lines. Standard measure: U_t/P_t; U: $(U_t^{ST} + 0.48U^{LT})/P_t$; U + Want: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t)/P_t$; U + Nonparticipants: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t)/P_t$; and all: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t + 0.07E_t)/P_t$. Series normalized to 1 in 2006.

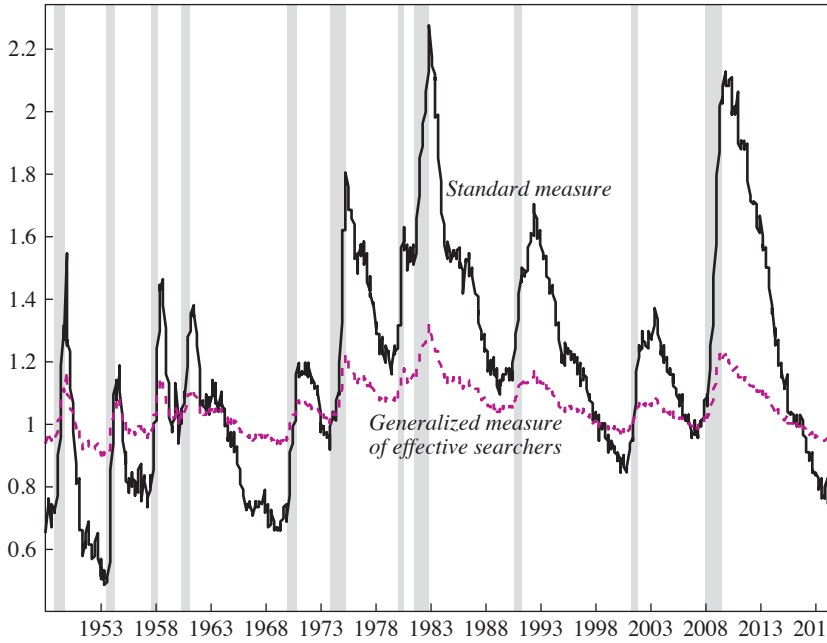
measures.[1] The intuition is clear: while job-finding rates vary substantially by detailed labor market state, most of these groups are too small to affect the aggregate measure. Therefore I conclude that focusing on five broad labor market states does not change the essence of the authors' findings but simplifies their analysis considerably by making it possible to compute the index without using the CPS micro data.

HISTORICAL PERSPECTIVE Given that the five-state simplification does remarkably well, I now turn to historical data and calculate the measure of effective searchers starting in 1948. This requires consolidating the two groups of nonparticipants to only one, but even this abstraction does not change the behavior of the generalized measure. Figure 2 plots the generalized measure of effective searchers along with the standard measure

1. While these weights are not exact, following my discussion at the panel, the authors repeated their analysis with only five groups and found very similar results to mine.

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

Figure 2. Generalized Measure of Effective Searchers and the Standard Measure in the 1948–2018 Period



Source: Author's calculations based on the Current Population Survey (CPS).

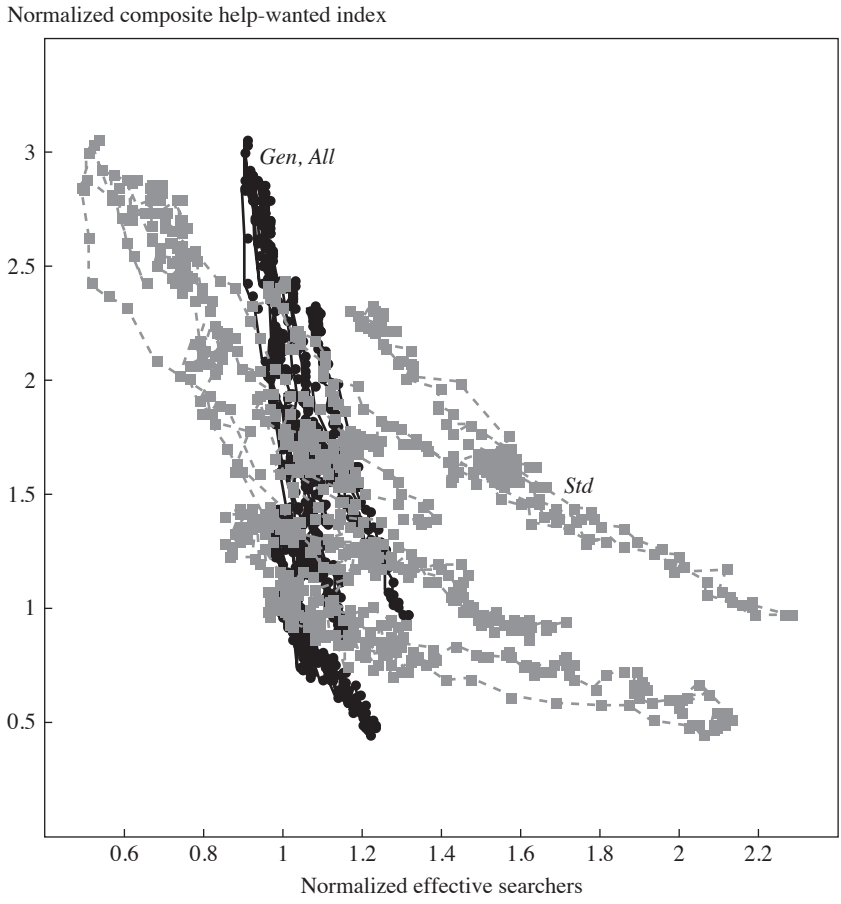
Notes: AHR measures indicated by dotted lines. Standard measure: U_t/P_t; generalized measure of effective searchers: $(U_t^{ST} + 0.48U^{LT} + 0.09N_t + 0.07E_t)/P_t$.

starting in 1948 and shows that the deviation between the two measures is not specific only to the most recent expansion. The effective searchers measure exhibits more muted countercyclical behavior than the standard measure throughout the sample—a distinct feature of the generalized measure that I will discuss below in detail.

Interestingly, despite differences in the magnitude of their cyclicality, measures based on unemployment and effective searchers provide a very similar assessment of the historical business cycles. Both measures imply that the 1973–1975, 1980–1982, and 2007–2009 recessions were the deepest downturns and the second halves of the 1960s, 1990s, and 2010s were the tightest labor markets of the postwar period.

I also examine the historical Beveridge curve using the standard measure and the measure of effective searchers combined with Barnichon's (2010) composite help-wanted vacancy index in figure 3. The Beveridge

Figure 3. Historical Beveridge Curves using the Generalized Measure of Effective Searchers and the Standard Measure

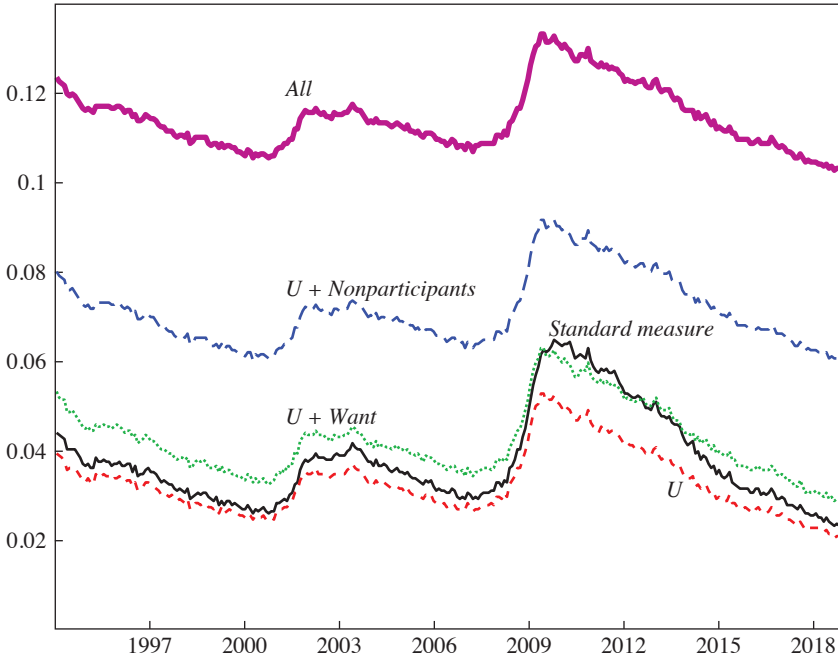


Source: Author's calculations based on the Current Population Survey (CPS) and composite help-wanted vacancy index developed by Barnichon (2010).

Note: Historical Beveridge curve using the unemployment-to-population ratio (gray dots) and the effective measure of searchers (black dots).

curve constructed using unemployment exhibits substantial horizontal shifts over time as is well documented in Diamond and Şahin (2015). The shifts are much less pronounced when the measure of effective searchers is used. However, the Beveridge curve also becomes much steeper when it is constructed with the new measure. The steepness of the Beveridge curve suggests that fluctuations in vacancies have very little effect on labor

Figure 4. Measures of Effective Searchers in Levels



Source: Author's calculations based on the Current Population Survey (CPS).

Notes: Alternative searcher measures in levels. Standard measure: U_t/P_t; U: $(U_t^{ST} + 0.48U^{LT})/P_t$; U + Want: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t)/P_t$; U + Nonparticipants: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t)/P_t$; and all: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t + 0.07E_t)/P_t$.

market underutilization. This interesting observation deserves further discussion and analysis since it suggests that the link between labor demand and labor underutilization has been considerably weaker than is typically assumed in the postwar period.

WHY IS THE COUNTERCYCLICALITY DAMPENED? Abraham, Haltiwanger, and Rendell focus only on normalized measures of effective searchers and do not report or interpret the *levels* of their measures. While this type of normalization is useful in the context of the matching function framework, I argue that the difference in the levels of alternative measures of searchers is important for the interpretation of the authors' findings.

Figure 4 plots different measures of searchers in levels without normalizing them to 1 in 2006. Comparison of figures 1 and 4 shows that measures plotted in levels look very different than their normalized counterparts. Expanding the pool of searchers by adding employed

workers and nonparticipants increases the level substantially despite the low relative job-finding rates of these groups, and the normalization reduces cyclicality of the measures substantially.

This comparison is very helpful in understanding the reasons for dampened countercyclicality of the effective searchers measure relative to unemployment. First, the level of the generalized measure is higher than the level of the standard measure due to the addition of the high number of employed workers and nonparticipants. Second, there are offsetting changes in the composition of searchers over the business cycle. During recessions, unemployment increases but due to the decline in employment, the number of employed searchers declines, dampening the rise in unemployment. Similarly, during expansions as the number of unemployed searchers declines, the number of employed searchers increases, moderating the decline in the number of unemployed searchers. As a result, the generalized measure of effective searchers fluctuates less than unemployment over the business cycle. This finding has an important implication about the hiring incentives of firms: firms do not find it that much easier to fill jobs during recessions despite high unemployment since the number of employed searchers declines. This is why the generalized tightness measure does better in capturing the behavior of hires in the matching function framework.

HOW SHOULD WE ASSESS LABOR MARKET CONDITIONS? The Great Recession and the subsequent period of sluggish recovery in the labor market triggered an important line of research, with numerous studies developing labor market indicators as alternatives to the unemployment rate. Most of these measures exploit one of the following two approaches. The first focuses on estimating a time-varying natural rate of unemployment and uses the unemployment gap—the difference between the actual and the natural rate of unemployment—as a measure of labor market tightness. While this line of research typically focuses on estimating the natural rate of unemployment using only simple aggregate data, more recent work unifies the macro approach with rich labor market data (Crump and others 2019). The main advantage of this approach is that it provides a unified framework that takes into account secular demographic trends, wage and price inflation, and inflation expectations. Therefore it directly connects to maximum employment and price stability objectives.

The second approach is to develop broader measures which take into account additional margins of labor market underutilization. A common practice is to weight different groups of workers depending on their demographic characteristics, wages, search activity, or job-finding rates,

such as in Perry (1970), Hornstein, Kudlyak, and Lange (2014), and this paper. While implementing this approach, choice of weights is an important issue since the interpretation of the aggregate measure depends heavily on the weights used. For example, the authors choose to use relative job-finding rates and show that their measure relying on realized transitions is informative about the hiring process in the economy. However, this does not necessarily mean that their measure would be preferable to other measures for other purposes. For example, one fundamental issue in macroeconomics is to estimate the potential output of the economy. In that case, an alternative measure which weights workers by their average idle hours (desired hours minus actual hours worked), as in Faberman and others (2020), would be informative about potential output.

To conclude, the authors' generalized tightness measure is a valuable addition to our arsenal of labor market indicators. They convincingly demonstrated that the measure they develop captures the essence of the hiring process well. Future work calls for broadening the metric for success by considering the usefulness of the new metric in capturing wage growth and inflation. This extension would help to connect the new measure more directly to monetary policy implementation. With the unemployment rate at its highest level in the postwar period, we need as much information as possible about the labor market for real-time assessment of the state of the labor market and its medium-term evolution. The work of Abraham, Haltiwanger, and Rendell undoubtedly will contribute to our understanding of the labor market in the years to come.

REFERENCES FOR THE ŞAHIN COMMENT

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109, no. 3: 175–78.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring U*." *Brookings Papers on Economic Activity*, Spring, 143–214.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring." *Quarterly Journal of Economics* 128, no. 2: 581–622.

Diamond, Peter A., and Ayşegül Şahin. 2015. "Shifts in the Beveridge Curve." In *Research in Economics*, edited by Federico Etro. Amsterdam: North-Holland.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2020. "The Shadow Margins of Labor Market Slack." Working Paper 26852. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w26852.

Hall, Robert, and Sam Schulhofer-Wohl. 2018. "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers." *American Economic Journal: Macroeconomics* 10, no. 1: 1–32.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange. 2014. "Measuring Resource Utilization in the Labor Market." *Federal Reserve Bank of Richmond Economic Quarterly* 100, no. 1: 1–21.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity*, no. 3: 411–48.

COMMENT BY

JUSTIN WOLFERS This paper by Abraham, Haltiwanger, and Rendell follows a long tradition of Brookings papers in proposing a new measure of labor market tightness. The new measure aims to capture effective job search effort, and its key advantages are that it recognizes that the unemployed aren't the only group who search and it accounts for the reality that different groups of workers search with different degrees of success. To be precise, the proposed new measure—which the authors call the "generalized measure of effective searchers"—is constructed as a weighted average of the share of the population in twenty-two different labor market states, with the weights reflecting each group's baseline level of search effectiveness (measured by each group's historical job-finding rate). As a weighted average rather than simple average, it generalizes more standard metrics like unemployment (which effectively puts a weight of one on the unemployed and zero on all others).

This paper was written during what now seems a distant time in early 2020 when the official unemployment rate was lower than it had been in about half a century and one of the central macro policy questions of the day was whether unemployment could go any lower. The final draft was submitted a few weeks later, when the coronavirus pandemic had pushed the unemployment rate to levels higher than at any time since the Great Depression. A superficial assessment might argue that a paper titled "How Tight Is the US Labor Market?" has been overtaken by events which have rendered it less relevant. But in reality, this research question may actually be more relevant than ever before, as the recession caused by the coronavirus—which includes elements of demand shocks and supply shocks—has led to arguably greater uncertainty about the extent of slack in the labor market than the United States has experienced at any point throughout the postwar period. The present-day relevance simply requires a reframing from this being a paper assessing how best to measure labor market tightness to one assessing how best to measure labor market looseness.

In proposing a new measure of labor market conditions, this paper enters an already quite crowded marketplace. Between the various measures of labor market slack published by the Bureau of Labor Statistics (BLS), alternatives tracked by the Federal Reserve banks, the preferred metrics of an array of private-sector forecasters, and new measures proposed in back issues of *Brookings Papers on Economic Activity*, it's no exaggeration to say that dozens of measures have been proposed and are closely tracked. Whether it's worth tracking one more measure depends on whether this new measure yields additional useful information.

The main test that the paper offers is to ask whether this new measure does a better job in explaining time series variation in the job-filling and job-finding rates than standard metrics like the official BLS measure of unemployment. The authors argue that their measure does add useful additional information. The analysis that I present below disagrees. To preview, I find little evidence that their proposed new measure outperforms standard measures. While the authors claim the generalized measure yields a better fit, the tests reported reflect arbitrary auxiliary assumptions that tilt the playing field against the standard measure of unemployment. When these assumptions are relaxed, the proposed new measures no longer appear to be more predictive. Indeed, the proposed new measure appears to largely track the official measure of unemployment. By this telling, there is no problem with the newly proposed measure, but there is also little to recommend it, as it adds little information beyond that in the standard measure.

EXPLAINING TIME SERIES VARIATION IN THE JOB-FILLING RATE The main test that the paper implements is to ask which measure of labor market slack—the proposed generalized measure of effective searchers or the widely used official measure of unemployment—does a better job explaining time series variation in the job-filling rate, which is the ratio of hiring to vacancies. To motivate their analysis, the authors start with a Cobb-Douglas matching function of the form:

$$H_t = \mu U_t^{\alpha} V_t^{1-\alpha}.$$

This can be understood as a production function, where the output is the number of people hired, H_t, and the production of these matches is a function of the aggregate effort workers put into searching for jobs (U_t) and the aggregate effort that firms put into searching for workers (V_t). In addition, μ is an index of the efficiency of this production function, and the empirical analysis in this paper holds this parameter constant, which is why it lacks a time subscript.

It is common for empirical analyses of the matching function to use the level of unemployment (which I'll denote U_t^s) as a proxy for total search effort by workers (so $U_t = U_t^s$) and the level of vacancies as a proxy for total search effort by firms (V_t). The authors propose that their generalized measure of effective searchers (which I'll denote U_t^g) may be a better proxy for the level of total search effort by workers (so $U_t = U_t^g$, instead). In a later extension, they also argue that accounting for recruiting intensity may yield a better proxy of the search effort by firms.

To test this, the paper asks which competing proxy measure of U_t—the standard measure, which is the number of unemployed people U_t^s, or their generalized measure of effective searchers U_t^g—better explains the observed time series movements in the job-filling rate. The specific empirical exercise it implements comes from a simple rearrangement of the matching function so that the dependent variable is now the job-filling rate:

$$\frac{H_t}{V_t} = \mu \left[\frac{V_t}{U_t} \right]^{-\alpha}.$$

The variable to be forecast, H_t/V_t, is readily measured as monthly hires divided by monthly vacancies, both from the Job Openings and Labor Turnover Survey (JOLTS) data. Likewise, the two candidate inputs into this forecasting equation—U_t/V_t, measured as either U_t^S/V_t, the ratio of the unemployed (from the Current Population Survey [CPS]) to the number of vacancies, or as U_g^g/V_t, the ratio of the generalized measure of effective searchers (constructed by the authors from CPS data) to the number of vacancies—are easily measured.

So far, this sounds like a standard forecast evaluation exercise. But there's a twist. No statistical agency publishes estimates of either μ (the efficiency of the matching function) or α (the Cobb-Douglas coefficient in the matching function). Think of the parameter μ as determining the position of the Beveridge curve, while the parameter α determines its slope. The values of these parameters will shape how well any measure of slack will fit the data. The values that are imposed reflect an auxiliary set of assumptions, and as I will show below, these assumptions largely drive the empirical findings that follow.

Unfortunately, the paper does not contain much discussion of where these parameters come from nor much of a defense of the specific parameter values that are imposed. In order to appropriately assess the quality of the evidence presented in the paper, it's worth being more transparent

about these assumptions. First, μ reflects an arbitrary normalization, in which the time series of forecasts of the job-filling rate (and the inputs into that forecast) are indexed to be equal to 1, on average, during 2006. This effectively imposes

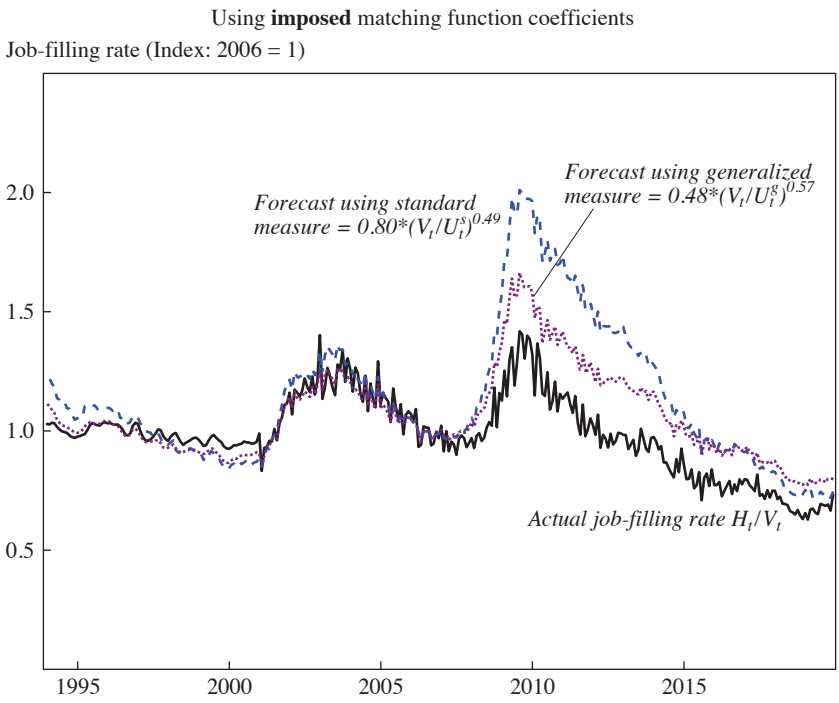
$$\mu = \frac{\overline{H_t/V_t}^{2006}}{\overline{(V_t/U_t)^{-\alpha}}^{2006}},$$

which would make sense if μ were constant and the matching function exactly fit the data (with an error term equal to zero) on average through 2006. When U_t is proxied by the level of unemployment, this yields $\widehat{\mu^s} = 0.80$, but when it is proxied by the generalized measure of effective searchers, it yields $\widehat{\mu^g} = 0.48$. Next, α comes from a structural estimation procedure described only in online appendix B of the paper, in which the evolution over time of the job-finding rates for each of the twenty-two groups are estimated as a function of a group fixed effect, a common time trend, and the aggregate vacancy-to-hires ratio. In this framework, α is recovered as a nonlinear function of the coefficient on the vacancy-to-hires ratio, and the various α's for each of the twenty-two groups are averaged to get an aggregate $\widehat{\alpha^g}$. (I append the superscript g to denote that this is the α estimated for the *g*eneralized measure). However, this equation is not estimated directly but rather using instrumental variables, where employment from the payroll survey is an instrument for the vacancy-to-hires ratio. For the generalized measure of effective searchers, this yields $\widehat{\alpha^g} = 0.57$. For the standard measure, this procedure collapses so that there is only one group of searchers (the unemployed), and so the dependent variable is the ratio of unemployment-to-employment flows to unemployment, and the dependent variables are a constant, a time trend, and the aggregate vacancy-to-hires ratio (which again is instrumented using payroll employment). This yields $\widehat{\alpha^s} = 0.49$ (where the superscript s denotes that this is the α that applies to the *s*tandard measure).

Imposing these assumptions yields the forecasts of the job-filling rate shown in figure 1 (which replicates exactly the corresponding lines from figure 5 in the paper). The top (dashed) line is the forecast generated using the standard measure of U_t, the next (dotted) line is the forecast generated using the generalized measure, while the actual job-filling rate, which is shown as a solid line, lies below both of these forecasts for most of the sample. The authors emphasize that the forecast generated using the standard

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

Figure 1. Predicting the Job-Filling Rate



Source: Author's calculations, from data in the paper.

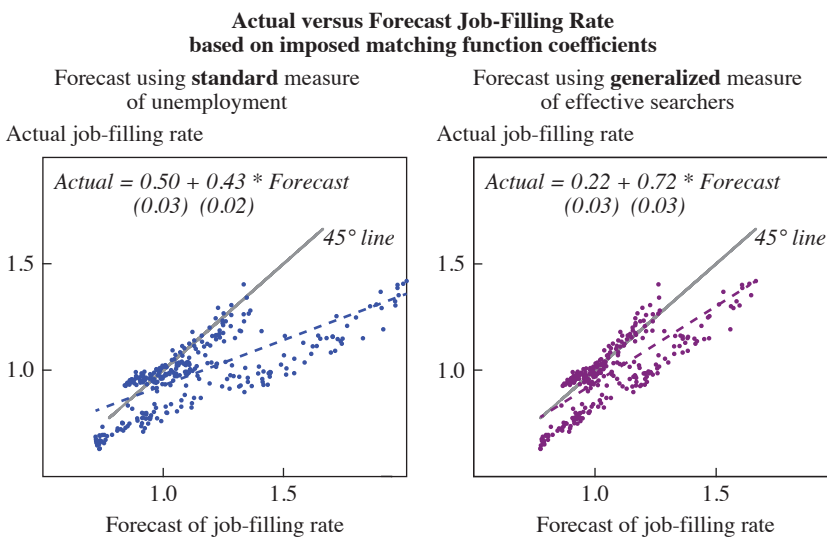
measure is further from the actual outcomes than the forecast generated using the generalized measure. This difference is the basis of the claim that the standard measure is outperformed by the authors' preferred generalized measure.

But eyeballing this figure reveals something more troubling: both lines look like problematic forecasts. Both are (substantially) higher on average than the actual job-filling rate. And both forecasts rise and fall substantially more over the business cycle than the actual job-filling rate. Neither forecast appears to fit the data well.

This is easy to confirm using standard forecast evaluation methods. In particular, figure 2 shows scatter plots of the actual job-filling rate against each of these forecasts. Neither fits the data well. Figure 2 also reports the results of regressions of the following form:

$$Actual_t = a + b \times Forecast_t + \epsilon_t.$$

Figure 2. Forecast Evaluation



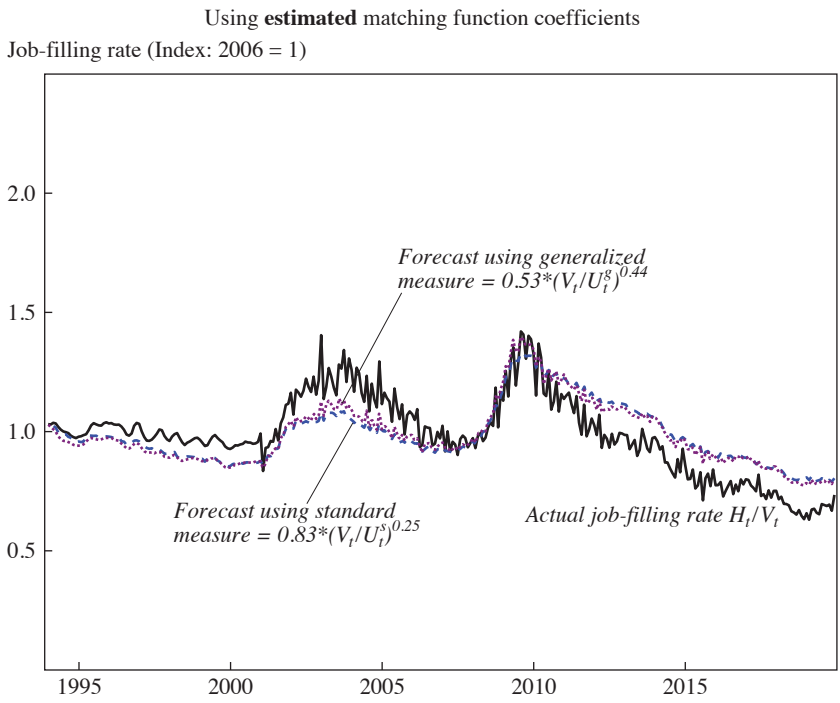
Source: Author's calculations.

An unbiased and efficient forecast would yield a constant term of zero and a coefficient on the forecast of one. (Failing this test implies that forecast errors are forecastable, and hence statistically inefficient.) The left panel shows that the forecast generated using the standard measure of unemployment fails both of these tests, and the right panel shows that the forecast generated using the generalized measure of effective searchers also fails the same test. (F-tests of the joint null that a = 0 and b = 1 are overwhelmingly rejected in both cases.)

These artifacts reflect the fact that the μ and α coefficients that were imposed appear to be inappropriate if one is interested in forecasting the job-filling rate. In addition, as we shall see in a moment, the differences between the two forecasts are largely attributable to these imposed coefficients. Correcting for these artifacts reveals that both the standard and the generalized measure yield very similar forecasts.

A more transparent (and arguably less arbitrary) approach to figuring out what values of μ and α to apply for this forecast evaluation exercise would be to simply estimate each of these parameters directly and then use those estimates to assess the accuracy of the ensuing forecasts. This alternative would be more consistent with the authors' claim that their analysis aims to estimate α in a way that would "give each set of matching function

Figure 3. Predicting the Job-Filling Rate Based on Estimated Matching Function Coefficients



Source: Author's calculations.

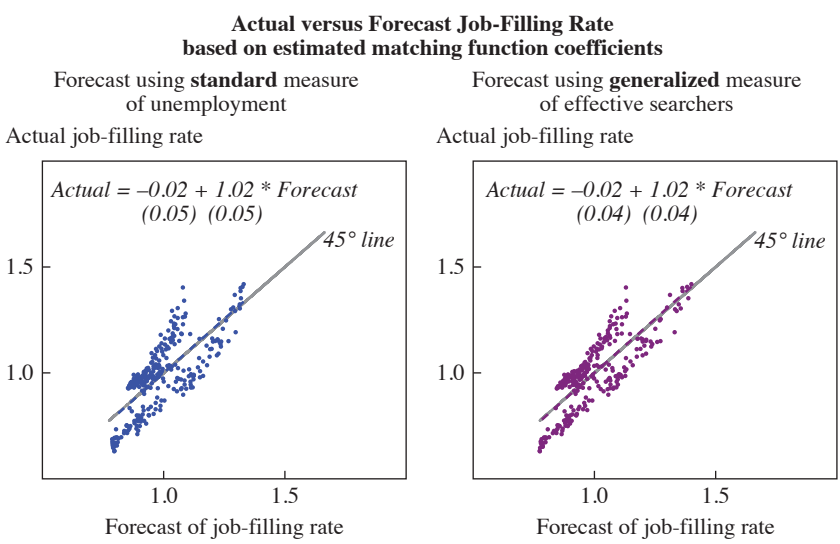
arguments the best possible chance to fit the data well." As such, I ran the following simple nonlinear least squares regression:

$$\frac{H_t}{V_t} = \mu\left[\frac{V_t}{U_t}\right]^{-\alpha}.$$

Initially, I ran this regression using the standard measure of U_t^s (the number of unemployed people), which yielded estimates of $\widehat{\mu^s} = 0.83$ (se = 0.01) and $\widehat{\alpha^s} = 0.25$ (se = 0.01). Running this regression, but using the generalized measure of effective searchers U_t^s as the measure of worker search instead, yielded estimates of $\widehat{\mu^g} = 0.53$ (se = 0.014) and $\widehat{\alpha^g} = 0.44$ (se = 0.02). (Running the regression in logs: $\log(H_t) = \log(\mu) + \alpha\log(U_t) + (1-\alpha)\log(V_t)$ yielded relatively similar estimates of α in each case.)

Figure 3, which is drawn on the same scale as figure 1, plots each of the resulting forecasts of the job-filling rate, along with the actual job-filling

Figure 4. Forecast Evaluation Based on New (Estimated) Forecasts



Source: Author's calculations.

rate. This yields a dramatically different picture than figure 1 (or figure 5 in the paper under discussion). Both sets of forecasts of the job-filling rate are now much more accurate.

Most importantly, the dashed and the dotted lines are now almost identical. Once one abandons the arbitrarily imposed matching function coefficients, the forecast for the job-filling rate generated using the standard measure of unemployment is almost identical to that forecast using the generalized measure of effective searchers. Indeed, the correlation between these two series is 0.991.

This analysis yields an interpretation of the paper's findings that is largely at odds with the interpretation offered by the authors. It suggests that any differences in the performance of these measures is not due to richer information embedded in the generalized measure of effective searchers but rather is due to the specific coefficients imposed on the matching function in generating forecasts. As such, the authors' claim that "the relatively poor performance of the standard model cannot be rescued with an alternative estimate of the matching function elasticity" appears to be wrong.

Not surprisingly, estimating the coefficients that we use to generate forecasts of the job-filling rate yields better behaved forecasts. Figure 4 plots each of our new forecasts of the job-filling rate against the actual

Table 1. Formal Forecast Evaluation

	Standard deviation	*Correlation with actual*	*RMSE*	*Ratio of RMSE to standard*
Panel A: Generating forecasts using imposed coefficients (from table 3 of the paper)				
Actual job-filling rate	0.17	1.00	—	—
Forecast using standard measure of unemployment	0.29	0.75	0.25	1.00
Forecast using generalized measure of search activity	0.19	0.81	0.13	0.54
Panel B: Generating forecasts using estimated coefficients (preferred alternative)				
Actual job-filling rate	0.17	1.00		
Forecast using standard measure of unemployment	0.12	0.76	0.11	1.00
Forecast using generalized measure of search activity	0.13	0.81	0.10	0.90

Source: Author's calculations.

job-filling rate based on the coefficients estimated above. Forecasts generated using either the standard or the generalized measure of labor market slack are unbiased, and they both rise and fall in proportion with the business cycle. Of course, this occurs largely because these forecasts are generated using coefficients that were estimated with an eye to closely matching the actual outcomes. (In the regression of actual outcomes on forecasts, while the constant is close to zero, it is not precisely zero—and the slope is close to one, but not precisely one—because these forecasts were generated from an equation that is not linear in the parameters.)

At this point, it is worth revisiting the authors' evaluation of the relative performance of each measure. Table 1 presents a formal forecast evaluation. Panel A replicates the findings the authors present in their table 3, showing that using the imposed matching function coefficients leads to the conclusion that the root-mean-square error (RMSE) of forecasts generated using the standard measure is much higher (that is, worse) than that of the forecasts generated using the generalized measure of search activity. Indeed, the authors' preferred measure yields forecast errors that tend to be nearly half that when using the standard measure. A Diebold-Mariano test reveals that this difference in RMSE is (highly!) statistically significant.

But this finding no longer survives when one is no longer willing to impose the specific values that the authors impose on the matching function coefficients. To see this, panel B has a parallel structure to panel A, but it reports on the accuracy of the forecasts generated using each measure of labor

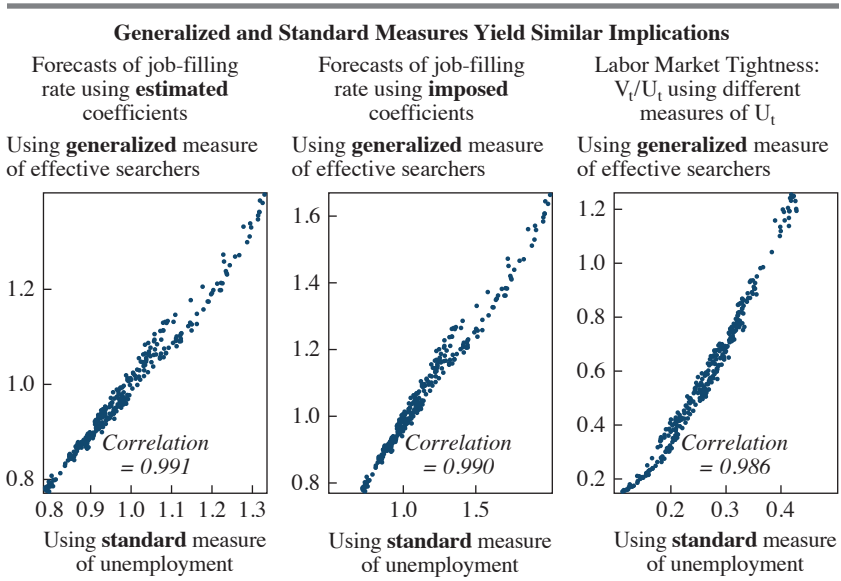
market slack, but using the matching function coefficients that I estimated above. The fit of both sets of forecasts—generated using either measure of labor market slack—is now much better. Moreover, the performance of the forecasts generated using the competing measures of slack—whether measured as the correlation between the forecast and the actual outcome or as the RMSE—are now remarkably similar. The RMSE of the two alternative forecasts evaluated in panel B are no longer economically meaningful, and a Diebold-Mariano test fails to reject the null that they're equal. All told, the evidence in panel B is that there's little to recommend one measure over the other. In turn, this reflects the reality that the alternative forecasts are remarkably similar.

THE REMARKABLE SIMILARITY OF ALTERNATIVE MEASURES OF SLACK The most striking finding so far is that using the estimated matching function coefficients yields forecasts of the job-filling rate that are almost identical whether using the standard measure of unemployment or the generalized measure of effective searchers as an input. It's worth pausing a bit to see where this comes from.

The left panel of figure 5 plots the forecast of the job-filling rate generated using the generalized measure of effective searchers versus the forecast generated using the standard measure of unemployment. This panel uses my estimated matching function coefficients, and as reported above, it reveals a correlation between these measures of 0.991. The middle panel turns to a similar plot, but this time the forecasts are based on the matching function coefficients imposed by the authors. The correlation remains very high. This suggests that the high correlation of the forecasts generated by the two competing measures of labor market slack is not driven by my preference for estimating the matching function coefficients rather than imposing them, as this high correlation between these two sets of forecasts is a feature even in the authors' preferred measures. The third panel shows the source of this similarity, plotting the basic measure of labor market tightness, V_t/U_t, when estimated using the two alternative measures of U_t. The measure of slack based on the generalized measure of effective searchers, V_t/U_t^s, is remarkably similar to that generated using the standard measure of unemployment, V_t/U_t^s.

This, in turn, implies that the differences in the forecasts of these two measures aren't driven by one measure being more informative than the other. Both rise and fall almost in tandem. Rather, they are driven by difference in the scale and location of the two measures of labor market slack. Indeed, this is the key to reconciling the evidence in figure 1 which shows that the authors' forecasts of the job-filling rate based on their generalized

Figure 5. Comparing Generalized and Standard Measures



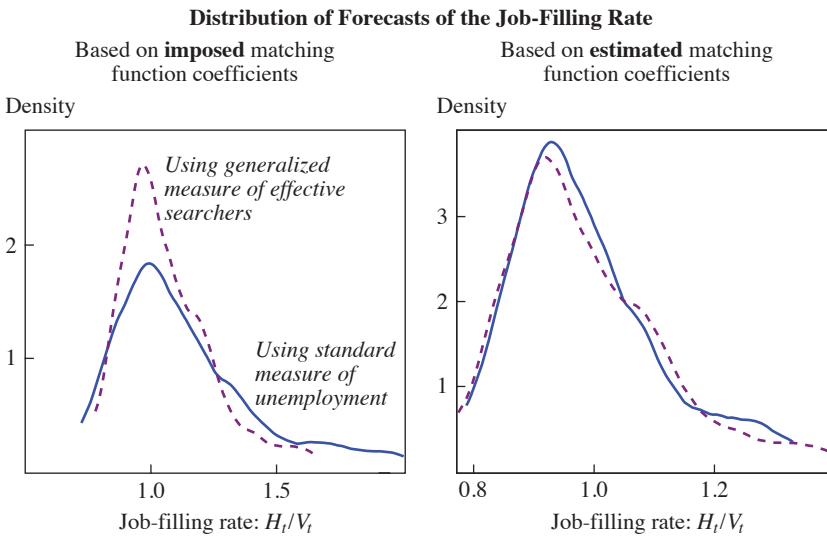
Source: Author's calculations.

measure of effective searchers is quite different from their forecasts based on the standard measure of unemployment, with the evidence in figure 5 that these two sets of forecasts are still very highly correlated. As figure 1 shows, both measures rise and fall in lockstep, but using the coefficients imposed by the authors yields a series of forecasts of the job-filling rate based on the standard measure that are typically located higher than those based on the generalized measure, and they are scaled so that they are more variable.

The location and scale of each series is determined by the coefficients α and μ. Thus, nearly *all* of the difference in the forecasts generated by these series reflects the choices of these auxiliary parameters. Remarkably little of the difference is due to the generalized measure bringing extra information about the state of the labor market.

Figure 6 illustrates the role that the matching function coefficients play in determining the scale and location of each series. It should be read in concert with the right panel of figure 5, which illustrates that the inputs into these forecasts are remarkably highly correlated. Yet as the left panel of figure 6 shows, the matching function coefficients imposed by the authors yields a series of forecasts based on the standard measure of unemployment

Figure 6. Scale and Location of Alternative Forecasts of the Job-Filling Rate



Source: Author's calculations.

that is both larger on average and much more variable than the forecasts based on the generalized measure of effective searchers. The right panel shows that these differences in scale and location are no longer meaningful when using estimated matching function coefficients.

DO THE IMPOSED COEFFICIENTS MAKE SENSE? At this point, it is worth assessing whether there is any reason to prefer the tests based on the imposed matching function coefficients presented in the paper versus the estimated coefficients emphasized here. On this score, there are two key points to make, one methodological and the other quantitative.

First, the methodological point yields a prima facie case that the matching function coefficients imposed in the paper are not appropriate for this sort of analysis. The exercise the paper performs is effectively a forecast evaluation exercise. But the source of the imposed α's comes from an instrumental variables regression, designed to recover the "structural" matching function coefficient. While structural coefficients are useful for constructing policy counterfactuals, in the presence of measurement error (and this paper is motivated by the idea that there is measurement error in our measures of labor market slack) a structural coefficient will not generally yield an efficient or unbiased forecast. (Recall, while ordinary least squares may yield biased coefficients, those coefficients are still the best

linear unbiased estimates.) Indeed, figure 3 illustrates that the imposed coefficients yield forecasts of the job-filling rate that are both biased and inefficient. Even if policymakers were to rely on the generalized measure of effective searchers, it is hard to believe that they would rely on these imposed coefficients to evaluate the likely implications for the job-finding rate.

To further muddy the methodological waters, the μ coefficients that the authors impose are neither structural nor reduced form, but the result of an arbitrary normalization in which it is essentially assumed that the unemployment and vacancy data for 2006 determine the location of the Beveridge curve. As far as I can tell, the authors never defend this choice, but it turns out not to be an innocuous normalization, as it also helps determine how strongly the job-filling rate responds to changes in labor market tightness.

Second, there is a quantitative point to be made. The authors are correct to argue that they need to use different estimates of α and μ when generating forecasts based on the two different measures of slack. Different estimates are needed because the generalized measure of effective searchers is scaled in a way that does not permit direct comparison to the standard measure, which is the level of unemployment. (The generalized measure is the sum over twenty-two population groups of the product of their share in the population at that point in time and their relative job-finding rate in 2006. Its precise scale depends on the normalization of these job-finding rates, or γ_i, in the authors' terminology.) As such, it makes sense to use different estimates of α_s and α_g to generate forecasts based on either the standard or generalized measures of labor market tightness.

But while the precise numeric values of α_s and α_g used to generate each forecast may differ, for this to be an apples-to-apples comparison, they should have a similar *economic* interpretation. They should each predict a similar response of the job-filling rate (H_t/V_t) to business cycle fluctuations in labor market tightness ($\theta_t = V_t/U_t$). After all, only by ensuring that both comparisons embed a similar sensitivity to the state of the business cycle can we ensure that any observed differences are due to the extra information embedded in one measure rather than the other.

To assess this, start by taking the first derivative of the earlier job-filling rate equation to obtain:

$$\frac{d\frac{H}{V}}{d\theta} = -\alpha\mu\theta^{-(\alpha+1)}.$$

[128.239.126.20] Project MUSE (2026-08-10 18:21 GMT) College of William & Mary & (Viva)

Next, recall that the measure of labor market tightness based on the standard measure of unemployment ($\theta_t^s = V_t/U_t^s$) is scaled differently from that based on the generalized measure of effective searchers ($\theta_t^g = V_t/U_t^g$). To be precise, the former has a mean of 0.58 and a standard deviation of 0.26, while the latter has a mean of 0.26 and a standard deviation of 0.71. So rather than comparing the responsiveness of the job-filling rate to a one-unit change in each measure, it would be more of an apples-to-apples comparison to evaluate how the job-filling rate responds to a one standard deviation change in each measure of labor market tightness, σ_θ. I call this the "cyclical sensitivity of the job-filling rate" and evaluate it at the mean level of tightness, $\bar{\theta}$, as follows:

$$\textit{Cyclical sensitivity of job-filling rate} = \sigma_\theta \times \left.\frac{d\frac{H}{V}}{d\theta}\right|_{\theta=\bar{\theta}} = -\sigma_\theta \alpha \bar{\theta}^{-(\alpha+1)}.$$

Table 2 shows how I calculate the cyclical sensitivity of the job-filling rate. The first four columns show the inputs into my calculations, and the final column shows the calculated sensitivity.

Importantly, notice that in panel A, the estimated sensitivity of the forecasts generated using the standard measure of unemployment is quite different from that calculated using the generalized measure of search activity. The point is that the relationship between job-filling, workers' search, and firm search used to compare these two measures is quite different in each case. And that in turn is the sense in which the comparisons calculated by the authors are not apples-to-apples comparisons.

By contrast, the panel B shifts the analysis to using the estimated matching function coefficients that I focus on in this comment. While the precise α's are quite different (that is α_s is different from α_g—reflecting differences in the scaling of the two competing measures of worker search, their implications), if instead we evaluate their consequences in terms of the effect of a one standard deviation change in search, the effects are quite similar. This suggests that the comparison based on the estimated coefficients is more of an apples-to-apples comparison. That comparison, in turn, found that the forecast of the job-filling rate based on the generalized measure of search effectiveness was almost identical to the forecast based on the standard measure of unemployment.

Finally, I should add that I'm not quite sure why the process by which the authors arrived at the α coefficients they impose are so different when

Table 2. Evaluation of the Cyclical Sensitivity of the Job-Filling Rate

	α *coefficient*	μ *coefficient*	*Average tightness* $\bar{\theta} = \overline{(V/U)}$	*Standard deviation of tightness* σ_θ	*Cyclical sensitivity of job-filling rate* $= -\sigma_\theta \alpha \bar{\theta}^{-(\alpha+1)}$
Panel A: Generating forecasts using imposed coefficients					
Using standard measure of unemployment	0.49	0.80	0.58	0.26	–0.23
Using generalized measure of search activity	0.57	0.48	0.26	0.71	–0.16
Panel B: Generating forecasts using estimated coefficients					
Using standard measure of unemployment	0.25	0.83	0.58	0.26	–0.10
Using generalized measure of search activity	0.44	0.53	0.26	0.71	–0.11

Source: Author's calculations.

using a measure of worker search based on the standard versus the generalized measure. The estimation is relatively opaque, and it is surely worth exploring why highly correlated series yield such different estimates of the cyclical sensitivity of the job-filling rate.

CONCLUSION This comment has dived pretty deep into the weeds, so it's worth panning back to the big picture. Abraham, Haltiwanger, and Rendell propose a new measure of labor market search that at a conceptual level has a lot to recommend it. The reality, however, is that it yields a measure that is remarkably highly correlated with a more standard measure like unemployment. The particular empirical exercise carried out in the paper asks which measure does a better job at predicting the job-filling rate. (It also asks which does a better job in matching the job-finding rate. For the sake of space, I've not dug into that measure here, but similar conceptual issues arise in analyzing competing predictions of that measure.)

Of course, one needs not just a measure of labor market slack but also a model if one is to predict the job-filling rate. And any model brings with it a set of auxiliary assumptions This comment has argued that these auxiliary assumptions are responsible for much of the difference in the forecasts generated by the two competing measures of labor market slack. My alternative set of tests are based on a different—and I would argue simpler, more

transparent, more statistically coherent, and economically meaningful—set of auxiliary assumptions. And my alternative tests reveal that both measures of labor market slack yield almost identical forecasts of the job-filling rate. As such, there is little evidence that the new generalized measure of active searchers includes much useful information that's not already embedded in a more standard measure like unemployment. The new measure is neither meaningfully better nor worse than a more standard alternative.

GENERAL DISCUSSION Olivier Blanchard noted that the authors clearly made progress on the matching function but he was struck by the lack of discussion around the Phillips curve. He also commented that the analysis directly speaks to the paper in this volume by Del Negro, Lenza, Primiceri, and Tambalotti, which finds that wages are less responsive to measures of slack.[1] However, Abraham, Haltiwanger, and Rendell's paper suggests that slack is likely mismeasured. Blanchard continued by noting the importance of capturing the heterogeneity of workers in the matching function. It may also be that wage bargaining associated with these different groups of searchers is not the same, resulting in a more complex relationship between slack and wage setting.

Giorgio Primiceri commented that to understand if this new measure of slack performs better with the standard Phillips curve correlation, one needs to examine data from the 1960s to the 1980s. Primiceri continued by noting that Ayşegül Şahin's discussion showed that the movement of the authors' new measure of slack across the business cycle is likely very similar before and after the 1990s. So it is unlikely that this new measure of slack would solve the puzzle of the post-1990s Phillips curve correlation.

James Stock said he is unsure about the paper's argument that the Beveridge curve is more stable and not shifting up with this new measure of slack. While roughly vertical in the way it is graphed, it is not obvious whether it has less instability and appears to shift to the right during the relevant period.

Ricardo Reis wondered if the authors could relate their paper to the Stansbury and Summers paper in this volume, which finds a decline in worker bargaining power.[2] In a strict Diamond-Mortensen-Pissarides model,

1. Marco Del Negro, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti, "What's Up with the Phillips Curve?," *Brookings Papers on Economic Activity*, Spring (2020).

2. Anna Stansbury and Lawrence H. Summers, "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy," *Brookings Papers on Economic Activity*, Spring (2020).

one would think of the Beveridge curve separately from the bargaining of the surplus.[3] However, in a McCall model of the labor market where a change in the bargaining power of workers can impact their search decisions, one may find, as a result, a different relationship between the amount of vacancies and the amount of job seekers.[4]

Katharine Abraham and John Haltiwanger made several points in response to Justin Wolfer's discussion of the relative performance of forecasts based on generalized versus standard measures of labor market tightness. (Further details regarding the authors' response to these comments can be found in online appendix D.) Abraham and Haltiwanger observe, first, that Wolfers's analysis does not recognize the trade-off in how well the standard model tracks the job-filling rate versus how well it tracks the job-finding rate among the unemployed. Using Wolfers's ordinary least squares (OLS) estimates of the matching function elasticity, the standard model performs almost as well as the generalized model with regard to the job-filling rate. But, using those same elasticity estimates, the relative performance of the standard model for tracking the job-finding rate among the unemployed is significantly *worse*—applying Wolfers's OLS estimates to tracking the job-finding rate among the unemployed yields an RMSE for the standard model that is three times as large as that for the generalized model.

Another issue Abraham and Haltiwanger note with Wolfers's analysis is that it uses the same data to estimate model parameters and then evaluates their performance through the lens of a forecasting perspective. As discussed in more detail in online appendix D, when Wolfers's approach is modified so that models are fit with data through 2007 and then used to forecast outcomes in later years, the result is a generalized matching function that significantly outperforms the standard matching function both for predicting the job-filling rate and for predicting the job-finding rate among the unemployed. Finally, Abraham and Haltiwanger argue, the OLS approach Wolfers suggests for estimating the elasticity of the matching function yields estimates that are inherently biased. This is not true of the instrumental variables (IV) estimates developed in the paper.

3. Peter A. Diamond, "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies* 49, no. 2 (1982): 217–27; Dale T. Mortensen, "The Matching Process as a Noncooperative Bargaining Game," in *The Economics of Information and Uncertainty*, ed. J. J. McCall (Chicago: University of Chicago Press, 1982); Christopher A. Pissarides, "Short-Run Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review* 75, no. 4 (1985): 676–90.

4. John Joseph McCall. "Economics of Information and Job Search," *Quarterly Journal of Economics* 84, no. 1 (1970): 113–26.

Using the IV estimates, the generalized model again significantly outperforms the standard model in tracking both the job-filling and the job-finding rates.

Haltiwanger addressed Şahin's and Blanchard's comments noting no incorporation of the Phillips curve in the analysis. Haltiwanger said that the paper produced an alternative measure to better capture slack in the labor market. Hopefully, there is consensus that the standard matching function is mis-specified. According to the data, a large portion of job seekers come out of the labor force and from among the employed, which is not captured by the standard matching function. Given this fact, their paper builds an alternative generalized matching function where the heterogeneity of job seekers is captured. Haltiwanger continued by saying that they are sympathetic to working out the details around their new measure and taking the next step to incorporate it into the Phillips curve framework. The new measure may have important implications for wage and price dynamics that should be explored in the future.

Haltiwanger responded to Reis's question by saying that if one were to pursue a Diamond-Mortensen-Pissarides model in the context of the Phillips curve, one would need to take a generalized approach like they do in the paper.

Haltiwanger continued by addressing Stock's question and part of Şahin's comment around normalizing measures. The authors normalized their measures because they think that the intercept of the hiring function should be different across specifications. The standard model should not have the same intercept as the generalized model. The authors are interested in removing those intercept differences, so they normalized their measures.

Haltiwanger then turned to another part of Şahin's comment around the number of subgroups needed to calculate a generalized matching function. He acknowledged that it's not clear whether one needs twenty-two groups for the sample period covered by the paper. The most important subgroups in this sample period are decomposing the unemployed into short-term and long-term unemployed and including and decomposing the group of those out of the labor force into want a job and other. However, over a longer sample period (e.g., including the entire post-WWII period), it is likely that additional subcategories matter. In considering this issue, it is noteworthy that there are large differences in relative search intensities across the detailed groups but for this to matter there also need to be changes in composition across the groups over time.

STAFF
REPORTS

NO. 1063
JUNE 2023

Measuring the Natural Rate of Interest after COVID-19

Kathryn Holston | Thomas Laubach | John C. Williams

FEDERAL RESERVE BANK *of* NEW YORK

Measuring the Natural Rate of Interest after COVID-19

Kathryn Holston Thomas Laubach, and John C. Williams

Federal Reserve Bank of New York Staff Reports, no. 1063
June 2023
JEL classification: C32, E43, E52, O40

Abstract

We modify the Laubach-Williams and Holston-Laubach-Williams models of the natural rate of interest to account for time-varying volatility and a persistent COVID supply shock during the pandemic. Resulting estimates of the natural rate of interest in the United States, Canada, and the Euro Area at the end of 2022 are close to their respective levels estimated directly before the pandemic; that is, we do not find evidence that the era of historically low estimated natural rates of interest has ended. In contrast, estimates of the natural rate of output have declined relative to those projected before the pandemic.

Key words: natural rate of output, time-varying volatility, Kalman filter, trend growth, COVID-19 pandemic

---

Williams: Federal Reserve Bank of New York (email: john.c.williams@ny.frb.org). Holston: Harvard University (email: kathrynaholston@gmail.com). This paper builds upon and extends the work of the three authors that was initially contained in the note "Adapting the Laubach and Williams and Holston, Laubach, and Williams Models to the COVID-19 Pandemic," May 27, 2020. The authors thank participants at the Thomas Laubach Research Conference for valuable comments. They also thank Logan Casey for outstanding research assistance.

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Open Market Committee, or the Federal Reserve System. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1063.html.

1 Introduction

The downward trend in estimates of the natural rate of interest to historically low levels observed in many countries has garnered considerable attention and debate about their sources and consequences (Laubach and Williams, 2016, Gourinchas and Rey, 2019, Rachel and Summers, 2019).[1] The events of the past few years, including the COVID-19 pandemic and subsequent policy actions, have renewed the debate over whether historically low natural rates of interest will persist in the post-pandemic era (International Monetary Fund, 2023, Obstfeld, 2023). Answering this question using empirical models of the natural rate of interest has been challenging owing to the unprecedented macroeconomic volatility across the globe during the pandemic. This paper develops and implements a data-driven approach that addresses the extraordinary effects of the pandemic using the Holston, Laubach, and Williams (HLW, 2017) and Laubach and Williams (LW, 2003) models of the natural rate of interest. Our approach preserves to the greatest extent the basic structure and flexibility of the original models, while providing consistent model estimates of natural rates before, during, and after the pandemic period. It also has broader application to models that estimate latent variables using frequentist (as in HLW and LW) and Bayesian methods.

The HLW and LW models apply the Kalman filter to translate movements in real GDP, inflation, and short-term interest rates into estimates of trend growth, the natural rate of output, and the natural rate of interest. The model's structure is flexible and incorporates transitory and permanent shocks to supply and demand and dynamic endogenous behavior of inflation and output. However, like other models that use the Kalman or other statistical filters, identifying assumptions regarding the nature of the shock processes are imposed. In particular, the shocks are assumed to be serially uncorrelated and described by time-invariant Gaussian distributions. The COVID-19 pandemic generated extraordinary swings in macroeconomic data that are dramatically at odds with both of these assumptions.

First, the assumption of a time-invariant Gaussian distribution for the shock processes is clearly contradicted by the data. Relative to the historical experience of the prior half century, the COVID-19 pandemic is an extreme tail event in terms of its

[1] There is a related literature on the history of real interest rates over a very long time span. See, for example, Rogoff, Rossi, and Schmelzing (2022) and references therein.

effects on economies around the world. Figure 1 plots the normalized quarterly log changes in GDP in the United States, Canada and the Euro Area over 2020-2022. Each series is normalized relative to their respective means and standard deviations from 1970-2019 (1972-2019 for the Euro Area). The fluctuations in GDP growth during the first three quarters of the pandemic (2020:Q2-2020:Q4) are enormous relative to pre-pandemic behavior. In addition, an unusually high number of observations in 2021 are outliers relative to historical behavior.

Second, the assumption of serially uncorrelated shocks is inconsistent with the highly negatively-correlated sequence of swings in output associated with the shutdowns and re-openings caused by COVID-19. For example, in all three economies, GDP declined sharply in the second quarter of 2020, then rose sharply in the following quarter. If uncorrected, these two stark violations of the assumptions significantly distort the estimation of model parameters and latent variables.

We make two modifications to the HLW and LW models that address the two violations of the original models' assumptions. First, we allow for time-varying volatility of the shocks to output and inflation during the pandemic period consistent with the appearance of extreme outliers in the data. We build on the insight from Lenza and Primiceri (2022) that if the timing of increased volatility is known, one can introduce time-varying volatility in the model directly by applying a scale factor to the innovation variances. Due to the presence of shocks to both observed and latent variables, we follow Harvey and Koopman (1992) and Harvey et al (1999) in measuring outliers in terms of auxiliary residuals. These have the advantage of providing a direct interpretation and test for outliers.

Second, we incorporate a proxy for a persistent, but not permanent, supply shock that is designed to capture the effects of COVID-19 and related policy responses. Specifically, we create a COVID-adjusted measure of the natural rate of output, where the magnitude of the adjustment is proportional to the country-specific COVID-19 Stringency Index from the Oxford COVID-19 Government Response Tracker (Hale et al., 2021). The magnitude of the effect of COVID policies on the natural rate of output is estimated for each economy. This COVID shock is in addition to the full set of innovations already present in the models.

The HLW model with these two modifications is estimated on data from the United States, Canada, and the Euro Area through 2022.[2] The estimation results

[2]In addition, we make some technical adjustments to the HLW model and the estimation pro-

demonstrate that these two modifications effectively address the two econometric issues associated with the pandemic.[3] The estimation procedure yields parameter estimates consistent with the model structure, and the pre-pandemic estimates of the natural rates of output and interest and the trend growth rate are very similar to those estimated on data ending in 2019.

The pattern of historically low estimates of trend GDP growth and the natural rate of interest experienced before the pandemic persist after the COVID-19 pandemic. In all three economies, the estimates of trend growth and the natural rate of interest in 2022 are within a few tenths of a percentage point of the corresponding estimates for 2019. In particular, these estimates provide no evidence of a reversal of the trend decline in estimates of the natural rate of interest based on data through 2022.

In all three economies, estimates of the COVID-adjusted natural rate of output in 2022 are significantly lower than what the model predicts based on pre-pandemic data. These declines reflect both the estimated effects of COVID-related restrictions and permanent negative shocks to the natural rate of output. According to the model, these declines in the natural rate of output are the most economically significant lasting effects of the COVID era.

This paper is organized as follows. Section 2 describes the HLW model and the evidence of significant departures from the model's assumptions brought on by the COVID-19 pandemic. Section 3 describes the modifications to the model to address the pandemic-related effects. Section 4 reports estimation results. Section 5 reports results from robustness exercises. Section 6 concludes.

2 Results from the Pre-Pandemic Model

In this section, we provide a short description of the original HLW model. We show that the data during the pandemic generate large outliers that are inconsistent with the assumptions of the model.

cedures for both models. Specifically, consistent with the specification of the LW model, we now estimate the relationship between trend growth and the natural rate of interest in the HLW model, instead of restricting it to unity. We also make minor technical adjustments that align the assumptions used in the various stages of the model estimation procedure, described in Appendix A1.

[3] In HLW (2017), the model was also estimated using data from the United Kingdom. Extending the sample to include the most recent years has weakened the estimated relationship between the output gap and real interest rates in the UK data, making estimates of the natural rate of interest highly unreliable. For that reason, we no longer estimate the model for the United Kingdom.

2.1 The Original HLW Model

In the HLW model, the natural rate of interest, r_t^*, is the real interest rate consistent with output equaling its natural rate, y_t^*, and stable inflation. As is standard in this literature (e.g., see Woodford, 2003), we model the output gap and inflation dynamics as a function of the real interest rate gap, $r_t - r_t^*$, using an intertemporal IS equation and Phillips curve relationship, in line with the New Keynesian framework:

$$\widetilde{y}_t = a_{y,1}\widetilde{y}_{t-1} + a_{y,2}\widetilde{y}_{t-2} + \frac{a_r}{2}\sum_{j=1}^{2}(r_{t-j} - r_{t-j}^*) + \epsilon_{\widetilde{y},t} \quad (1)$$

$$\pi_t = b_\pi \pi_{t-1} + (1 - b_\pi)\pi_{t-2,4} + b_y \widetilde{y}_{t-1} + \epsilon_{\pi,t} \quad (2)$$

The output gap is defined as $\widetilde{y}_t = 100 \cdot (y_t - y_t^*)$, where y_t and y_t^* are logarithms of real GDP and the unobserved natural rate of output, respectively, r_t is the real short-term interest rate, π_t denotes consumer price inflation, and $\pi_{t-2,4}$ is the average of the second to fourth lags of the inflation rate.[4] The stochastic disturbances $\epsilon_{\widetilde{y},t}$ and $\epsilon_{\pi,t}$ are transitory shocks to the output gap and inflation equations, respectively.

We use the Kalman filter to estimate the latent variables, which are the natural rate of output, its trend growth rate, and a process capturing other low-frequency determinants of the natural rate of interest. In keeping with the standard Kalman filter approach, the stochastic innovations to the measurement equations – the IS and Phillips curve equations – are assumed to follow a Gaussian distribution with standard deviations $\sigma_{\tilde{y}}$ and σ_π, respectively, and to be mutually and serially uncorrelated.

In contrast to the transitory shocks to the output gap and inflation equations, movements in r_t^* reflect highly persistent, or permanent, shifts in the relationship between the real short-term interest rate and the output gap (Williams, 2003). The law of motion for the natural rate of interest is given by

$$r_t^* = c \cdot g_t + z_t \quad (3)$$

where g_t is the trend growth rate of the natural rate of output, and z_t captures other determinants of r_t^*.[5] We specify the three latent variables in our state-space model

[4] See HLW (2017) Section 2 and Appendix A for details of the model specification. We take as a starting point the open-economy New Keynesian model specification as in Galí (2008) and relax two standard restrictions to work with reduced-form IS and Phillips curve equations.

[5] Note that, consistent with LW (2003), we relax the assumption in HLW (2017) of a one-for-one

as follows. The logarithm of the natural rate of output follows a random walk with a stochastic drift, g_t, that itself follows a random walk,

$$y_t^* = y_{t-1}^* + g_{t-1} + \epsilon_{y^*,t} \quad (4)$$

$$g_t = g_{t-1} + \epsilon_{g,t} \quad (5)$$

and the component z_t capturing other determinants of r_t^*, which is assumed to follow a random walk as well,

$$z_t = z_{t-1} + \epsilon_{z,t} \quad (6)$$

We assume that the disturbances $\epsilon_{y^*,t}$, $\epsilon_{g,t}$, and $\epsilon_{z,t}$ are normally distributed with standard deviations σ_{y^*}, σ_g, and σ_z, respectively, and are serially and contemporaneously uncorrelated with all other disturbances.

Equations 1 and 2 make up the measurement equations in our state-space model and can be expressed as

$$\mathbf{y}_t = \mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \xi_t + \epsilon_t \quad (7)$$

with stochastic innovations ϵ_t. Equations 4, 5, and 6 make up the state equations in our state-space model, written as

$$\xi_t = \mathbf{F} \cdot \xi_{t-1} + \eta_t \quad (8)$$

where ξ_t is the state vector of latent variables and η_t is the vector of stochastic innovations. See Appendix A1 for the full state-space representation of the model.

2.2 Outliers in Estimation with 2019:Q4 Model Parameters

We now analyze how the extreme movements in GDP and inflation during the COVID-19 pandemic yield large outliers in the standard HLW model. We then show that estimates of the latent variables are heavily affected by these sizable outliers, even when we constrain the model parameters at their pre-pandemic values, and demonstrate that modifications to the HLW and LW models are necessary.

relationship between trend growth and the natural rate of interest and estimate this relationship. See Appendix A1 for details on changes to the HLW (2017) model.

Before making any adjustments to the models, we begin by estimating the standard HLW model in Section 2.1 with data through 2019:Q4, prior to the onset of the COVID-19 pandemic.[6] For reference, the upper three panels of Figure 2 show the full samples of data. We fix the model parameters at their estimated values and re-estimate the latent variables through 2022:Q4 using the Kalman filter, taking all parameter values as given from the 2019:Q4 estimated model. We also fix the initial vector of unobserved states and its covariance matrix at the 2019:Q4 values.[7] This exercise is equivalent to dropping observations beginning in 2020:Q1 through the end of the sample during the maximum likelihood estimation of model parameters, while allowing the Kalman updating procedure to continue without modification through the end of the sample. In other words, we make no modifications to the state-space model, except that the model coefficient matrices and covariance matrices in the Kalman filtering procedure are fixed at their 2019:Q4 values. This would be a suitable approach if we take the view that the pandemic period is not informative for the model parameters, such as the slopes of the IS and Phillips curve equations, but *is informative* for the latent variables.

The final step of the Kalman filtering procedure to estimate the vector of unobserved state variables at time t (denoted as $\hat{\xi}_{t|t}$, and conditional on the information set at time t) is given by the Kalman updating equation,

$$\hat{\xi}_{t|t} = \hat{\xi}_{t|t-1} + \mathbf{K}_t \cdot \underbrace{(y_t - \mathbf{A}' \cdot x_t - \mathbf{H}' \cdot \hat{\xi}_{t|t-1})}_{\text{one-step-ahead prediction error}} \tag{9}$$

where $\hat{\xi}_{t|t-1}$ is the initial estimate of the state vector during the period, conditional on the information set at time $t-1$, and $\mathbf{K}_t$ is the Kalman gain matrix. The final term contains the one-step-ahead prediction errors (or forecast errors) corresponding to the measurement equations in the model. These one-step-ahead prediction errors are the residuals to the IS and Phillips curve equations, using the forecast of $\mathbf{y}_t$ (the vector of contemporaneous endogenous variables, that is, the output gap and inflation) based

[6] Throughout the paper, we use the current data vintage at the time of publication, regardless of the sample period.

[7] We store the estimated parameter vector θ from the final (stage 3) model as well as the signal-to-noise ratios λ_g and λ_z from the median unbiased estimation procedures following stages 1 and 2, respectively. See HLW (2017) for a description of the estimation procedure and footnote 6 for the initialization process of the vector of unobserved states, its conditional expectation $\xi_{1|0}$ in the first period, and the covariance matrix $P_{1|0}$.

on the data at time t and information at time $t-1$, corresponding to the state vector $\hat{\xi}_{t|t-1}$:

$$\mathbf{y}_t - \mathbb{E}[\mathbf{y}_t|\mathbf{x}_t, \zeta_{t-1}] = \mathbf{y}_t - (\mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \hat{\xi}_{t|t-1}) \tag{10}$$

In Equation 9, the Kalman gain matrix dictates the weight placed on the one-step-ahead prediction errors during the latent variable estimation. A larger Kalman gain $\mathbf{K}_t$ indicates that the final estimates of the latent variables are more heavily influenced by the gap between the realized data and the model's prediction, relative to $\hat{\xi}_{t|t-1}$, the prior estimate of ξ_t conditional on information in the preceding period.

In this initial exercise, the coefficient matrices $\mathbf{H}'$ (on the state vector $\hat{\xi}_{t|t-1}$) and $\mathbf{A}'$ (on the data $\mathbf{x}_t$) are fixed at their 2019:Q4 values. The resulting one-step-ahead prediction errors are very large during much of the pandemic period. These large forecast errors translate directly to the estimated vector of unobserved latent variables (y_t^*, g_t, z_t), so that the data during this period has large effects on estimates of these latent variables.

Because the HLW model is an unobserved-components model, there exists another set of model residuals in addition to the one-step-ahead prediction errors that are commonly used for diagnostic testing. These auxiliary residuals are smoothed estimators of the disturbances to the measurement equations, ϵ_t, and to the state equations, η_t, meaning that they incorporate all available information over the full sample period and provide a different interpretation of the stochastic innovations (Harvey and Koopman, 1992; Harvey et al. 1999). They have the advantage in applying a test for outliers: under the assumption that the stochastic innovations are from a Gaussian distribution, standardized auxiliary residuals greater than 2 (in absolute value) indicate either the presence of outliers or structural change.

We use the algorithm from Koopman and Durbin (2000) to obtain the standardized auxiliary residuals to the measurement equations, $\widetilde{\epsilon}_t/\hat{\sigma}_{\epsilon,t}$, and to the state equations, $\widetilde{\eta}_t/\hat{\sigma}_{\eta,t}$. The gold lines in Figure 3 shows that the standardized auxiliary residuals to the measurement equations, given by

$$\frac{\widetilde{\epsilon}_t}{\hat{\sigma}_{\epsilon,t}} = \frac{\mathbb{E}[\epsilon_t|\mathbf{y}_T, \mathbf{x}_T, \zeta_T]}{SD[\epsilon_t|\mathbf{y}_T, \mathbf{x}_T, \zeta_T]}, \tag{11}$$

indicate extreme outliers to the output gap equation in all economies in our sample, under the standard HLW model using 2019:Q4 parameter values. In the United States and Canada, standardized auxiliary residuals to the IS equation are 9 to 10 times the

outlier threshold for a Gaussian model in the second quarter of 2020, and 17 times the threshold in the Euro Area. While not as extreme, the model residuals also detect outliers to the inflation equation in each of the three economies, with standardized auxiliary residuals to the Phillips curve equation reaching double the outlier threshold in the United States.

Standardized auxiliary residuals to the unobserved state equations are given by

$$\frac{\widetilde{\eta}_t}{\hat{\sigma}_{\eta,t}} = \frac{\mathbb{E}[\eta_t|\mathbf{y}_T, \mathbf{x}_T, \zeta_T]}{SD[\eta_t|\mathbf{y}_T, \mathbf{x}_T, \zeta_T]} \tag{12}$$

As shown by the gold lines in Figure 4, these residuals demonstrate the presence of extreme outliers to the natural rate of output in the standard HLW model across all of the economies in our sample.

This observation is not unique to the HLW model. Figure 1 shows that GDP realizations during the pandemic are outliers with respect to historical data. Macroeconomic fluctuations of this magnitude would result in outliers in any standard macroeconomic model. As expected, even when pandemic-era GDP and inflation data are excluded during parameter estimation, using the Kalman filter with these extreme outliers present in the sample significantly distorts estimates of the latent variables during the pandemic period. The gold lines in Figure 5 displays large swings in the estimates of the latent variables using the model with 2019:Q4 parameter values. The extreme volatility in these estimates is in conflict with the specification of these latent variables as reflecting lower-frequency movements.

3 COVID-adjusted Model

The objective of this paper is to estimate the natural rate of interest following the COVID-19 pandemic in a way that is consistent with the HLW model outlined in Section 2.1. The large movements in economic activity and the persistent supply shocks during the pandemic period violate two standard, but important, model assumptions in HLW. As we show in Section 2.2, dropping the observations from this period during the maximum likelihood estimation of model parameters is insufficient to overcome these violated model assumptions, with model residuals indicating the presence of large outliers. The resulting estimates of the natural rates of output and interest are extremely volatile and inconsistent with our specification of r_t^* as a medium-run

concept that is driven by low-frequency movements.

This section details two adjustments to the HLW model that, taken together, address the violations to the original model. The first is the introduction of time-varying volatility in the model, which we implement by allowing the variances of the stochastic innovations to the output gap and inflation equations to be higher in the COVID era relative to the non-pandemic period. The second is the introduction of a persistent, but ultimately temporary, COVID supply shock, in addition to the transitory and permanent demand and supply shocks that are already present in the model. Each of these modifications alone is insufficient to overcome the estimation challenges posed by the pandemic, but in conjunction with each other they allow for continued estimation of the natural rate of interest.

Importantly, because we are estimating latent variables that are specified as random walks, simply dropping the observations during the pandemic period would not only understate the true uncertainty associated with this period, but would also necessitate interpolating these latent variables from their pre-pandemic values. Instead, rather than imposing that these variables do not change as a result of the COVID-19 pandemic, we are able to let the data inform the estimated natural rates of interest and output after the pandemic has abated. While the introduction of time-varying volatility has a similar effect on the latent variables as dropping the observations from 2020, our approach provides more flexibility in the later years of the pandemic and does not require the strong assumption that several years of data has no effect on the latent variables. Additionally, the HLW model with time-varying volatility but without serially correlated shocks to supply would constrain how the natural rate of output could evolve in response to the pandemic data. Modeling the persistent COVID supply shock is necessary in order to capture the effects of the pandemic on the natural rate of output.

3.1 Time-Varying Volatility during COVID-19

COVID-19 represents an extreme tail event relative to the assumption of Gaussian disturbances. As we show in Section 2.2, the resulting outliers contaminate estimates of the latent variables even when we exclude them from the estimation of model parameters. This is a statistical problem that is not unique to our model or to estimation of the natural rate of interest. We present a straightforward approach to

account for the substantial increase in volatility during this period by introducing time-varying volatility in the model during the window of time associated with the COVID-19 pandemic. We build on an insight from Lenza and Primiceri (2022): if the timing of increased volatility is known – as is the case for the COVID pandemic – we can introduce time-varying volatility in the model directly by applying a scale factor to the innovation variances during the period of increased volatility. We apply this insight to our unobserved-components model in order to estimate the natural rate of interest, but our approach can generalize to any state-space model with latent variables.

In particular, we introduce three new model parameters, κ_{2020}, κ_{2021}, and κ_{2022}. These are the variance scale parameters for 2020, 2021, and 2022, respectively, which multiply the variances of the innovations to the output gap and inflation equations. We define the vector κ_t of variance scale parameters at time t, that takes the values

$$\kappa_t = \begin{cases} \kappa_{2020} & \text{2020:Q2} \leq t \leq \text{2020:Q4} \\ \kappa_{2021} & \text{2021:Q1} \leq t \leq \text{2021:Q4} \\ \kappa_{2022} & \text{2022:Q1} \leq t \leq \text{2022:Q4} \\ 1 & \text{otherwise} \end{cases} \tag{13}$$

We estimate the three variance scale parameters by maximum likelihood together with the other model parameters, with the constraints $\kappa_{2020} \geq 1$, $\kappa_{2021} \geq 1$, and $\kappa_{2022} \geq 1$[8]. κ_t takes the value of 1 before the pandemic period and in 2023 and beyond. Section 5.1 considers alternative specifications of time-varying volatility.

The covariance matrix of the stochastic innovations to the output gap and inflation equations is now time-varying and is given by

$$\mathbf{R}_t = \kappa_t^2 \cdot \mathbf{R} = \begin{bmatrix} (\kappa_t \sigma_{\tilde{y}})^2 & 0 \\ 0 & (\kappa_t \sigma_\pi)^2 \end{bmatrix} \tag{14}$$

with time-varying innovation variances to the IS curve and Phillips curve equations of $(\kappa_t \sigma_{\tilde{y}})^2$ and $(\kappa_t \sigma_\pi)^2$, respectively.

[8]The restriction that $\kappa_t \geq 1$ is necessary to ensure that the likelihood estimation cannot *down-weight* the variance of certain observations, which would in effect allow it to place more weight on favorable observations. Instead, the estimated κ_t factors can only increase the innovation variances during the pandemic period.

Outside of the pandemic period, the innovation variances are specified exactly as in HLW (2017). That is, the innovation variances to the output gap equation, $\sigma_{\tilde{y}}^2$, and inflation equation, σ_{π}^2, are constant over the sample prior to 2020 and after 2022. During the year 2021, for example, the innovation variances take the values $(\kappa_{2021} \cdot \sigma_{\tilde{y}})^2$ and $(\kappa_{2021} \cdot \sigma_{\pi})^2$, respectively. Therefore κ_t is a ratio of the standard deviations of the disturbances to the measurement equations (the output gap and inflation equations) at time t relative to the standard deviations in the non-pandemic sample. When $\kappa_t > 1$, as we find for 2020 through 2022 in all economies in our sample, the innovation variances are greater than in the non-pandemic sample.

Introducing time-varying volatility to the stochastic innovations via the variance scale factors in κ_t has the effect of down-weighting extreme outlier observations in the maximum likelihood estimation of model parameters as well as in estimation of the latent variables via the Kalman filter. When $\kappa_t > 1$, the diagonal covariance matrix $\mathbf{R}_t$ of the disturbances to the output gap and inflation equations is larger relative to the case where $\kappa_t = 1$, and the resulting Kalman gain is smaller. As shown in Equation 9, the Kalman gain dictates the weight placed on the one-step-ahead prediction errors – the difference between realized values of the output gap and inflation in a given period and the model's predicted values based on information in the prior period – in updating the filtered estimates of the latent variables. As the innovation variances in a given period become large, the Kalman gain shrinks, so that the Kalman filter places relatively little weight on these new observations and the estimates of the latent variables in the state vector (that is, y_t^*, g_t, z_t and therefore r_t^*) remain close to the estimates from the prior period.

In the limit as the innovation variances tend toward infinity, the Kalman gain approaches zero, so that no weight is placed on the time-t observations in estimating the state vector. In effect, the model does not make use of time-t information, so that the forecast of the state vector at time t given the time-t information set is unchanged from the forecast given the information set at time $t-1$. This limiting case is equivalent to dropping the COVID-19 observations when estimating the latent variables. The same holds for parameter estimation: when κ_t is large, the model forecast error in this period is down-weighted when computing the log likelihood function, such that the data in this period have relatively little impact on the set of parameters that maximize the log likelihood function.

We see our approach as preferable to outright discarding the COVID-19 outliers

by treating them as missing data for several reasons. It is not our primary goal to provide estimates of r_t^* during the COVID-19 pandemic, and we treat these estimates with extreme caution. Rather, our objective is to deliver a framework for estimation of the natural rate of interest that is consistent with our approach in HLW, so that we are able to parse permanent changes to r_t^* from transitory shocks once the pandemic has abated. While the timing of the onset of the pandemic is clear, selecting an end date for the set of observations to discard would not be straightforward, and estimation of r_t^* may be sensitive to this choice. Additionally, a binary decision to drop or keep pandemic-related observations necessitates treating the entire period universally. Our approach is flexible in that we *allow* for increased volatility during the three years following the onset of the pandemic, but we do not *impose* higher innovation variances. We also do not impose any relationship between κ_{2020}, κ_{2021}, and κ_{2022}. By estimating these parameters together with the remaining model parameters, including the innovation variances during the non-pandemic period, our approach instead allows the data to inform the choice of variance scale factors, so that more extreme outliers are more heavily down-weighted. Indeed, when we allow the model to treat the later quarters of the pandemic differently from the earlier quarters, we find that it chooses to do so. Finally, it is well-known that estimates of the natural rate of interest are highly uncertain. Excluding data associated with the COVID-19 pandemic would understate the true uncertainty about future r_t^* estimates (Lenza and Primiceri, 2022). Our approach preserves estimation of the model standard errors over the full sample.

Lenza and Primiceri (2022) first implement this approach of applying scale factors to the covariance matrix of model forecast errors in order to introduce time-varying volatility within a fixed period to estimation of a vector autoregression. In the VAR setting, multiplying the innovations by a scale factor is equivalent to transforming the (observed) data directly by dividing by the scale factor. One can estimate the VAR with transformed data with no other modifications. In contrast, estimating a model with latent variables requires an additional modeling choice: treatment of the stochastic innovations to the state equations. To see this, note that the measurement equations in a model with latent variables, written in state-space form as Equation 7, include the vector of unobserved state variables ξ_t. Because these latent variables appear in the measurement equations, we cannot simply transform the observed variables without also transforming the latent variables. Transforming the

latent variables would have the effect of applying the same variance scale factor to the stochastic innovations to the latent variables. This embeds the assumption that the variances of the permanent shocks to the natural rate of output, its trend growth rate, and the other determinants of the natural rate of interest have increased in the same proportion as the transitory shocks to the output gap and inflation. Instead, we model the increase in shock volatility within the covariance matrix of the transitory shocks only, so that we do not introduce this assumption.

The choice to allow for increased volatility in the output gap and inflation innovations during the pandemic, but *not* increased volatility in the innovations to the latent variables, is consistent with the original HLW approach. We maintain the ability of the model to distinguish transitory shocks that do not affect the latent variables from low-frequency movements that do. The outsized movements in GDP and inflation during the early pandemic period, while extraordinarily large, were ultimately short-lived. Rather than explicitly modeling an increase in the volatility of the unobserved latent variables, we let the data speak by modeling an increase in the volatility of the transitory shocks only and leaving the specification of the latent variables unchanged. In the event that pandemic-induced movements in GDP and inflation are long-lasting, the Kalman filter will infer this accordingly and will ascribe permanent changes to r_t^* and y_t^*.

3.2 COVID-19 Supply Shock

The second feature of the COVID-19 pandemic that is at odds with the HLW model specification is the persistence of the associated supply shock. The HLW model incorporates transitory shocks to supply as innovations in the Phillips curve equation. Stochastic innovations to the IS and Phillips curve equations are assumed to be mutually and serially uncorrelated. Because the effects on supply of the sequence of shutdowns and re-openings associated with COVID-19 are highly serially correlated during the pandemic, they are not adequately captured by a sequence of serially uncorrelated transitory shocks. This is the case even once we have accounted for the increased variances of the stochastic innovations during the pandemic. We modify the model to incorporate a persistent, but ultimately temporary, COVID supply shock, in addition to the transitory and permanent demand and supply shocks already present in HLW.

The direct effects of COVID-19 on the economy are incorporated in the model as an adjustment to the natural rate of output in the output gap specification. We introduce one new variable, denoted d_t, as a proxy for the direct effects of the government restrictions and shutdowns implemented in response to the pandemic. We set this COVID indicator variable equal to the quarterly average of the COVID-19 Stringency Index from the Oxford COVID-19 Government Response Tracker (OxCGRT) for each country or region, as shown in the lower panel of Figure 2.[9] The stringency index, which ranges between 0 and 100 with larger numbers indicating stricter restrictions, aggregates measures of government containment and closure policies such as school and workplace closures, travel restrictions, bans or limits on public gatherings, and shutdowns of public transportation.

We choose this indicator because it is comprehensive and publicly available for all of the economies in our sample. We recognize that such an index of government responses cannot capture the full set of behavioral responses or compliance; nonetheless, it should provide a reasonable first-order approximation to the time-series properties of the direct effects of the pandemic and associated public health actions on economies. For the Euro Area, we use a GDP-weighted stringency index with 2019 GDP weights. As the OxCGRT project suspended data collection at the end of 2022, we assume each indicator variable declines linearly beginning in 2023:Q1, reaching zero in 2024:Q4. The COVID indicator is set equal to zero up to and including 2019:Q4 in all economies. It is assumed to be an exogenous variable.

We incorporate the COVID variable as an adjustment to the natural rate of output, within the output gap specification. In particular, the COVID-adjusted natural rate of output is given by

$$y^*_{t,COVID} = y^*_t + \frac{\phi}{100} d_t \tag{15}$$

where y^*_t is the standard natural rate of output, d_t is the COVID-19 indicator, and ϕ is an estimated parameter that translates the COVID variable d_t into effects on output. The output gap is correspondingly modified, with the COVID-adjusted natural rate of output replacing the standard natural rate of output:

$$\widetilde{y}_{t,COVID} = 100(y_t - y^*_{t,COVID}) = 100(y_t - y^*_t) - \phi d_t \tag{16}$$

[9]See Hale et al. (2021). We use the national weighted average of the stringency indices for vaccinated and unvaccinated populations, and our results are robust to using the index for vaccinated individuals only.

where y_t is the logarithm of real GDP. We estimate the parameter ϕ together with the remaining model parameters by maximum likelihood, including the variance-scaling parameters $\kappa_{2020}, \kappa_{2021}$, and κ_{2022}.

The COVID-adjusted output gap replaces the standard output gap in the measurement equations, which are the IS and Phillips curve equations. The state equations describing the law of motion for the unobserved state variables, including those for the natural rate of output and its trend growth rate, are unchanged.

4 Empirical Findings

In this section, we describe the main empirical findings using data through the end of the sample. We start with the estimation results and then analyze the behavior of key latent variables over the pandemic period. We estimate the COVID-adjusted HLW model for the United States, Canada, and the Euro Area. In previous work, we estimated the HLW model for the United Kingdom. Extending the sample to include pandemic-era data has weakened the estimated relationship between the output gap and the real interest rate gap in the UK data, making estimates of the natural rate of interest highly unreliable. For that reason, we no longer estimate the model for the United Kingdom. We present estimates from the COVID-adjusted LW model in Appendix A4.

4.1 Estimation Results

Table 1 reports the estimates of model parameters for the three economies. For comparison, the corresponding estimates from the model estimated through 2019:Q4 are reported in the Appendix Table A1. The parameter estimates are broadly similar to estimates from the pre-pandemic sample. Worth noting is that for Canada and the United States, the estimated values of λ_g (λ_z) are somewhat higher (lower) than in the pre-pandemic sample. The estimated values of the parameter c, linking the trend growth rate to the natural rate of interest, are close to unity (the value imposed in HLW 2017).

Figure 6 shows key estimated latent variables for the United States, Canada, and the Euro Area, respectively. The upper panel of each figure shows the time series of the estimated COVID-adjusted output gap and model-based real interest rate gap,

defined to be the model definition of the real interest rate less the estimated value of the natural rate of interest. The lower panel of each figure shows the estimates of the natural rate of interest and the trend growth rate, g_t. All three economies display the secular decline in estimates of the trend growth rate and the natural rate of interest highlighted in HLW (2017).

The model parameters introduced to address COVID-related effects on the economy are generally statistically significant at the 5 percent level. As discussed below, the estimated values of the parameter on the COVID shock variable, ϕ, have economically large effects. Consistent with the observations of enormous outliers early in the pandemic, the estimated values of the parameter κ_{2020} are sizable for 2020:Q2-Q4 and statistically significant. For the Euro Area and the United States, the estimated values of κ_{2021} are near 2 and κ_{2022} are around 1-1/2, and are generally statistically significant. For Canada, the estimated κ_{2021} is close to unity and κ_{2022} is around 1-1/2, and both are less precisely estimated than for the other economies.

The effects of the model modifications can be seen by comparing the standardized auxiliary residuals from the unadjusted model to those from the COVID-adjusted model described in Section 3. The blue lines in Figure 3 shows the standardized auxiliary residuals, described in Section 2.2, to the output gap and inflation equations, with the horizontal lines indicating two standard deviations. The gold lines show residuals from the version of the model with parameters held fixed at 2019:Q4 values. In all three economies, IS equation residuals in 2020 indicate extreme outliers to the model. The magnitude of the estimated κ_{2020} parameters is consistent with these extreme outliers. The blue lines show auxiliary residuals from the modified HLW model. With the two modifications, these residuals are of similar magnitude to the pre-pandemic period, and no longer indicate the presence of outliers. A similar, albeit less extreme, pattern is seen in the comparison of standardized auxiliary residuals from the inflation equation. The blue lines in Figure 4 displays the standardized auxiliary residuals from the equation for the natural rate of output. As with the output gap equation, the massive outliers in the standard version of the model are no longer present in the COVID-adjusted model, despite the fact that the model specifications of the natural rate of output and of the associated stochastic processes are unchanged.

4.2 The Effects of the COVID-19 period on g, r^*, and y^*

We now trace the evolution of key latent variables since the onset of the pandemic. This analysis reveals three key findings. First, the modified estimation procedure yields results that are overall quite similar to those from the original model during the pre-pandemic period. Second, the current estimates of the natural rate of interest are similar to those estimated directly before the pandemic. Third, the estimates of the natural rate of output at the end of 2022 are much lower than predicted before the pandemic.

Over the pre-pandemic sample, the estimates of the output gap, trend growth, and the natural rate of interest from the modified model are close to the corresponding estimates from the model without the COVID modifications (i.e. $\phi = 0, \kappa_t = 1$ for all periods). Figure 5 compares the estimates for the United States, Canada, and the Euro Area, respectively, from the modified model estimated through 2022:Q4 with estimates from the model with parameter values fixed their 2019:Q4 values. While the two sets of estimates are very similar through 2019, they differ sharply during the acute period of the pandemic, when the estimates from the model without COVID adjustments exhibit large swings due to the presence of extreme outliers.

The estimates of the trend growth rate of output in all three economies are slightly lower in 2022 than in 2019. Table 2 reports the estimates of the trend growth rate, g_t, in selected years as well as changes over sub-samples. It also reports forecasts of trend growth rates from various sources, which are generally little changed from 2019 to 2022. This represents a continuation of the pattern of declining estimates of trend growth observed over the preceding three decades, as seen in the columns in the table that report the changes in estimates over 1990-2007 and 2007-2019. Interestingly, for each economy, the 2022 estimates from the HLW model are reasonably close to outside forecasts.

In all three economies, the estimates of the natural rate of interest in 2022 are within a few tenths of a percentage point of the corresponding estimates in 2019. Table 3 reports annual averages of the estimates for selected years, along with changes over sub-samples. For comparison, it reports selected forecasts and market-based measures of longer-run real interest rate expectations. Note that market-based measures include term premia that can change over time, so they are not directly comparable to estimates of natural rates. Market-based measures of longer-run real interest rates rose from 2019 to 2022, but in each case, this brought them closer to the HLW

estimates.

The HLW estimates of r_t^* for Canada, the Euro Area, and the United States show no evidence of a quantitatively meaningful reversal of the decline in estimated natural rates of interest evident in prior decades. This result that the estimates of the natural rate of interest have not changed much since the start of the pandemic runs counter to some commentary that very large fiscal stimulus and rising levels of government debt, alongside evidence of large output gaps and high inflation, points to a higher level of the natural rate than before the pandemic. One reason why the HLW estimates of the natural rate of interest are little changed over the past three years is that lower estimates of trend growth work to offset any increase in the estimates of z.

The largest differences between model estimates pre- and post-pandemic relate to the level of each economy's natural rate of output. Figure 7 compares the model projections of the natural rate of output based on estimates using data through 2019:Q4 (the black lines) with estimates from 2022:Q4 (the blue lines). The figure also shows the effect of the COVID adjustment over this period (the difference between the blue and gold lines). As seen in the figure, the estimated shortfalls in the COVID-adjusted natural rate of output relative to pre-pandemic projections are economically sizable. For example, at the end of 2022, the COVID-adjusted level of the natural rate of output in the United States is about 4 percent below the pre-pandemic projection, with nearly half of that shortfall explained by the COVID shock measure and the remainder a permanent change in the natural rate of output.

5 Robustness

This section examines the sensitivity of the empirical findings to alternative specifications of the model. We consider two alternative specifications of time-varying volatility by adjusting the start and end dates over which we introduce scale factors to the variances of the stochastic innovations. Next, we estimate a version of the model without the COVID supply shock. Finally, we consider two alternative specifications for output. Our finding that estimates of the natural rate of interest remain low following the COVID-19 pandemic is robust to these alternative specifications.

5.1 Specification of Time-Varying Volatility

In order to account for the extreme GDP and inflation observations during the COVID-19 pandemic, we introduce time-varying volatility in the HLW model by applying a scale factor to the innovation variances of the model's measurement equations during the period of increased volatility associated with the pandemic, as described in Section 3. In our baseline specification, we allow the innovation variances to the output gap and inflation equations to increase during the years 2020, 2021, and 2022. The corresponding vector of variance scale parameters κ_t is given by Equation 13. Our baseline specification is informed by the observed data, the model one-step-ahead prediction errors, and the model auxiliary residuals. In this section, we consider two alternative specifications of κ_t. Because the onset of the pandemic occurred later within the first quarter of 2020 in the economies we study, the period of increased volatility begins in 2020:Q2 in our main specification. However, the auxiliary residuals to the IS equation, shown in Figure 3, indicate the presence of sizable outliers in 2020:Q1. Our first alternative specification includes 2020:Q1 when estimating and applying the variance scale parameters, with κ_t given by

$$\kappa_t = \begin{cases} \kappa_{2020} & \text{2020:Q1} \leq t \leq \text{2020:Q4} \\ \kappa_{2021} & \text{2021:Q1} \leq t \leq \text{2021:Q4} \\ \kappa_{2022} & \text{2022:Q1} \leq t \leq \text{2022:Q4} \\ 1 & \text{otherwise} \end{cases} \tag{17}$$

Although the timing of the pandemic onset is known, it is more difficult to determine the end date of the COVID-19 pandemic. In fact, many metrics indicate that the pandemic is ongoing at the time of publication. However, the GDP and inflation realizations as well as model auxiliary residuals in 2022 are more in line with the pre-pandemic sample. While our main specification allows for increased innovation variances in the year 2022 (and subsequently fixes $\kappa_t = 1$ in 2023 and beyond), our second alternative specification assumes that the innovation variances return to their

non-pandemic levels in 2022, with κ_t given by

$$\kappa_t = \begin{cases} \kappa_{2020} & \text{2020:Q2} \leq t \leq \text{2020:Q4} \\ \kappa_{2021} & \text{2021:Q1} \leq t \leq \text{2021:Q4} \\ 1 & \text{otherwise} \end{cases} \tag{18}$$

In all three economies, estimates of the latent natural rate of output under the alternative specifications of time-varying volatility are close to the baseline specification.[10] Estimates of y_t^* at the end of 2022 lie below the projected level of the natural rate of output based on 2019:Q4 data, regardless of the κ_t specification, with the gap between projected y_t^* as of 2019:Q4 and current model estimates generally widening over the past three years. Our finding that the pandemic had a persistently negative effect on the estimated natural rates of output is robust across specifications.

We find that the estimated coefficient ϕ, which translates the COVID indicator variable into effects on the output gap, is sensitive to the specification of κ_t. As a result, estimates of the COVID-adjusted natural rate of output are more sensitive, but key results hold across all three specifications. In all cases, the COVID-adjusted natural rate of output is below the standard natural rate of output during the pandemic period. In the United States and the Euro Area, the COVID-adjusted natural rate of output is significantly lower than the estimated natural rate of output in 2020 and 2021 and remains below y_t^* at the end of 2022. In Canada, $y_{t,COVID}^*$ lies below y_t^* in 2020 and 2021 in all three specifications, but the gap relative to the baseline is narrower and closes in 2022 as the COVID indicator variable declines. In the Euro Area, the COVID-adjusted natural rate of output remained below its 2019 level at the end of the sample, while the United States and Canada had recovered to their pre-pandemic levels. As of the sample end, the COVID-adjusted natural rate of output had not recovered to the pre-pandemic trend in any of the three economies.

Accordingly, the COVID-adjusted output gap deviates substantially from the baseline estimates in some cases. Despite a wider range of output gap estimates across the κ_t specifications, estimates of r_t^* and g_t are relatively close to the baseline estimates. In the United States, estimates of r_t^* under alternative specifications are within 50 basis points of the baseline specification at the sample end, well within

[10]Note that the equation for the natural rate of output is unchanged from the standard HLW model; the variance scale parameter κ_t does not enter the y_t^* equation.

the estimated standard error bands. Estimates in Canada are very close across all three specifications. Estimates of r_t^* and g_t in the Euro Area are more sensitive to the specification of κ_t throughout the sample, but remain within 40 basis points of the baseline specification at the sample end. We find that the COVID pandemic has a large effect on the estimated natural rates of output, though the magnitude of the effect varies across specification, while estimates of trend growth and the natural rate of interest are generally close to their pre-pandemic levels and robust to alternative specifications of time-varying volatility.

5.2 Model without COVID Supply Shock

We now consider a version of the model that includes time-varying volatility but does not include the COVID supply shock. This is equivalent to imposing $\phi = 0$ in the COVID-adjusted model. In all economies, we find that the output gap is much deeper and rebounds much more slowly in the specification without the COVID supply shock. However, in all economies, the level of the (unadjusted) output gap under this alternative specification reaches that of the baseline COVID-adjusted output gap by the end of 2022. Estimates of the natural rate of interest are largely similar to, or slightly below, the baseline model at the end of the sample, as shown in Figure 10. In the United States, estimates of trend growth and the natural rate of interest from the model without the COVID supply shock are lower than the baseline model during the pandemic period, with an estimate of r_t^* that is 50 basis points below the baseline estimate in final quarter of 2022. When we impose $\phi = 0$ in the model, we find that estimates of κ_{2020} are larger in all three economies.

5.3 Alternative Measures of Output

Finally, we consider two alternative specifications of output for the United States. Figure 11 displays estimates of the natural rate of interest under these alternative specifications. The first is a version of the model in which potential output is an exogenous variable, using estimates from the Congressional Budget Office as data. While the HLW model jointly estimates the natural rates of output and interest and models the trend growth rate as a determinant of r_t^*, there is no explicit role for trend GDP growth in this specification. The natural rate of interest follows a random walk and is the only latent variable in the model; since we do not model trend

growth as a determinant of r_t^* in this extension, there is no separate specification of other determinants z_t, and we estimate r_t^* directly via the Kalman filter. Because we take the CBO's measure of potential output as given, we do not introduce the COVID supply shock to the natural rate of output described in Section 3.2 (i.e. $\phi = 0$). Otherwise, we use the COVID-adjusted HLW model, which includes time-varying volatility during the pandemic period. The output gap based on estimates of potential output from the CBO lies below the estimated output gap from the HLW model, with substantial differences between the two measures emerging in 2008 and persisting through the end of the sample. Estimates of the natural rate of interest under this alternative specification, shown in Figure 11, are substantially negative and significantly lower than the baseline HLW estimates in the decade preceding the pandemic as well as during the pandemic period, with the gap widening at the end of the sample.

In the second alternative specification, output is defined as the average of real GDP and real Gross Domestic Income, the latter of which is an alternative measure of economic output constructed from income data and production costs by the Bureau of Economic Analysis. Switching to this alternative output measure has a minor effect on estimates of the latent variables, with r_t^* estimates that are about 50 basis points below the baseline HLW model (which uses real GDP data alone) from 2020 through 2022. In each of the alternative specifications we consider, estimates of the natural rate of interest lie below r_t^* estimates from the baseline HLW model for the entirety of the pandemic period. Across these alternative specifications, estimates of the natural rates of interest remain at historically low levels.

6 Conclusion

This paper develops and implements an approach to address the unusual effects of the pandemic on the economy and allows continued estimation of the natural rates of output and interest in the post-pandemic era using the Holston, Laubach, and Williams and Laubach and Williams models. These methods can be applied to a broad set of models that estimate latent variables over the pandemic period or other periods of large-scale disruptions to economic activity. According to the model estimates, the main longer-term consequence from the pandemic period is a reduction in the natural rate of output, but the imprint on the natural rate of interest appears to be relatively

modest. We do not find evidence that the era of historically low estimated natural rates of interest has come to an end. These findings are shown to be robust to a range of alternative model specifications.

7 References

Buncic, Daniel, "Econometric Issues with Laubach and Williams' Estimates of the Natural Rate of Interest," July 2021.

Buncic, Daniel, "On a Standard Method for Measuring the Natural Rate of Interest," April 2022.

Fagan, Gabriel, Jerome Henry, and Ricardo Mestre. 2001. "An Area-Wide Model (AWM) for the Euro Area," European Central Bank Working Paper Series No. 42, January.

Galí, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton, NJ: Princeton University Press.

Gourinchas, Pierre-Olivier and Hélène Rey. 2019, "Global real rates: a secular approach," BIS Working Papers, No 793, June.

Hale, Thomas, Noam Angrist, Rafael Goldszmidt, Beatriz Kira, Anna Petherick, Toby Phillips, Samuel Webster, Emily Cameron-Blake, Laura Hallas, Saptarshi Majumdar, and Helen Tatlow. 2021. "A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)." *Nature Human Behaviour.*

Hamilton, James D. 1986. "A Standard Error for the Estimated State Vector of a State-Space Model," *Journal of Econometrics,* 33, 387-397.

Harvey, Andrew C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter,* Cambridge: Cambridge University Press.

Harvey, Andrew C, and Siem Jan Koopman. 1992. "Diagnostic Checking of Unobserved-Components Time Series Models," *Journal of Business & Economic Statistics,* 10(4), 377-389.

Harvey, Andrew C, Siem Jan Koopman, and Jeremy Penzer. 1999. "Messy Time Series: A Unified Approach," in Fomby, T.B. and Carter Hill, R. (Eds.) *Messy Data (Advances in Econometrics)*, Vol. 13, Emerald Group Publishing Limited, Bingley, 103-143.

Holston, Kathryn, Thomas Laubach, and John C. Williams. 2017. "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, S59-S75.

International Monetary Fund. 2023. *World Economic Outlook: A Rocky Recovery,* Washington, DC, April.

Koopman, Siem Jan and James Durbin. 2000. "Fast Filtering and Smoothing for Multivariate State Space Models," *Journal of Time Series Analysis*, 21(3), 281-296.

Laubach, Thomas, and John C. Williams. 2003. "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, 85(4), November, 1063-1070.

Laubach, Thomas, and John C. Williams. 2016. "Measuring the Natural Rate of Interest Redux," *Business Economics*, 51, 257-267.

Lenza, Michele, and Giorgio E. Primiceri. 2022. "How to Estimate a Vector Autoregression After March 2020," *Journal of Applied Econometrics*, 37(4), 688-699.

Obstfeld, Maurice. 2023. "Natural and Neutral Real Interest Rates: Past and Future," Prepared for the 10th Asian Monetary Policy Forum, Singapore, May 25-26, 2023.

Rachel, Lukasz, and Lawrence H. Summers. 2019. "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity*, Spring 2019, 1–63.

Rogoff, Kenneth S., Barbara Rossi, and Paul Schmelzing. 2022. "Long-Run Trends in Long-Maturity Real Rates 1311-2021," NBER Working Paper No. 30475, October 2022.

Stock, James H. 1994. "Unit Roots, Structural Breaks, and Trends," in R. Engle and D. MacFadden (eds.) *Handbook of Econometrics*, Vol. 4, Amsterdam, Elsevier Science, 2739-2841.

Stock, James, and Mark Watson. 1998. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model," *Journal of the American Statistical Association*, 93, March, 349–358.

Wicksell, Knut. 1936. *Interest and Prices* (tr. of 1898 edition by R.F. Kahn). London: Macmillan.

Williams, John C. 2003. "The Natural Rate of Interest," *Federal Reserve Bank of San Francisco Economic Letter*, 2003-32, October 31, 2003.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

Table 1: Parameter Estimates

Parameter	United States	Canada	Euro Area
Sample	1961-2022	1961-2022	1972-2022
λ_g	0.073	0.061	0.039
λ_z	0.021	0.011	0.023
$\sum a_y$	0.936	0.946	0.945
a_r	−0.079	−0.077	−0.037
	(4.215)	(2.947)	(1.757)
b_y	0.073	0.049	0.088
	(3.003)	(1.821)	(2.310)
c	1.128	1.109	0.953
	(3.574)	(2.794)	(1.155)
ϕ	−0.085	−0.062	−0.135
	(2.199)	(3.442)	(6.557)
$\kappa_{2020Q2-Q4}$	9.033	9.377	18.609
	(2.351)	(2.134)	(2.384)
κ_{2021}	1.791	1.087	1.960
	(2.941)	(1.835)	(2.150)
κ_{2022}	1.676	1.613	1.557
	(2.060)	(1.293)	(1.231)
$\sigma_{\tilde{y}}$	0.452	0.579	0.318
σ_π	0.787	1.343	0.965
σ_{y^*}	0.500	0.448	0.382
σ_g	0.145	0.109	0.060
σ_z	0.118	0.085	0.201
$\sigma_{r^*} = \sqrt{c^2\sigma_g^2 + \sigma_z^2}$	0.202	0.148	0.209
S.E. (sample ave.)			
r^*	1.140	1.603	2.739
g	0.428	0.413	0.264
y^*	1.612	2.474	1.662
S.E. (final obs.)			
r^*	1.565	1.858	3.805
g	0.689	0.639	0.405
y^*	2.914	3.679	2.878

Notes: t statistics are in parentheses; σ_g is expressed at an annual rate.

Table 2: Trend Growth Estimates

					Change		
	1990	2007	2019	2022	1990-2007	2007-2019	2019-2022
United States							
HLW g estimates	3.3	2.8	2.1	1.8	−0.5	−0.7	−0.2
Consensus Forecasts	2.4	2.8	2.0	1.9	0.4	−0.8	−0.1
IMF World Economic Outlook	2.6	3.0	1.6	1.8	0.4	−1.5	0.2
Blue Chip Financial Forecasts	2.5	2.8	2.0	1.9	0.4	−0.8	−0.1
Canada							
HLW g estimates	3.4	2.5	1.8	1.5	−0.9	−0.7	−0.3
Consensus Forecasts	2.8	2.4	1.8	1.8	−0.4	−0.7	0.0
IMF World Economic Outlook	3.2	2.7	1.7	1.7	−0.5	−1.0	0.0
Euro Area							
HLW g estimates	2.7	2.0	1.2	1.0	−0.7	−0.9	−0.1
Consensus Forecasts	*n/a*	1.9	1.3	1.2	*n/a*	−0.6	−0.1
IMF World Economic Outlook	*n/a*	2.0	1.3	1.4	*n/a*	−0.7	0.1

Notes: Consensus Forecasts data are the mean of panelists' trend estimates of expected GDP growth in the subsequent 6 to 10 years. IMF estimates are the 5-year-ahead forecast for real GDP growth. Blue Chip estimates are the mean long-run forecasts for real GDP growth. For these forecasts, reported numbers are averages of the Spring and Fall publications. HLW estimates are annual averages. Numbers may not sum due to rounding.
Sources: Consensus Economics Inc, London; IMF World Economic Outlook; Blue Chip Financial Forecasts.

Table 3: Natural Rate of Interest Estimates

					Change		
	1990	2007	2019	2022	1990-2007	2007-2019	2019-2022
United States							
HLW r^* estimates	3.6	2.5	0.9	0.8	−1.1	−1.5	−0.2
TIPS yields (5-10 years ahead)	*n/a*	2.5	0.6	0.8	*n/a*	−1.9	0.2
Blue Chip Financial Forecasts	3.1	2.4	0.5	0.5	−0.7	−1.9	0.0
Canada							
HLW r^* estimates	3.5	2.7	1.7	1.5	−0.8	−1.0	−0.2
Real Return Bond yields	*n/a*	2.0	0.4	1.0	*n/a*	−1.5	0.5
Euro Area							
HLW r^* estimates	2.6	2.3	0.6	0.7	−0.3	−1.6	0.0
Swap-implied real yields (5-10 years ahead)	*n/a*	2.1	−0.7	0.0	*n/a*	−2.8	0.7

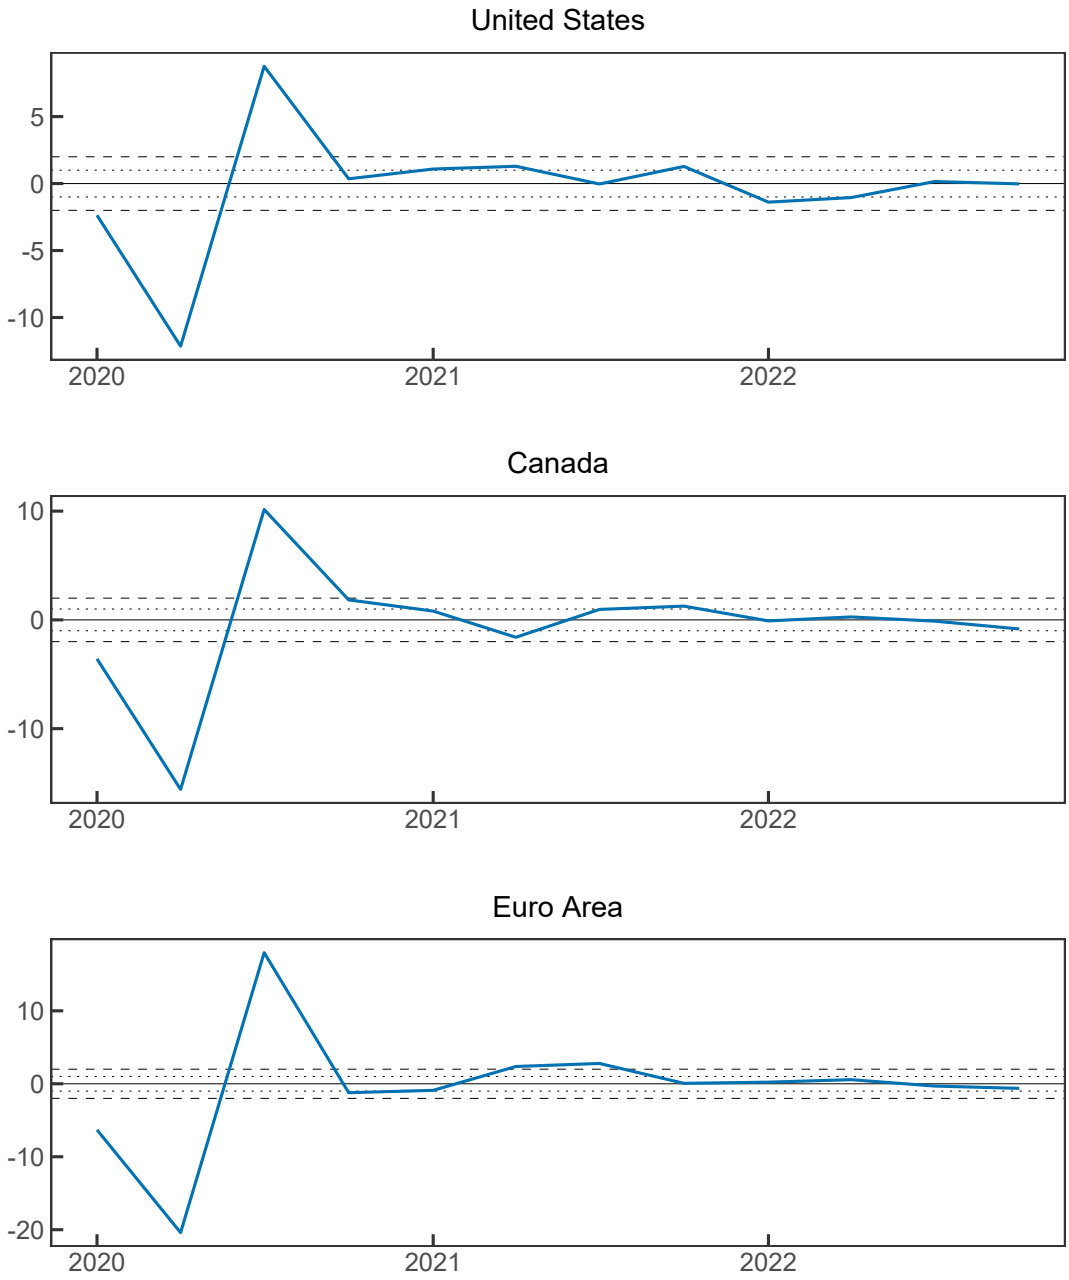
Notes: All numbers are annual averages. Blue Chip estimates are the difference in the mean long-years-ahead forecasts of the federal funds rate and CPI (average of Spring and Fall publications). Canadian Real return bond yields are long-term benchmark yields from the Bank of Canada. Euro Area yields are calculated using French nominal yields and data on zero-coupon inflation swaps linked to Euro Area HICP inflation. Numbers may not sum due to rounding.

Table 4: Robustness to κ_t Specification

	ϕ	κ_{2020}	κ_{2021}	κ_{2022}
United States				
Baseline	−0.09	9.03	1.79	1.68
With 2020:Q1	−0.16	5.37	1.76	2.28
$\kappa_{2022} = 1$	−0.05	10.64	1.83	1.00
Canada				
Baseline	−0.06	9.38	1.09	1.61
With 2020:Q1	−0.03	8.99	1.13	1.32
$\kappa_{2022} = 1$	−0.02	8.99	1.14	1.00
Euro Area				
Baseline	−0.13	18.61	1.96	1.56
With 2020:Q1	−0.11	17.60	1.99	1.32
$\kappa_{2022} = 1$	−0.13	18.09	1.97	1.00

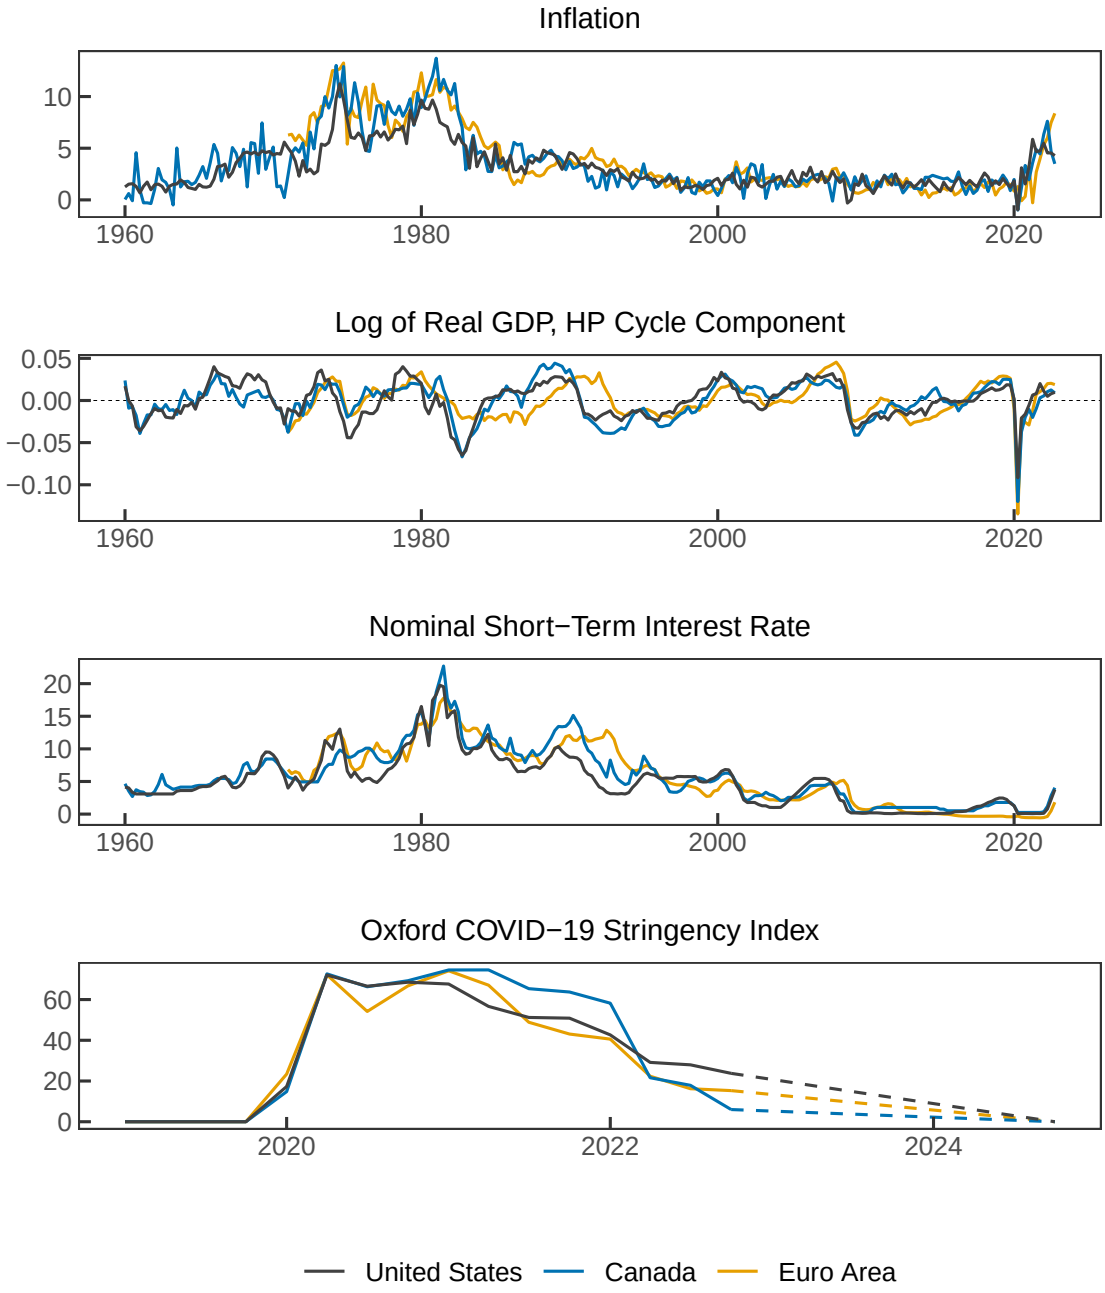
Notes: The rows labeled "with 2020:Q1" report parameter estimates from an alternative specification of the model in which κ_{2020} includes 2020:Q1. Rows labeled "$\kappa_{2022} = 1$" report parameter estimates from a specification of the model with variance scale parameters estimated for 2020 and 2021 only. "Baseline" parameter estimates are from the COVID-adjusted HLW model.

Figure 1: Normalized GDP Growth



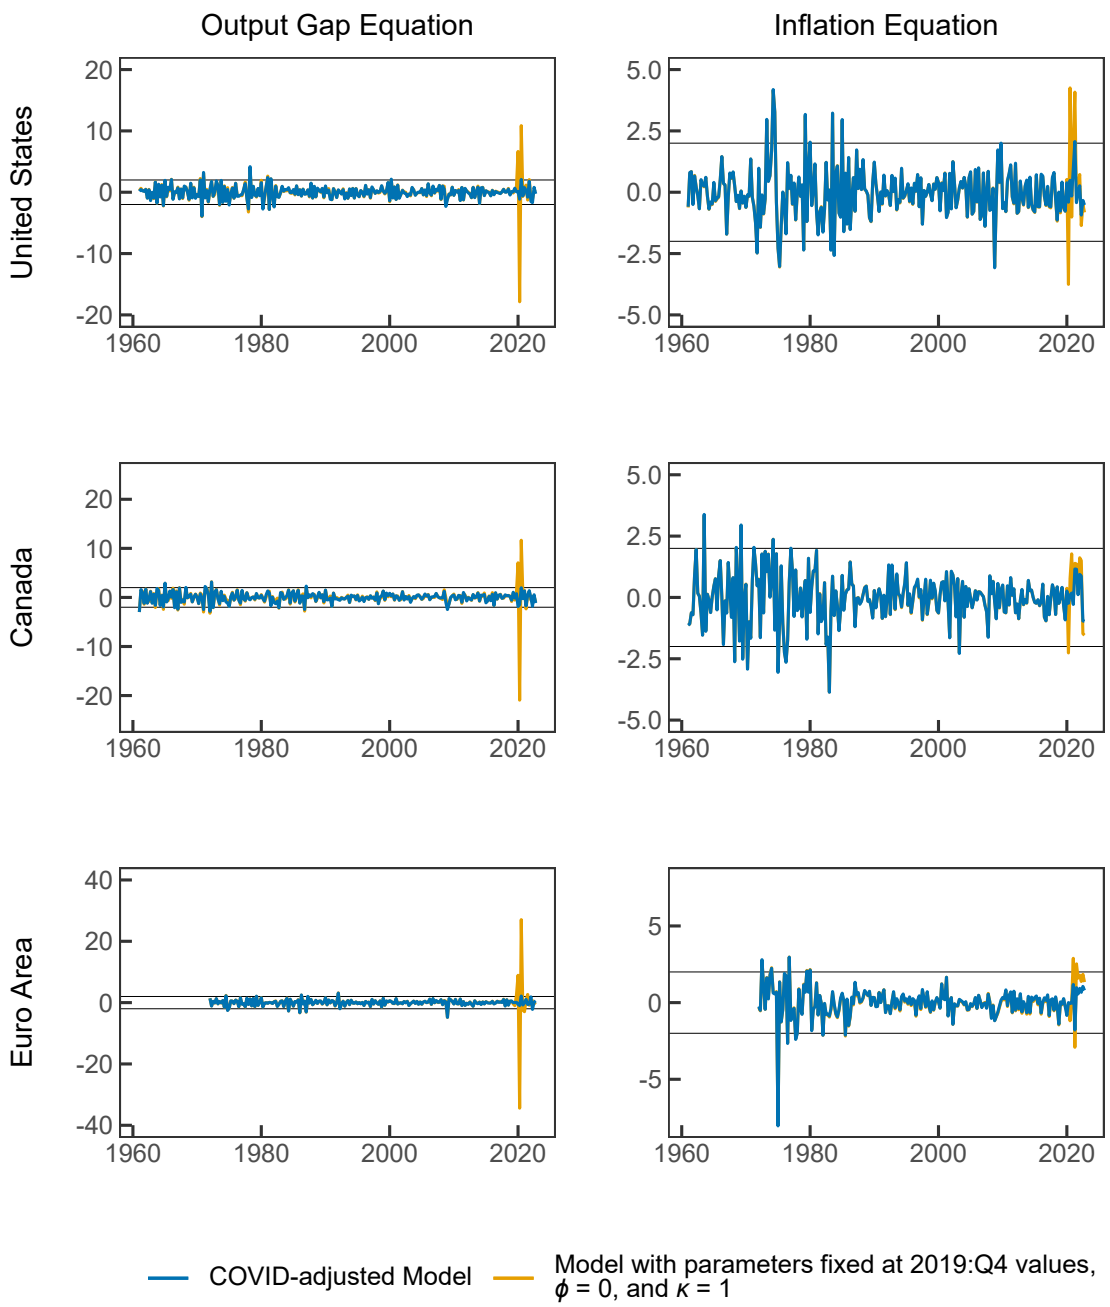
Notes: The chart shows the log-difference of quarterly real GDP less the pre-pandemic mean, divided by the pre-pandemic standard deviation. The pre-pandemic period refers to 1970-2019 for the United States and Canada, and 1972-2019 for the Euro Area. The solid horizontal line indicates zero, the dotted lines indicate the value one, and the dashed lines indicate the value two.

Figure 2: Comparison of Data Across Economies



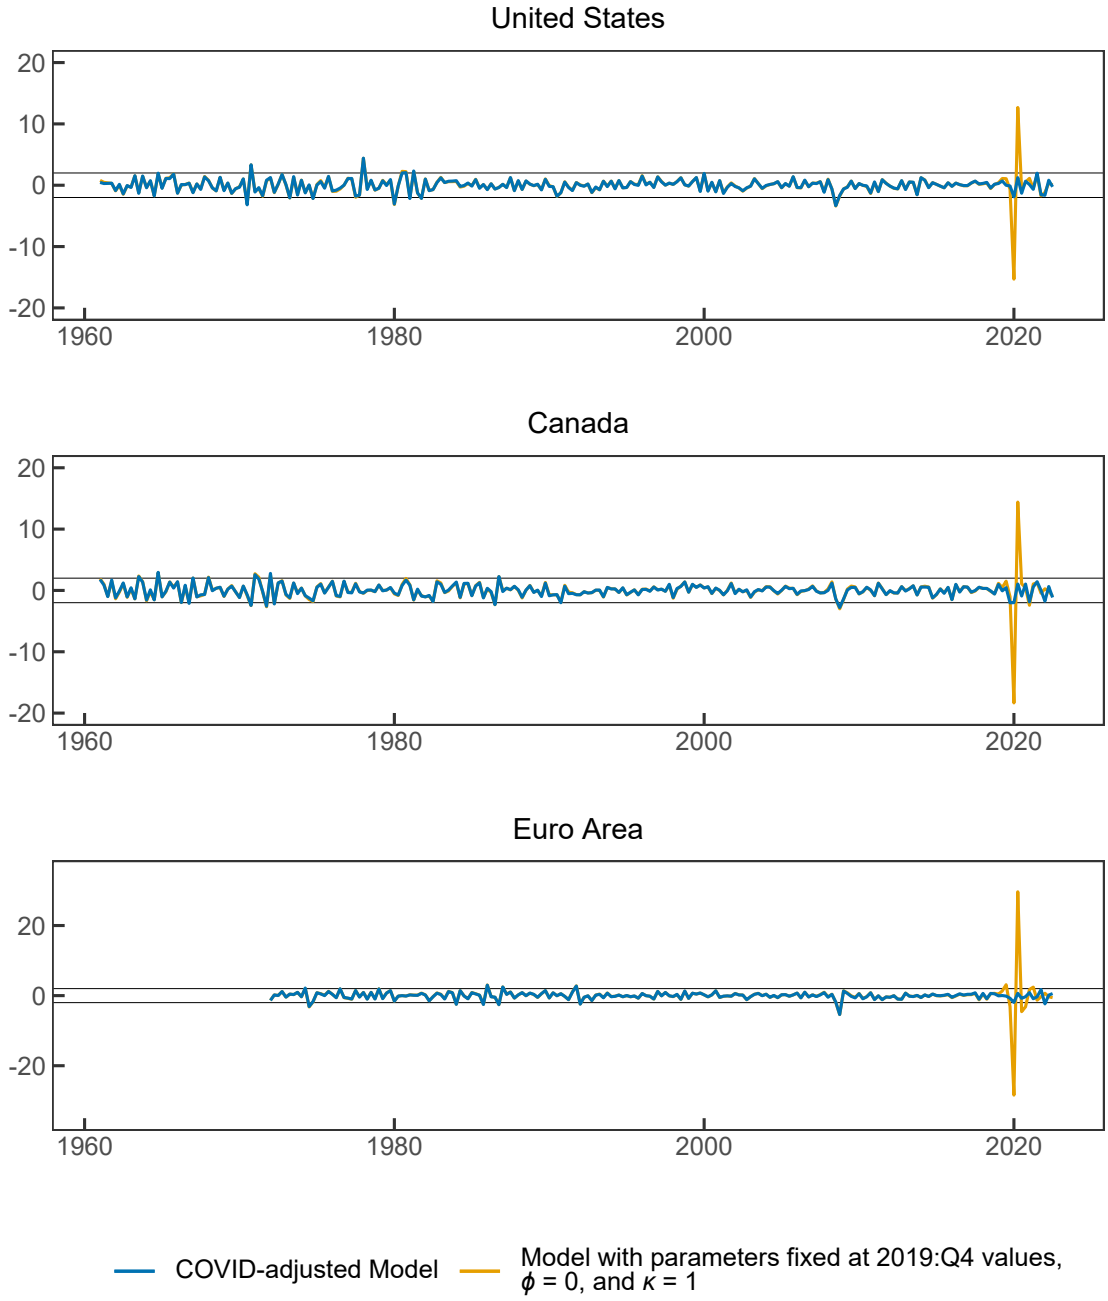
Notes: The upper three panels show the published data used in estimating the model. To ease comparison across economies, the second panel shows the cycle component of real GDP using an HP filter. The bottom panel shows the COVID index: the solid lines show the published values and the dashed lines show the assumed path over 2023-2024 when published numbers are no longer available.

Figure 3: Standardized Auxiliary Residuals: Output Gap and Inflation



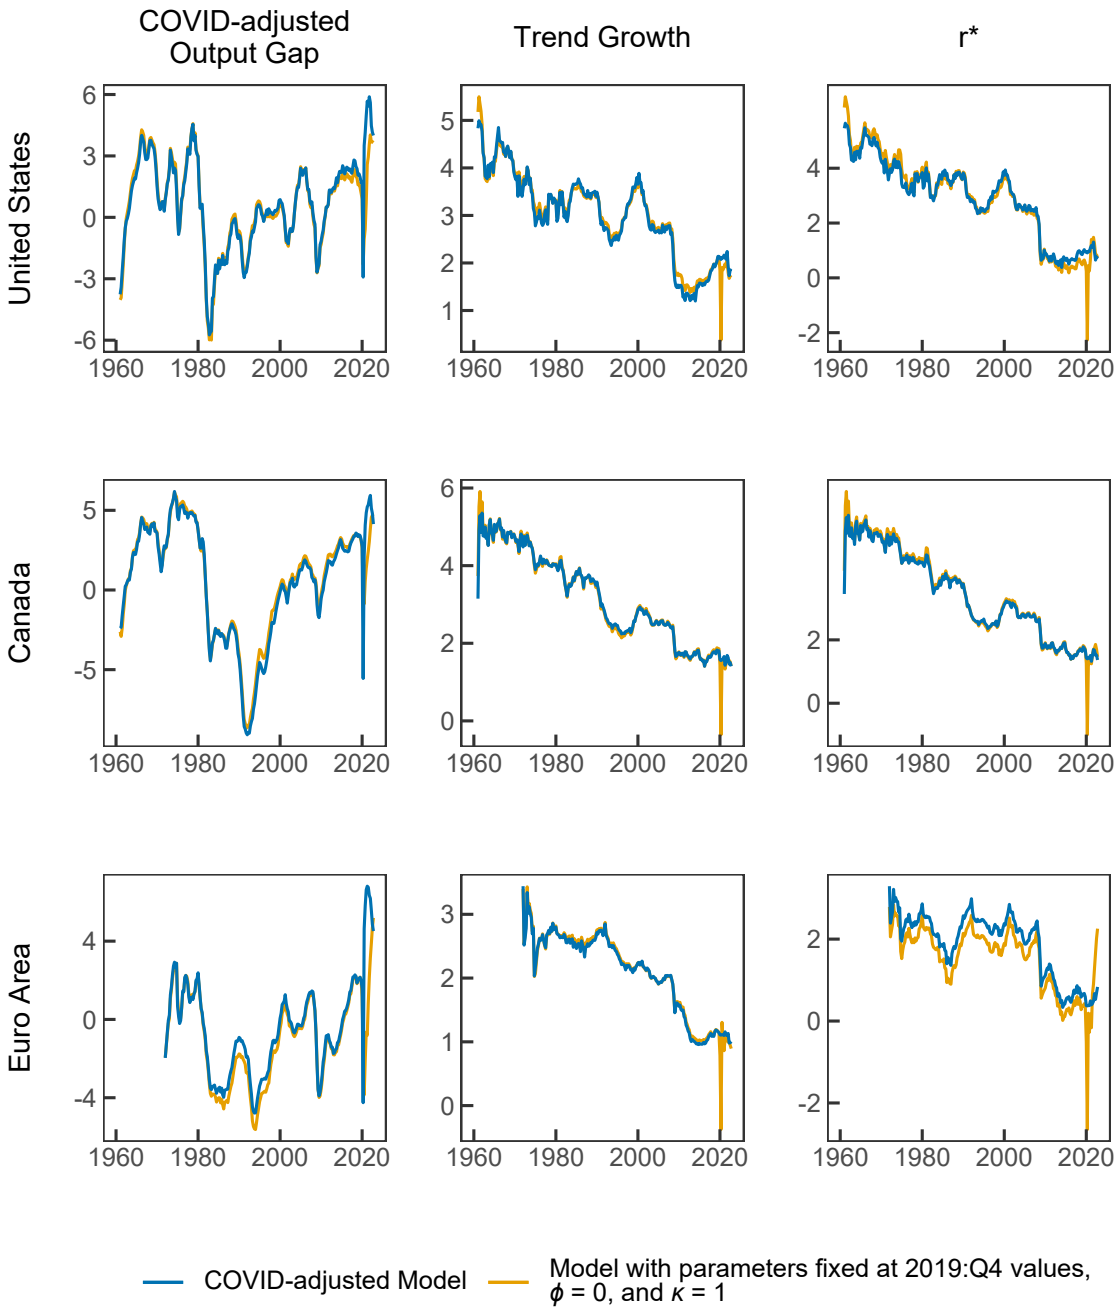
Notes: The left (right) panels show the standardized auxiliary residuals to the model IS equation (Phillips curve equation). The auxiliary residuals are smoothed estimators of the disturbances to the measurement equations, normalized by their standard deviations. The gold lines labeled "Model with parameters fixed at 2019:Q4 values" show the residuals using data through 2022:Q4, but where the model parameters are estimated on data through 2019:Q4. The blue lines labeled "COVID-adjusted model" show residuals from the full-sample estimation of the modified HLW model.

Figure 4: Standardized Auxiliary Residuals: Natural Rate of Output



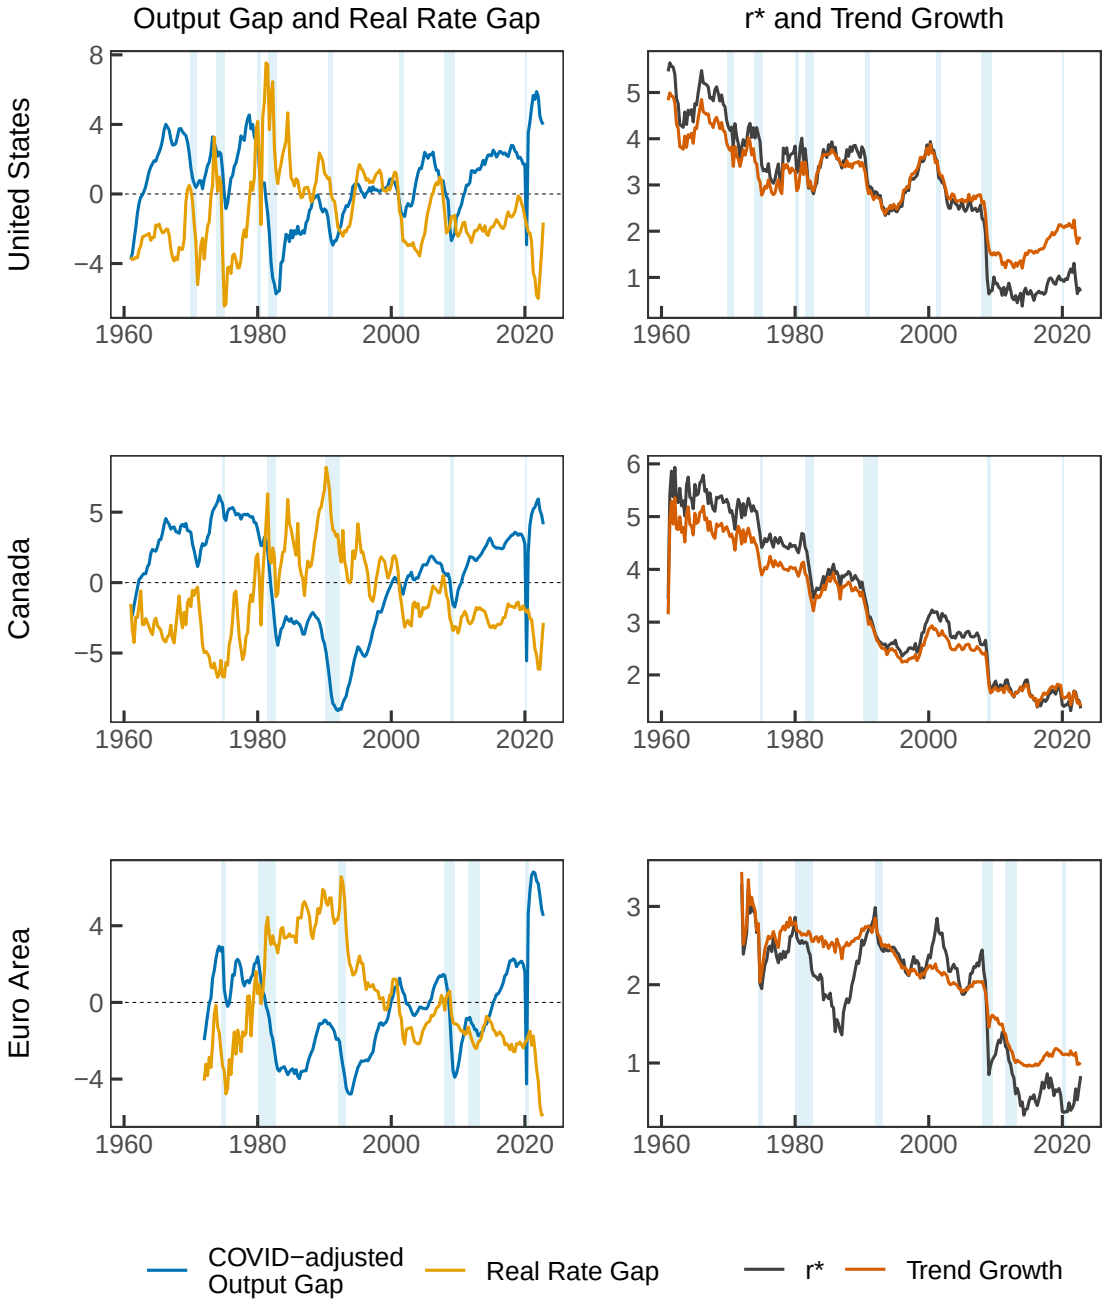
Notes: The panels show the standardized auxiliary residuals to the natural rate of output equation. The auxiliary residuals are smoothed estimators of the disturbances to the state equations, normalized by their standard deviations. The gold lines labeled "Model with parameters fixed at 2019:Q4 values" show the residuals using data through 2022:Q4, but where the model parameters are estimated on data through 2019:Q4. The blue lines labeled "COVID-adjusted model" show residuals from the full-sample estimation of the modified HLW model.

Figure 5: HLW Estimates of Key Latent Variables



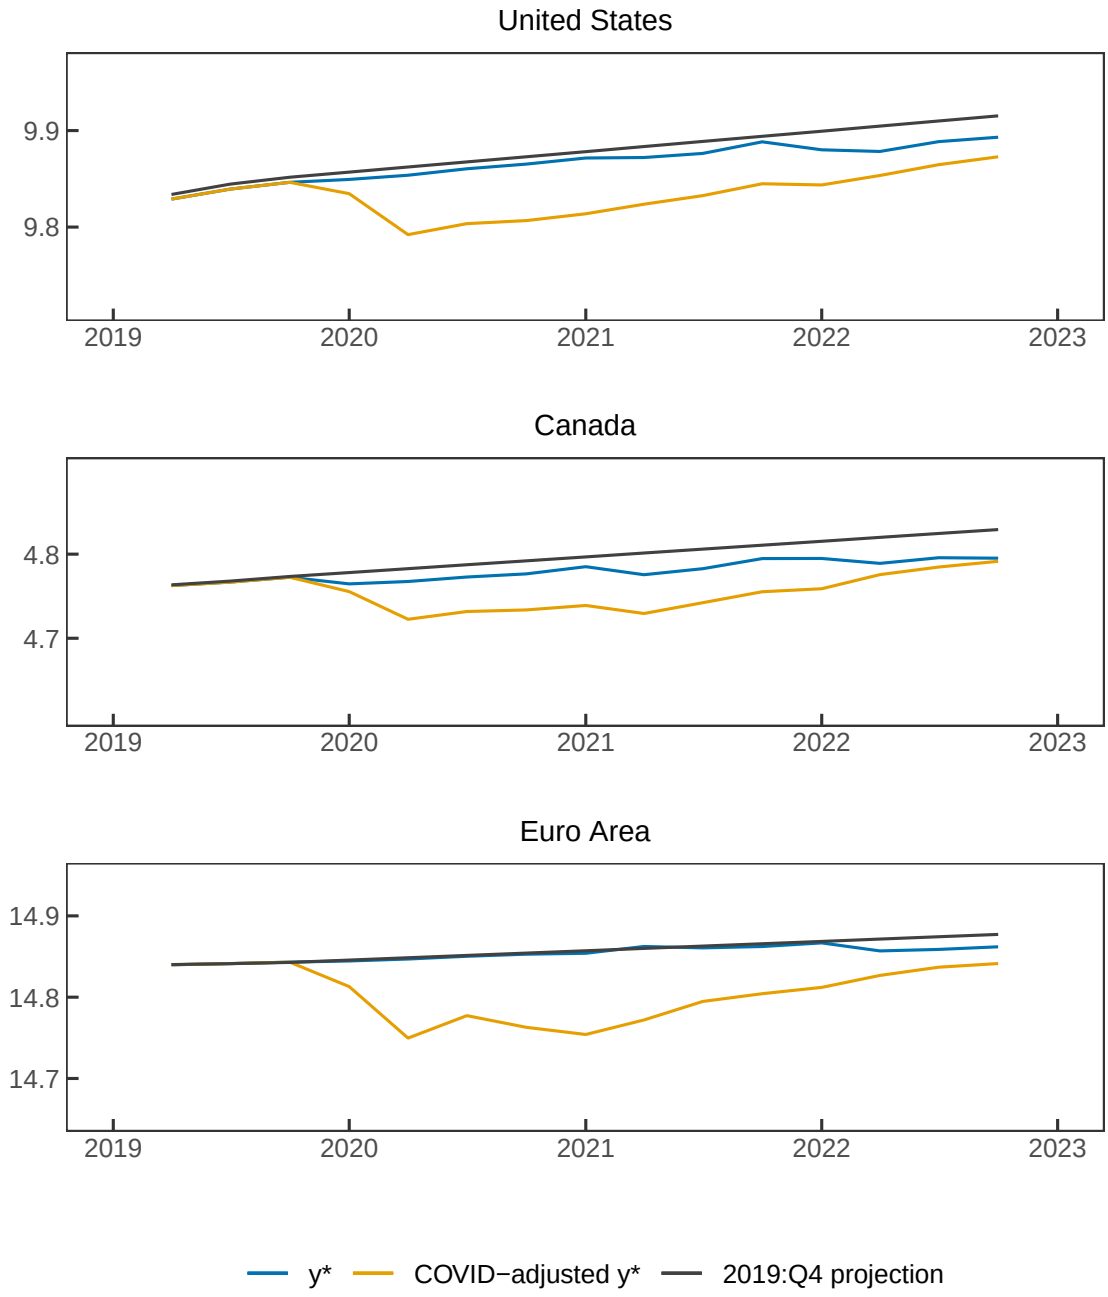
Notes: The gold lines labeled "Model with parameters fixed at 2019:Q4 values" show the estimates of the indicated variables using data through 2022:Q4, but where the model parameters are estimated on data through 2019:Q4. The blue lines labeled "COVID-adjusted model" show the full-sample estimates of the modified HLW model.

Figure 6: HLW Full-sample Estimation Results



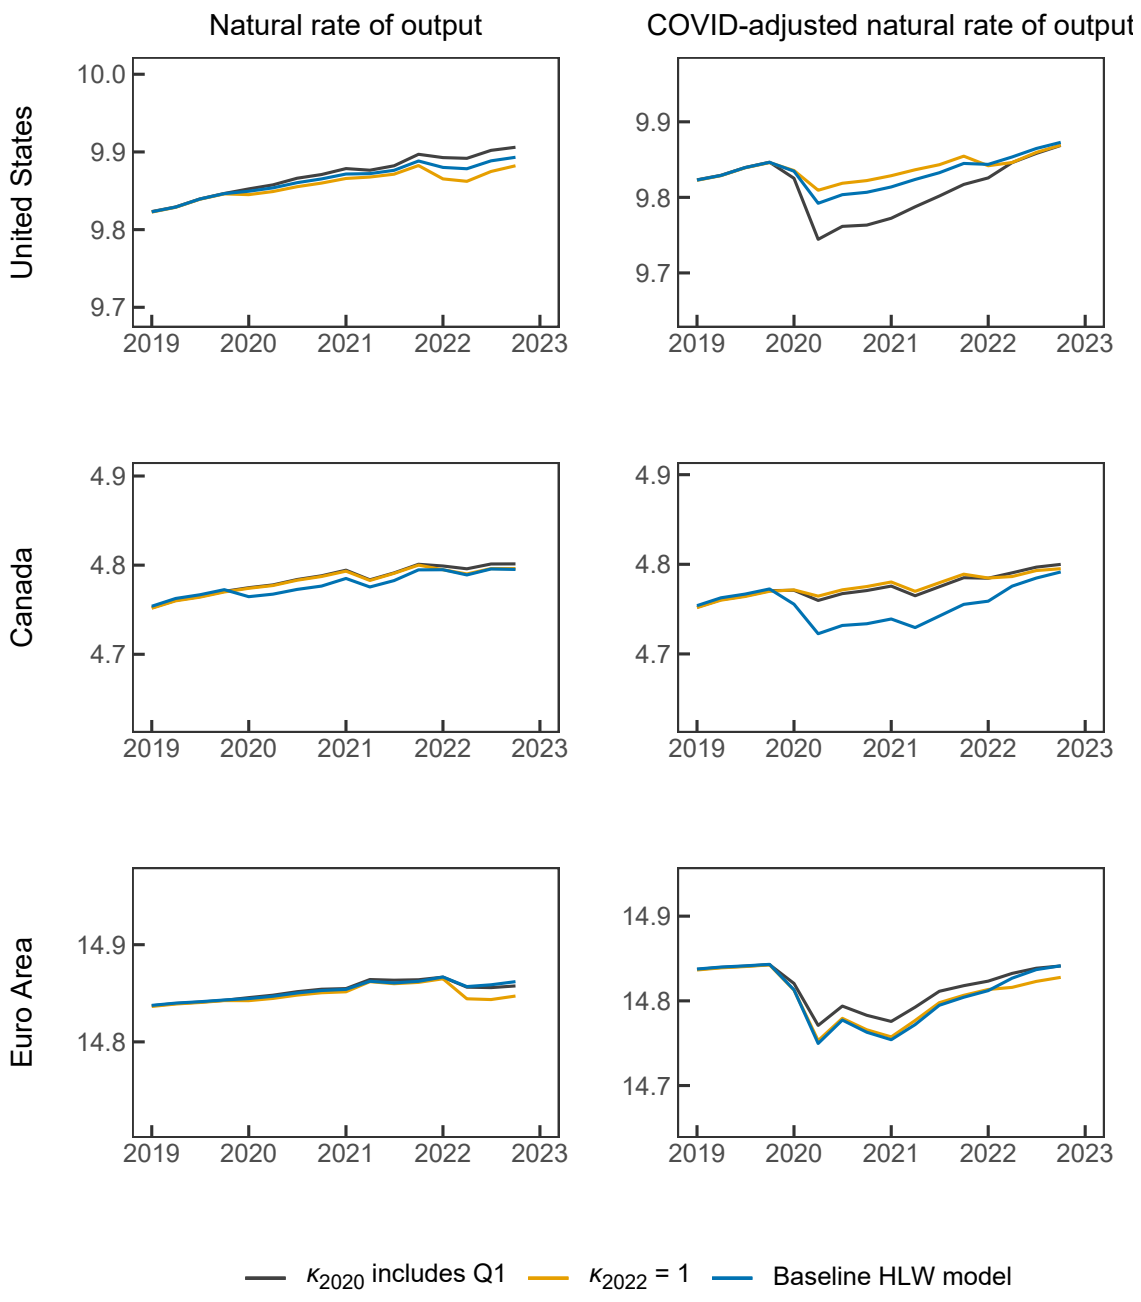
Notes: The lines show the full-sample estimates from the modified HLW model. The vertical shaded regions indicate periods when the economy was in a recession.

Figure 7: HLW Estimates of the Natural Rate of Output



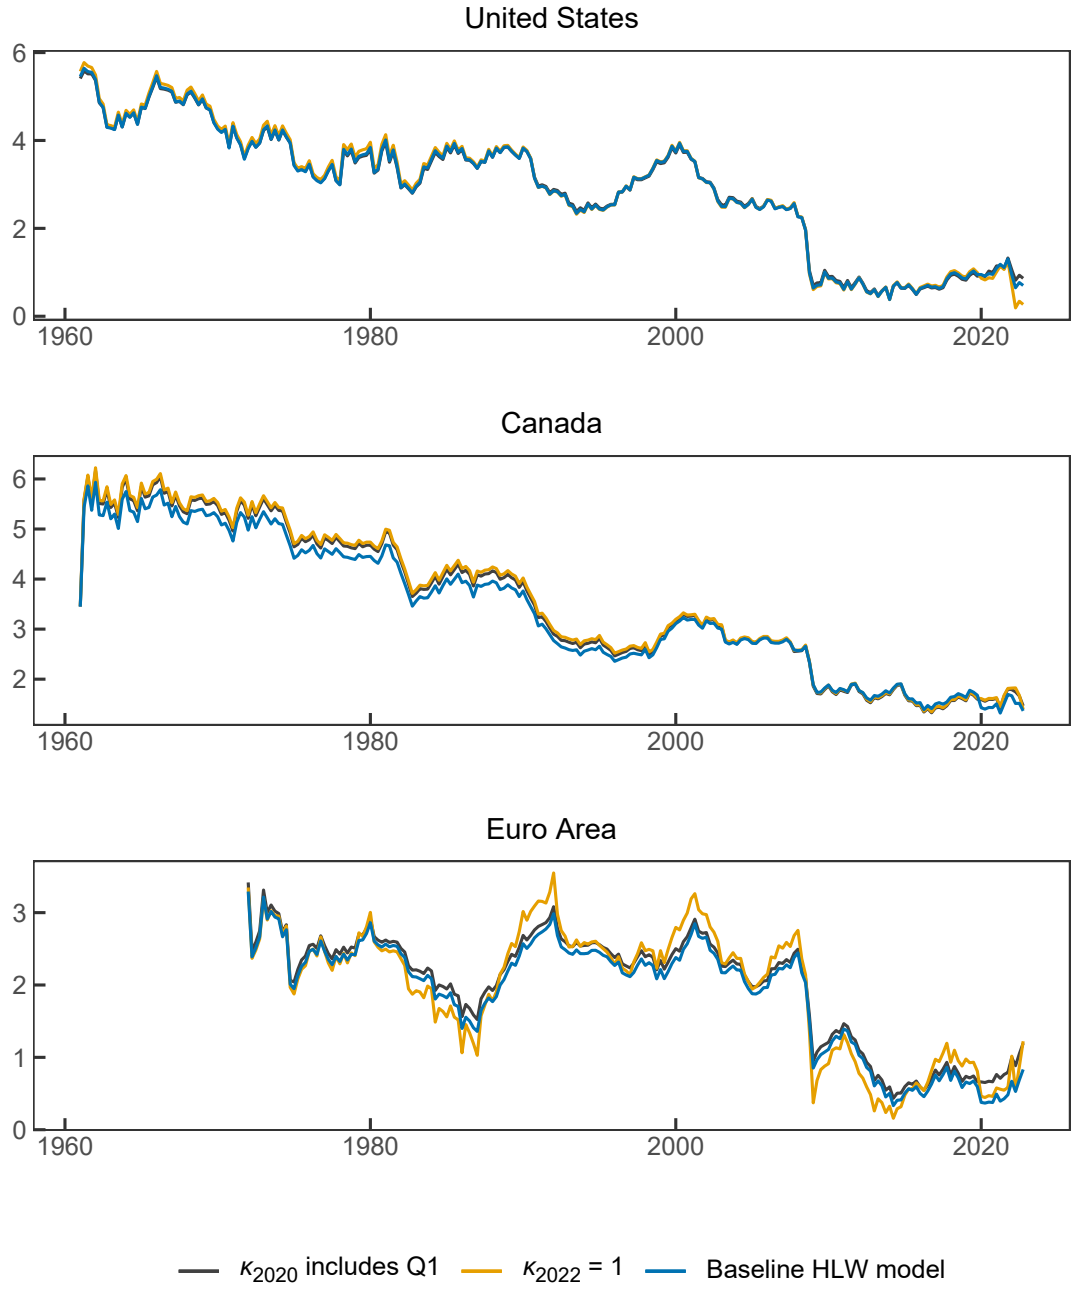
Notes: The black lines show the projected path of the natural rate of output based on model estimates using data through 2019:Q4. The blue lines show the estimates of the natural rate of output from the modified HLW model using data through 2022:Q4. The gold lines show the estimates of the COVID-adjusted natural rate of output from the modified HLW model using data through 2022:Q4.

Figure 8: Estimated Natural Rate of Output under Alternative Specifications of κ_t



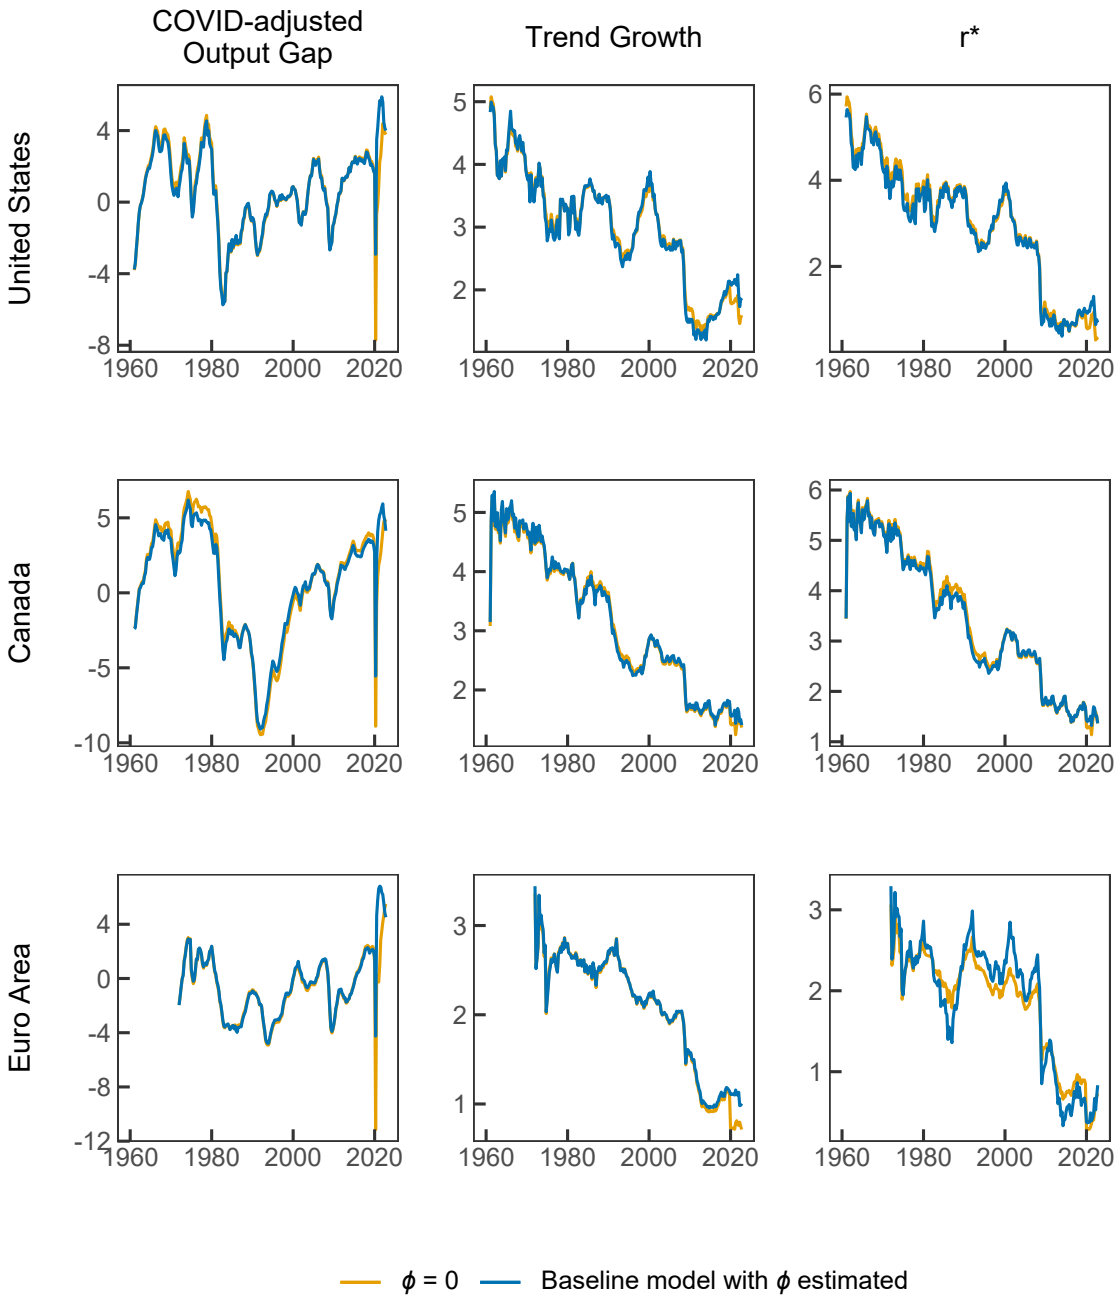
Notes: The left panels show estimates of the natural rate of output from the modified HLW model using data through 2022:Q4, while the right panels show estimates of the COVID-adjusted natural rate of output. The black lines labeled "κ_{2020} includes Q1" show estimates from a version of the model in which the variance scale parameter κ_{2020} includes 2020:Q1. The gold lines labeled "$\kappa_{2022} = 1$" show estimates from a specification of the model with variance scale parameters estimated for 2020 and 2021 only.

Figure 9: Natural Rate of Interest under Alternative Specifications of κ_t



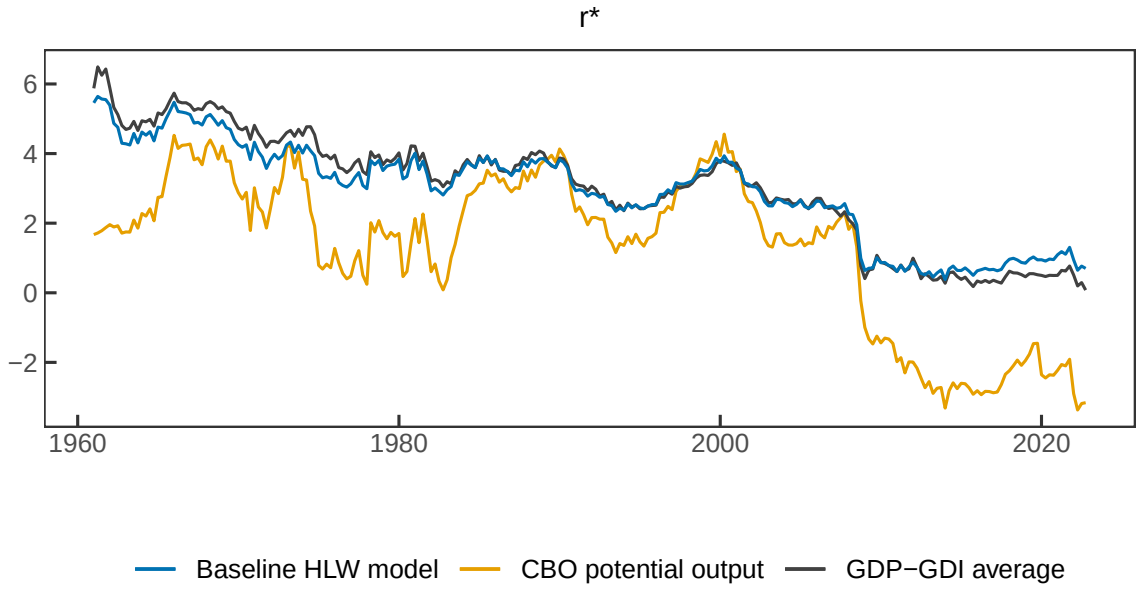
Notes: The panels show estimates of the natural rate of interest from the modified HLW model using data through 2022:Q4. The black lines labeled "κ_{2020} includes Q1" show estimates from a version of the model in which the variance scale parameter κ_{2020} includes 2020:Q1. The gold lines labeled "$\kappa_{2022} = 1$" show estimates from a specification of the model with variance scale parameters estimated for 2020 and 2021 only.

Figure 10: Model without COVID Supply Shock ($\phi = 0$)



Notes: The gold lines labeled "$\phi = 0$" show the estimates of the indicated variables from the COVID-adjusted model using data through 2022:Q4, but without the COVID supply shock in the model. The blue lines labeled "Baseline model with ϕ estimated" show the full-sample estimates of the modified HLW model, including the COVID supply shock.

Figure 11: Alternative Measures of Output



Notes: The gold line labeled "CBO potential output" shows the estimated natural rate of interest for the United States from an alternative specification of the model in which potential output is an exogenous variable, using data from the Congressional Budget Office. The black line labeled "GDP-GDI average" shows the estimated natural rate of interest from a version of the model in which output is specified as the average of GDP and Gross Domestic Income, using data from the Bureau of Economic Analysis. The blue line shows the estimates from the modified HLW model.

Appendix A1: State-Space Models

This section presents the COVID-adjusted and original HLW models in state-space form.[11] See HLW (2017) for a full description of our estimation procedure, which includes three stages. The first and second stage models represent versions of the final stage model, and each of the models can be cast in state-space form:

$$\mathbf{y}_t = \mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \xi_t + \epsilon_t \tag{19}$$

$$\xi_t = \mathbf{F} \cdot \xi_{t-1} + \eta_t \tag{20}$$

Here, $\mathbf{y}_t$ is a vector of contemporaneous endogenous variables, while $\mathbf{x}_t$ is a vector of exogenous and lagged endogenous variables. ξ_t the vector of unobserved state variables. In the HLW (2017) and Laubach and Williams (2003) models, the vectors of stochastic disturbances ϵ_t and η_t are assumed to be Gaussian and mutually uncorrelated, with mean zero and covariance matrices $\mathbf{R}$ and $\mathbf{Q}$, respectively. The covariance matrix $\mathbf{R}$ is assumed to be diagonal. In the COVID-adjusted model, we modify the covariance matrix $\mathbf{R}_t$ to be time-varying.

Each model has a corresponding vector of parameters to be estimated by maximum likelihood. Because maximum likelihood estimates of σ_g and σ_z, which are the standard deviations of the innovations to the g_t and z_t equations, are likely to be biased towards zero due to the pile-up problem (see Section 2.2 of HLW), we use Stock and Watson's (1998) median unbiased estimator to obtain estimates of two ratios, $\lambda_g \equiv \frac{\sigma_g}{\sigma_{y^*}}$ and $\lambda_z \equiv \frac{a_r \sigma_z}{\sigma_{\tilde{y}}}$. We estimate λ_g following the first stage model and λ_z following the second stage model, and impose these ratios in subsequent stages of the estimation, including when estimating the remaining model parameters by maximum likelihood. The COVID-adjusted model includes five additional parameters: the coefficient ϕ that translates the COVID indicator variable d_t into effects on output; the three variance scale parameters κ_{2020}, κ_{2021}, and κ_{2022}; and the coefficient c on trend growth g_t in the r_t^* equation, which appears only in the stage 3 model and is estimated in LW (2003) but fixed at unity in HLW (2017).

In addition to estimating the relationship between trend growth in the natural rate of output and r_t^*, we make two minor technical changes to the model that are not related to the COVID-19 pandemic.[12] First, we include a second lag of trend

[11] Notation follows Hamilton (1994) and is consistent with the corresponding R programs.

[12] We also explicitly include g_t and z_t in the vector of unobserved state variables in addition to

growth, g_{t-2}, in the stage 2 IS equation, consistent with the IS equation specification in the stage 3 model. Second, we correct the stage 2 state-space model so that the y_t^* equation is $y_t^* = y_{t-1}^* + g_{t-1} + \epsilon_{y^*,t}$ as expressed in the paper; previously, the second stage y_t^* equation included the second lag of trend growth, g_{t-2}, rather than the first lag in error (Buncic, 2021, 2022). These modifications to the model are highlighted in blue text in the following sections, while changes in response to the COVID-19 pandemic are highlighted in red. These technical changes have minor effects on our estimates of the unobserved state variables, including r_t^*.

7.1 The COVID-Adjusted State-Space Models

7.1.1 The COVID-adjusted Stage 1 Model

The first-stage model, which corresponds to the *rstar.stage1.R* program, can be represented by the following matrices:

$$\mathbf{y}_t = [y_t, \pi_t]' \tag{21}$$

$$\mathbf{x}_t = [y_{t-1}, y_{t-2}, \pi_{t-1}, \pi_{t-2,4}, d_t, d_{t-1}, d_{t-2}]' \tag{22}$$

$$\xi_t = \left[y_t^*, y_{t-1}^*, y_{t-2}^*\right]' \tag{23}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} \\ 0 & -b_y & 0 \end{bmatrix}, \quad \mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & 0 & 0 & \phi & -\phi a_{y,1} & -\phi a_{y,2} \\ b_y & 0 & b_\pi & 1-b_\pi & 0 & -\phi b_y & 0 \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{y^*}^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{R}_t = \begin{bmatrix} (\kappa_t \sigma_{\tilde{y}})^2 & 0 \\ 0 & (\kappa_t \sigma_\pi)^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_1 = [a_{y,1}, a_{y,2}, b_\pi, b_y, g, \sigma_{\tilde{y}}, \sigma_\pi, \sigma_{y*}, \phi, \kappa_{2020Q2-Q4}, \kappa_{2021}, \kappa_{2022}]$$

two lags of each variable as in HLW. This is purely an accounting change and has no effect on the estimates.

7.1.2 The COVID-adjusted Stage 2 Model

The second-stage model, which corresponds to the *rstar.stage2.R* program, can be represented by the following matrices:

$$\begin{aligned} \mathbf{y}_t &= [y_t, \pi_t]' && (24)\\ \mathbf{x}_t &= [y_{t-1}, y_{t-2}, r_{t-1}, r_{t-2}, \pi_{t-1}, \pi_{t-2,4}, 1, d_t, d_{t-1}, d_{t-2}]' && (25)\\ \xi_t &= \left[y_t^*, y_{t-1}^*, y_{t-2}^*, g_t, g_{t-1}, g_{t-2}\right]' && (26) \end{aligned}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} & 0 & \frac{a_g}{2} & \frac{a_g}{2} \\ 0 & -b_y & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & \frac{a_r}{2} & \frac{a_r}{2} & 0 & 0 & a_0 & \phi & -\phi a_{y,1} & -\phi a_{y,2} \\ b_y & 0 & 0 & 0 & b_\pi & 1-b_\pi & 0 & 0 & -\phi b_y & 0 \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{y^*}^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & (\lambda_g \sigma_{y^*})^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{R}_t = \begin{bmatrix} (\kappa_t \sigma_{\tilde{y}})^2 & 0 \\ 0 & (\kappa_t \sigma_\pi)^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_2 = [a_{y,1}, a_{y,2}, a_r, a_0, a_g, b_\pi, b_y, \sigma_{\tilde{y}}, \sigma_\pi, \sigma_{y*}, \phi, \kappa_{2020Q2-Q4}, \kappa_{2021}, \kappa_{2022}]$$

7.1.3 The COVID-adjusted Stage 3 Model

The third-stage model, which corresponds to the *rstar.stage3.R* program, can be represented by the following matrices:

$$\begin{aligned} \mathbf{y}_t &= [y_t, \pi_t]' && (27)\\ \mathbf{x}_t &= [y_{t-1}, y_{t-2}, r_{t-1}, r_{t-2}, \pi_{t-1}, \pi_{t-2,4}, d_t, d_{t-1}, d_{t-2}]' && (28)\\ \xi_t &= \left[y_t^*, y_{t-1}^*, y_{t-2}^*, g_t, g_{t-1}, g_{t-2}, z_t, z_{t-1}, z_{t-2}\right]' && (29) \end{aligned}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} & 0 & -4c \cdot \frac{a_r}{2} & -4c \cdot \frac{a_r}{2} & 0 & \frac{-a_r}{2} & \frac{-a_r}{2} \\ 0 & -b_y & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & \frac{a_r}{2} & \frac{a_r}{2} & 0 & 0 & \phi & -\phi a_{y,1} & -\phi a_{y,2} \\ b_y & 0 & 0 & 0 & b_\pi & 1-b_\pi & 0 & -\phi b_y & 0 \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{y^*}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & (\lambda_g \sigma_{y^*})^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\lambda_z \sigma_{\tilde{y}}}{a_r}\right)^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\mathbf{R}_t = \begin{bmatrix} (\kappa_t \sigma_{\tilde{y}})^2 & 0 \\ 0 & (\kappa_t \sigma_\pi)^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_3 = [a_{y,1}, a_{y,2}, a_r, b_\pi, b_y, \sigma_{\tilde{y}}, \sigma_\pi, \sigma_{y*}, \phi, c, \kappa_{2020Q2-Q4}, \kappa_{2021}, \kappa_{2022}]$$

The law of motion for the natural rate of interest is $r_t^* = c \cdot g_t + z_t$.

7.2 The HLW (2017) State-Space Models

7.2.1 The Stage 1 Model

The first-stage model, which corresponds to the *rstar.stage1.R* program, can be represented by the following matrices:

$$\mathbf{y}_t = [y_t, \pi_t]' \tag{30}$$

$$\mathbf{x}_t = [y_{t-1}, y_{t-2}, \pi_{t-1}, \pi_{t-2,4}]' \tag{31}$$

$$\xi_t = [y_t^*, y_{t-1}^*, y_{t-2}^*]' \tag{32}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} \\ 0 & -b_y & 0 \end{bmatrix}, \quad \mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & 0 & 0 \\ b_y & 0 & b_\pi & 1-b_\pi \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{y^*}^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} \sigma_{\tilde{y}}^2 & 0 \\ 0 & \sigma_\pi^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_1 = [a_{y,1}, a_{y,2}, b_\pi, b_y, g, \sigma_{\tilde{y}}, \sigma_\pi, \sigma_{y*}]$$

7.2.2 The Stage 2 Model

The second-stage model, which corresponds to the *rstar.stage2.R* program, can be represented by the following matrices:

$$\mathbf{y}_t = [y_t, \pi_t]' \tag{33}$$

$$\mathbf{x}_t = [y_{t-1}, y_{t-2}, r_{t-1}, r_{t-2}, \pi_{t-1}, \pi_{t-2,4}, 1]' \tag{34}$$

$$\xi_t = \left[y_t^*, y_{t-1}^*, y_{t-2}^*, g_{t-1}\right]' \tag{35}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} & a_g \\ 0 & -b_y & 0 & 0 \end{bmatrix}, \quad \mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & \frac{a_r}{2} & \frac{a_r}{2} & 0 & 0 & a_0 \\ b_y & 0 & 0 & 0 & b_\pi & 1-b_\pi & 0 \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{y^*}^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \left(\lambda_g \sigma_{y^*}\right)^2 \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} \sigma_{\tilde{y}}^2 & 0 \\ 0 & \sigma_\pi^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_2 = [a_{y,1}, a_{y,2}, a_r, a_0, a_g, b_\pi, b_y, \sigma_{\tilde{y}}, \sigma_\pi, \sigma_{y*}]$$

7.2.3 The Stage 3 Model

The third-stage model, which corresponds to the *rstar.stage3.R* program, can be represented by the following matrices:

$$\begin{aligned} \mathbf{y}_t &= [y_t, \pi_t]' && (36) \\ \mathbf{x}_t &= [y_{t-1}, y_{t-2}, r_{t-1}, r_{t-2}, \pi_{t-1}, \pi_{t-2,4}]' && (37) \\ \xi_t &= \left[y_t^*, y_{t-1}^*, y_{t-2}^*, g_{t-1}, g_{t-2}, z_{t-1}, z_{t-2}\right]' && (38) \end{aligned}$$

$$\mathbf{H}' = \begin{bmatrix} 1 & -a_{y,1} & -a_{y,2} & \frac{-a_r}{2} & \frac{-a_r}{2} & \frac{-a_r}{2} & \frac{-a_r}{2} \\ 0 & -b_y & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{A}' = \begin{bmatrix} a_{y,1} & a_{y,2} & \frac{a_r}{2} & \frac{a_r}{2} & 0 & 0 \\ b_y & 0 & 0 & 0 & b_\pi & 1-b_\pi \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \left(1+\lambda_g^2\right)\sigma_{y^*}^2 & 0 & 0 & \left(\lambda_g\sigma_{y^*}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \left(\lambda_g\sigma_{y^*}\right)^2 & 0 & 0 & \left(\lambda_g\sigma_{y^*}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \left(\frac{\lambda_z\sigma_{\widetilde{y}}}{a_r}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\mathbf{R} = \begin{bmatrix} \sigma_{\widetilde{y}}^2 & 0 \\ 0 & \sigma_\pi^2 \end{bmatrix}$$

The vector of parameters to be estimated by maximum likelihood is as follows:

$$\theta_3 = [a_{y,1}, a_{y,2}, a_r, b_\pi, b_y, \sigma_{\widetilde{y}}, \sigma_\pi, \sigma_{y*}]$$

The law of motion for the natural rate of interest is $r_t^* = g_t + z_t$.

Appendix A2: Data

For each economy, we require data for real GDP, inflation, and the short-term nominal interest rate, as well as a procedure to compute inflation expectations to calculate the ex ante real short-term interest rate r_t.[13] The variable y_t refers to the logarithm of real GDP. The inflation measure is the annualized quarterly growth rate of the specified consumer price series. With the exception of the United States, for which core personal consumption expenditure (PCE) price data are available over the entire sample, the inflation series is constructed by splicing the core price index with an all-items price index. We use a four-quarter moving average of past inflation as a proxy for inflation expectations in constructing the ex ante real interest rate. Short-term interest rates are expressed on a 365-day annualized basis.

For the United States, we use real GDP and core PCE data published by the Bureau of Economic Analysis. Inflation is constructed using the price index for PCE excluding food and energy, referred to as core PCE inflation. The short-term interest rate is the annualized nominal federal funds rate, available from the Board of Governors. Because the federal funds rate frequently fell below the discount rate prior to 1965, we use the Federal Reserve Bank of New York's discount rate, part of the IMF's International Financial Statistics Yearbooks (IFS), prior to 1965. All U.S. data can be downloaded from the St. Louis Fed's Federal Reserve Economic Data (FRED) website.[14]

Canadian real GDP data is taken from the IMF's IFS. The short-term nominal interest rate is the Bank of Canada's target for the overnight rate, taken as the end-of-period value for each month and aggregated to quarterly frequency. Since the Bank of Canada began treating the target rate as its key interest rate in May 2001, we use the bank rate as the short-term interest rate prior to that date. We use the Bank of Canada's core Consumer Price Index to construct our inflation series. Prior to 1984, we use CPI containing all items. With the exception of real GDP, all data is from

[13] A detailed description of our data and programs, as well as replication materials for the standard HLW model, is available on the Federal Reserve Bank of New York's website.

[14] Mnemonics are as follows. Real GDP: GDPC1; Core PCE: PCEPILFE; Federal Funds Rate: FEDFUNDS; FRBNY Discount Rate: INTDSRUSM193N.

Statistics Canada.[15]

Euro Area data is from the Area Wide Model (AWM), available from the Euro Area Business Cycle Network (Fagan et al., 2001). The inflation measure is based on the core price index, HICP excluding energy (series HEX) beginning in 1988; prior to 1988 we use the overall price index HICP. The nominal short-term interest rate is the three-month rate (series STN) and the real GDP mnemonic is YER. At the time of publication, the final update to the Area Wide Model was in 2017; we update the three series from the ECB's Statistical Data Warehouse.[16]

[15]Mnemonics from Statistics Canada are as follows. Core CPI: v41690926 (Table 326-0022); CPI: v41690914 (Table 326-0022); v41690973 (Table 326-0020); Bank Rate: v122530 (Table 176-0043); Target Rate: v39079 (Table 176-0048). Real GDP is IFS series "Gross Domestic Product, Real, Seasonally adjusted, Index".

[16]Mnemonics for the SDW are as follows.
Real GDP: MNA.Q.Y.I8.W2.S1.S1.B.B1GQ._Z._Z._Z.EUR.LR.N;
Core HICP: ICP.M.U2.N.XE0000.4.INX;
Nominal Short-term Rate: FM.Q.U2.EUR.RT.MM.EURIBOR3MD_.HSTA.
Because data availability is longer for the non-seasonally adjusted price series, we use those and seasonally adjust them.

Appendix A3: Sensitivity of Pre-Pandemic Estimates

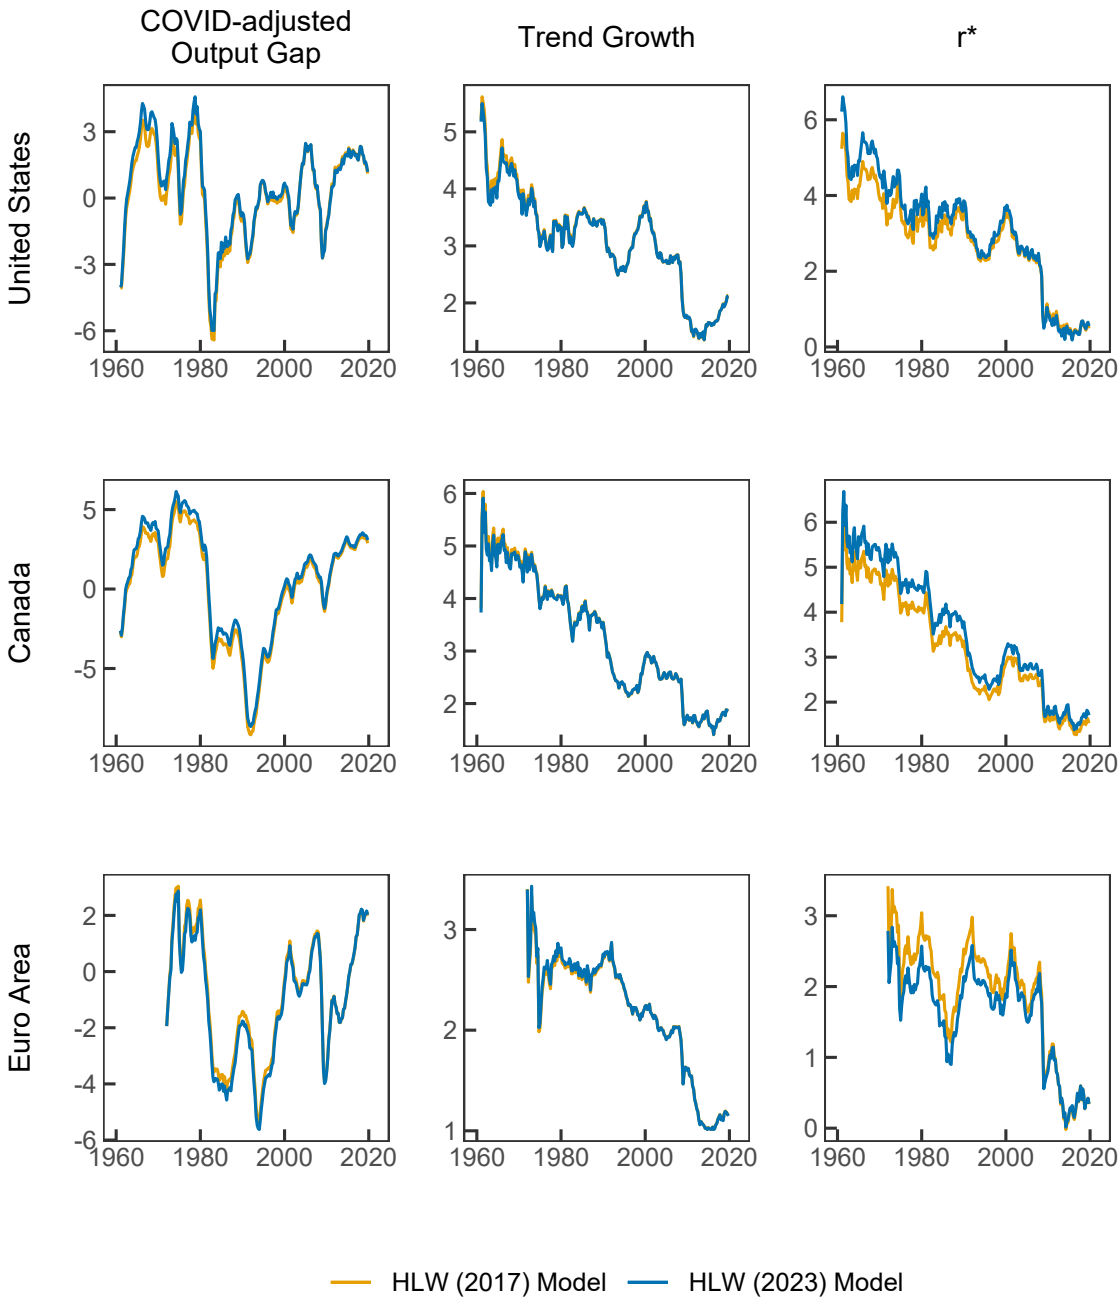
Table A1 reports parameter estimates from the model estimated through 2019:Q4. Figure A1 compares estimates from the original HLW (2017) model and the modified model through 2019:Q4. Both use the current data vintage at the time of publication.

Table A1: Parameter Estimates, Sample Ending 2019:Q4

Parameter	United States	Canada	Euro Area
Sample	1961-2019	1961-2019	1972-2019
λ_g	0.053	0.052	0.036
λ_z	0.031	0.016	0.032
$\sum a_y$	0.941	0.951	0.950
a_r	−0.067	−0.065	−0.037
	(3.973)	(2.969)	(1.798)
b_y	0.076	0.047	0.062
	(3.077)	(1.697)	(1.713)
c	1.198	1.130	0.816
	(3.484)	(2.716)	(0.888)
$\sigma_{\tilde{y}}$	0.344	0.395	0.289
σ_π	0.794	1.359	0.972
σ_{y^*}	0.568	0.581	0.397
σ_g	0.121	0.121	0.058
σ_z	0.157	0.096	0.246
$\sigma_{r^*} = \sqrt{c^2\sigma_g^2 + \sigma_z^2}$	0.213	0.167	0.251
S.E. (sample ave.)			
r^*	1.236	1.625	3.480
g	0.395	0.422	0.260
y^*	1.559	2.470	1.875
S.E. (final obs.)			
r^*	1.656	1.844	4.822
g	0.546	0.574	0.357
y^*	1.971	2.761	2.429

Notes: t statistics are in parentheses; σ_g is expressed at an annual rate. Sample through 2019:Q4, using current data vintage at time of publication.

Figure A1: Estimates through 2019:Q4, HLW (2023) vs. HLW (2017) Models



Notes: The gold lines labeled "HLW (2017) Model" show the estimates of the indicated variables using data through 2019:Q4 in the standard HLW model. The blue lines labeled "HLW (2023) Model" show the estimates of the modified HLW model on data through 2019:Q4. Both sets of estimates use the current data vintage at the time of publication.

Appendix A4: Laubach and Williams (2003) Model

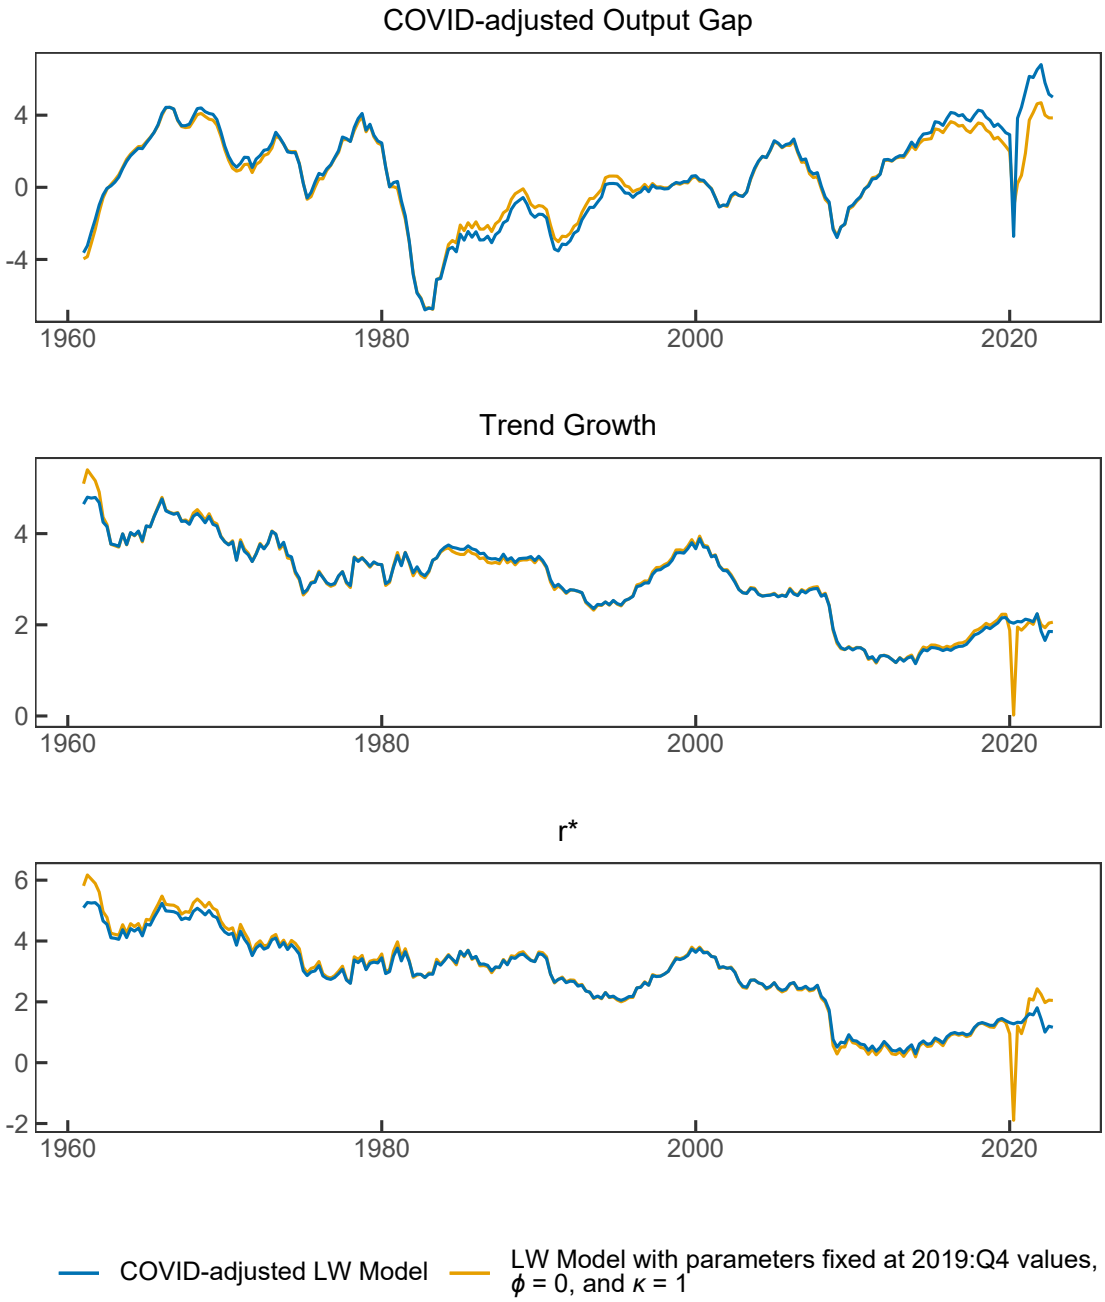
We apply the same model adjustments described in the paper to estimate the LW model for the United States. In addition to the pandemic-related modifications, we make the two technical adjustments described in Appendix A1. We make two additional changes that are already implemented in the HLW model: we initialize the state vector and its covariance matrix as described in HLW (2017), footnote 6, and impose the constraints that the slope of the IS equation (a_r) is negative and the slope of the Phillips curve (b_y) is positive in all three stages of the estimation. Table A2 reports parameter estimates from the LW model estimated through 2022:Q4. Figure A2 displays estimates from the COVID-adjusted LW model and from a version of the model with parameters fixed at 2019:Q4 values, analogous to Section 2.2.

Table A2: Parameter Estimates from LW Model

Parameter	LW	HLW
Sample	1961-2022	1961-2022
λ_g	0.070	0.073
λ_z	0.027	0.021
$\sum a_y$	0.946	0.936
a_r	−0.090	−0.079
	(4.718)	(4.215)
b_y	0.049	0.073
	(2.328)	(3.003)
c	1.097	1.128
	(3.256)	(3.574)
ϕ	−0.067	−0.085
	(2.212)	(2.199)
$\kappa_{2020Q2-Q4}$	9.572	9.033
	(2.182)	(2.351)
κ_{2021}	1.786	1.791
	(3.052)	(2.941)
κ_{2022}	1.000	1.676
	(1.766)	(2.060)
$\sigma_{\tilde{y}}$	0.457	0.452
σ_π	0.754	0.787
σ_{y^*}	0.503	0.500
σ_g	0.142	0.145
σ_z	0.139	0.118
$\sigma_{r^*} = \sqrt{c^2\sigma_g^2 + \sigma_z^2}$	0.208	0.202
S.E. (sample ave.)		
r^*	1.193	1.140
g	0.433	0.428
y^*	1.928	1.612
S.E. (final obs.)		
r^*	1.679	1.565
g	0.663	0.689
y^*	3.256	2.914

Notes: t statistics are in parentheses; σ_g is expressed at an annual rate.

Figure A2: Estimates from LW Model



Notes: The gold lines labeled "LW model with parameters fixed at 2019:Q4 values" show the estimates of the indicated variables using data through 2022:Q4, but where the model parameters are estimated on data through 2019:Q4. The blue lines labeled "COVID-adjusted LW model" show the full-sample estimates of the modified LW model.

NBER WORKING PAPER SERIES

CENTRAL BANK COMMUNICATION AND POLICY EFFECTIVENESS

Michael Woodford

Working Paper 11898
http://www.nber.org/papers/w11898

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2005

Presented at the Federal Reserve Bank of Kansas City Symposium \The Greenspan Era: Lessons for the Future," Jackson Hole, Wyoming, August 25-27, 2005. I would like to thank Charlie Bean, Ben Bernanke, Alan Blinder, Michael Ehrmann, Marcel Fratzscher, Charles Goodhart, Larry Meyer, and Anders Vredin for helpful comments on an earlier draft, without implicating any of them in the views expressed here. I would also like to thank Mauro Roca for research assistance and the National Science Foundation for research support through a grant to the NBER. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Michael Woodford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Central Bank Communication and Policy Effectiveness
Michael Woodford
NBER Working Paper No. 11898
December 2005
JEL No. E52, E58

ABSTRACT

A notable change in central banking over the past 15 years has been a world-wide movement toward increased communication by central banks about their policy decisions, the targets that they seek to achieve through those decisions, and the central bank's view of the economy's likely future evolution. This paper considers the role of such communication in the successful conduct of monetary policy, with a particular emphasis on an issue that remains controversial: to what extent is it desirable for central banks to comment on the likely path of short-term interest rates?

After reviewing general arguments for and against central-bank transparency, the paper considers two specific contexts in which central banks have been forced to consider how much they are willing to say about the future path of interest rates. The first is the experiment with policy signaling by the FOMC in the U.S., using the statement released following each Committee meeting, since August 2003. The second is the need to make some assumption about future policy when producing the projections (for future inflation and other variables) that are central to inflation-forecast targeting procedures, of the kind used by the Bank of England, the Swedish Riksbank, the Reserve Bank of New Zealand, and others. In both cases, it is argued that increased willingness to share the central bank's own assumptions about future policy with the public has increased the predictability of policy, in ways that are likely to have improved central bank's ability to achieve their stabilization objectives.

Michael Woodford
Department of Economics
Columbia University
1009B I.A.B., mc3308
New York, NY 10027
and NBER
mw2230@columbia.edu

One of the most notable changes at the Federal Reserve during the tenure of Alan Greenspan as Chairman of the Board of Governors has been a steady increase in the FOMC's willingness to talk openly about the policy decisions that it has made and those it is likely to make in the future. Before the 1990s, central banking was shrouded in mystery, at the Fed as elsewhere. The title of William Greider's 1987 bestseller about the Fed — *Secrets of the Temple* — gives an idea of the common perception of the institution at the beginning of the Greenspan era. This "mystique" of central banking was jealously guarded by central bankers — as the epigraph indicates — as essential to their success.

Things have changed rapidly over the past 15 years, both at the Fed and elsewhere. Indeed, St. Louis Fed President William Poole (2005) lists the increase in transparency, and the consequent increase in the predictability of monetary policy, as one of the four defining characteristics of "the Greenspan policy regime." Before 1994, the FOMC made no public announcement regarding its target for the federal funds rate following the meetings at which the target was determined; markets had to try to infer the target rate from the type and size of open-market operations that were subsequently conducted by the Trading Desk in New York to implement the policy. According to Poole, "before Greenspan many within the Fed believed that policy effectiveness depended on taking markets by surprise." But since February 1994, the FOMC has issued a public statement following each meeting at which the target has been changed, indicating the new target rate. The FOMC has also been increasingly willing to give advance signals of the likely future stance of policy. Beginning in December 1998, the FOMC began to include in the post-meeting statement an assessment of the FOMC's current "bias" with respect to possible changes in the stance of policy; in December 1999, the Committee decided that from then on it would issue a statement after every meeting, whether policy was changed or not, and that this would include a "balance of risks" assessment, understood to refer to a time horizon extending beyond the next Committee meeting. Since August 2003 — as is discussed further in section 2 — the post-meeting statements have included even more explicit statements about the likely future path of interest rates. This aspect of the statement now attracts considerable attention, in financial markets and in the financial press. Most recently, the FOMC has moved to expedite the release of the minutes of its deliberations, so that these are now available to the public before the next Committee meeting. This too has facilitated public understanding of current policy, and helped to increase the clarity with which the FOMC is able to explain its

view of the likely future path of policy.

Poole argues that the "improved predictability of policy [under Greenspan] has had much to do with improved effectiveness of policy." Is there reason to believe that this is true? And more specifically, does the Fed's bold recent experiment in greater explicitness about the future outlook for interest rates represent an innovation that should be expected to further enhance the effectiveness of policy, or does it represent a step too far?[1]

I shall begin by reviewing the general case for the importance of effective communication for effective monetary policy, and then ask, in the light of these general considerations, to what extent it makes sense for a central bank to be willing to make public statements about future policy. I then discuss in further detail two specific contexts in which central banks have recently given a great deal of attention to the question of how much they should talk about the future path of interest rates. The first is the Fed's experiment with policy signaling since August 2003, already mentioned. The second concerns the assumption about future policy that should be used in projections of the economy's likely future evolution that are made public. This has been a particularly crucial issue for the inflation-forecast targeting central banks, for reasons discussed further in section 3; but the issue is also being debated within the Federal Reserve System, especially among those considering the possibility of inflation targeting in the United States.

1 Why Communication Matters

The importance of communication strategy for policy effectiveness follows from a fundamental feature of the kind of problem that a central bank is called upon to solve. Central banking is not like steering an oil tanker, or even guiding a spacecraft, which follows a trajectory that depends on constantly changing factors, but that does *not* depend on the vehicle's own expectations about where it is heading. Because the key decisionmakers in an economy are forward-looking, central banks affect the economy as much through their influence on *expectations* as through any direct, mechanical effects of central bank trading in the market for overnight cash. As a consequence, there is good reason for a central bank to commit itself to a systematic approach to

[1] Even William Poole, in the remarks just cited, refrains from taking a stand on this last issue.

policy, that not only provides an explicit framework for decisionmaking within the bank, but that is also used to explain the bank's decisions to the public.

1.1 Central Banking as Management of Expectations

It is important for the public to understand the central bank's actions, to the greatest extent possible, not only for reasons of democratic legitimacy — though this is an excellent reason itself, given that central bankers are granted substantial autonomy in the execution of their task — but also in order for monetary policy to be most effective. For not only do expectations about policy matter, but, at least under current conditions, very little *else* matters. Few central banks of major industrial nations still make much use of credit controls or other attempts to directly regulate the flow of funds through financial markets and institutions. Increases in the sophistication of the financial system have made it more difficult for such controls to be effective, and in any event the goal of improvement of the efficiency of the sectoral allocation of resources stressed above would hardly be served by such controls, which (if successful) inevitably create inefficient distortions in the relative cost of funds to different parts of the economy.

Instead, banks restrict themselves to interventions that seek to control the overnight interest rate in an interbank market for central-bank balances (for example, the federal funds rate in the U.S.). But the current level of overnight interest rates *as such* is of negligible importance for economic decisionmaking; if a change in the overnight rate were thought to imply only a change in the cost of overnight borrowing for that one night, then even a large change (say, a full percentage point increase) would make little difference to anyone's spending decisions. The effectiveness of changes in central-bank targets for overnight rates in affecting spending decisions (and hence ultimately pricing and employment decisions) is wholly dependent upon the impact of such actions upon other financial-market prices, such as longer-term interest rates, equity prices and exchange rates. These are plausibly linked, through arbitrage relations, to the short-term interest rates most directly affected by central-bank actions; but it is the expected future path of short-term rates over coming months and even years that should matter for the determination of these other asset prices, rather than the current level of short-term rates by itself.[2]

[2]Gurkaynak (2005) finds that what he calls "timing surprises" — unexpected changes in the

Thus the ability of central banks to influence expenditure, and hence pricing, decisions is critically dependent upon their ability to influence market expectations regarding the *future path* of overnight interest rates, and not merely their current level. Better information on the part of market participants about central-bank actions and intentions should increase the degree to which central-bank policy decisions can actually affect these expectations, and so increase the effectiveness of monetary stabilization policy. Insofar as the significance of current developments for future policy are clear to the private sector, markets can to a large extent "do the central bank's work for it," in that the actual changes in overnight rates required to achieve the desired changes in incentives can be much more modest when expected future rates move as well.[3]

Thus the public's understanding, not only of what the central bank is currently doing, but of what it can be expected to do in the future, is critical for the effectiveness of policy. It might nonetheless be argued that it should be enough for a central bank to systematically *follow* a sound policy, without also needing to explain it to the public. If one assumes rational expectations on the part of the public, it would follow that any systematic pattern in the way that policy is conducted should be correctly inferred from the bank's observed behavior. Yet while it would be unwise to choose a policy the success of which depends on its *not* being understood by the public — which is the reason for choosing a policy rule that is associated with a desirable rational-expectations equilibrium — it is at the same time prudent not to rely too

current federal funds rate operating target that do not involve any change in market expectations regarding what the funds rate target will be after the *next* meeting, as when a change in the target that was already expected occurs sooner than some had expected it — have little effect on either bond yields or equity prices, while FOMC post-meeting statements that change expectations regarding the future path of the funds rate have significant effects on both.

[3]There is evidence that this is already happening, as a result both of greater sophistication on the part of financial markets and greater transparency on the part of central banks, the two developing in a sort of symbiosis with one another. Blinder *et al.* (2001, p. 8) argue that in the period from early 1996 through the middle of 1999, one could observe the U.S. bond market moving in response to macroeconomic developments that helped to stabilize the economy, despite relatively little change in the level of the federal funds rate, and suggest that this reflected an improvement in the bond market's ability to forecast Fed actions before they occur. Statistical evidence of increased forecastability of Fed policy by the markets is provided by Lange *et al.* (2001), who show that the ability of Treasury bill yields to predict changes in the federal funds rate some months in advance has increased since the late 1980s.

heavily on the assumption that the public will understand policy perfectly regardless of the efforts that are made to explain it. Insofar as explanation of the policy rule to the public does no harm under the assumption of rational expectations, but improves outcomes under the (more realistic) assumption that a correct understanding of the central bank's policy commitments does not occur automatically, then it is clearly desirable for the central bank to explain the rule that it follows.[4]

The advantages of a public target when the private sector must otherwise forecast future policy by extrapolating from experience are shown in a recent analysis by Orphanides and Williams (2005). In the Orphanides-Williams model, private agents forecast inflation using a linear regression model, the coefficients of which are constantly re-estimated using the most recent observations of inflation. The assumption of forecasting in this manner (on the basis of a finite time-window of historical observations) rather than a postulate of rational expectations worsens the tradeoff between inflation variability and output-gap variability that is available to the central bank.[5] Allowing inflation variations in response to "cost-push" shocks for the sake of output-gap stabilization is more costly than it would be under rational expectations, because temporary inflation fluctuations in response to the shocks can be misinterpreted as indicating different inflation objectives on the part of the central bank. Orphanides and Williams then show that a credible commitment to a long-run inflation target — so that private agents do not need to estimate the long-run average rate of inflation, but only the dynamics of transitory departures from it — allows substantially better stabilization outcomes, though still not quite as good as if private agents were to *fully* understand the equilibrium dynamics implied by the central bank's policy rule. This provides a nice example of theoretical support for the interpretation given by Mervyn

[4]King (2005b) proposes that it is more reasonable to expect the public to follow simple (but possibly fairly robust) "heuristics" in making decisions, of the kind discussed by Gigerenzer and Selten (2001), rather than behaving like the optimizing agents of economic theory. He argues that in this case central-bank communication can play an important role in leading people to choose heuristics of the right sort — *i.e.,* ones that lead to greater macroeconomic stability.

[5]Eusepi (2005) finds in the context of a model with more detailed microfoundations that requiring private agents to learn equilibrium patterns of fluctuations in inflation and the output gap by estimating atheoretical regressions can lead to instability of the learning dynamics and to persistent fluctuations driven by learning dynamics; transparency about the form of the central bank's policy rule (so that agents can estimate a correctly specified structural equation instead of a reduced-form econometric model) instead favors stability of the learning dynamics.

King (2005a) and others of practical experience with inflation targeting, which is that tighter anchoring of the public's inflation expecations has made possible greater stability of *both* real activity and inflation.

Nor is there any reason to suppose that it suffices for a central bank to make clear the long-run average inflation rate that it intends to maintain, while allowing the public to reach its own conclusions about the nature of transitory departures of the inflation rate from that long-run average. It is certainly true that anchoring expectations about the long-run average inflation rate is important, and that in itself is an important accomplishment. But the analysis of Orphanides and Williams also shows that even when private agents know the long-run average, but have to estimate the dynamics of transitory departures from it, the available tradeoff between inflation stabilization and output-gap stabilization is less favorable than it would be under rational expectations, *i.e.,* than it would be if one could rely on a correct understanding of the transitory dynamics. Thus there are in principle gains from an explicit commitment regarding this aspect of policy as well, and not simply trusting that people will be able to observe the pattern in one's behavior.

There is also a further, somewhat subtler, reason why explicit commitment to a target or policy rule is desirable, given the forward-looking behavior of the people in the economy that one seeks to stabilize. Even if one supposes that the private sector will fully understand whatever approach to policy the central bank takes, regardless of what it says about it, a public commitment to a rule can help policymakers to conduct policy in a way that achieves better outcomes. For is not enough that a central bank have sound objectives (reflecting a correct analysis of social welfare), that it make policy in a systematic way, using a correct model of the economy and a staff that is well-trained in numerical optimization, and that all this be explained thoroughly to the public. A bank that approaches its problem as one of optimization under *discretion* — deciding afresh on the best action in each decision cycle, with no commitment regarding future actions except that they will be the ones that seem best in whatever circumstances may arise — can still obtain a substantially worse outcome, from the point of view of its own objectives, than one that commits itself to follow a properly chosen policy *rule.* As Kydland and Prescott (1977) first showed, this can occur even when the central bank has a correct quantitative model of the policy tradeoffs that it faces at each point in time, and the private sector has correct expectations about the way that policy will be conducted.

At first thought, discretionary optimization might seem exactly what one would want an enlightened central bank to do. All sorts of unexpected events constantly occur that affect the determination of inflation and real activity, and it is not hard to see that, in general, the optimal level of interest rates at any point in time should depend on precisely what has occurred. It is plainly easiest, as a practical matter, to arrange for such complex state-dependence of policy by having the instrument setting at a given point in time be determined only after the unexpected shocks have already been observed. Furthermore, it might seem that the dynamic programming approach to the solution of intertemporal optimization problems provides justification for an approach in which a planning problem is reduced to a series of independent choices at each of a succession of decision dates.

But standard dynamic programming methods are valid only for the optimal control of a system that evolves mechanically in response to the current action of the controller. The problem of monetary stabilization policy is of a different sort, in that the consequences of the central bank's actions depend not only upon the sequence of instrument settings up until the present time, but also upon private-sector expectations regarding future policy. In such a case, sequential (discretionary) optimization leads to a sub-optimal outcome because at each decision point, prior expectations are taken as *given,* rather than as something that can be affected by policy. Nonetheless, the predictable character of the central bank's decisions, taken from this point of view, do determine the (endogenous) expectations of the private sector at earlier dates, under the hypothesis of rational expectations; a commitment to behave differently, that is made credible to the private sector, could shape those expectations in a different way, and because expectations matter for the determination of the variables that the central bank cares about, in general outcomes can be improved through shrewd use of this opportunity. This is illustrated concretely in section 2, when I discuss the way in which policy should be conducted when the lower bound on short-term nominal interest rates constrains the way that policy can be conducted.

In general, the most effective policy (the best outcome, from among the set of possible rational-expectations equilibria) requires that policy be conducted in a *history-dependent* way, so that policy at any time depends not only on conditions then (and what it is considered possible to achieve from then on), but also on *past* conditions, even though these no longer constrain what it is possible to achieve in the present. While there is no benefit, at the time, from conducting policy in a way that is condi-

tioned by the past, the *anticipation* that one would do so, at an earlier date, can have important beneficial effects on what policy can achieve at the earlier date. These benefits can make the subsequent losses worthwhile, as the example in the next section shows.

It is furthermore desirable, not simply that a central bank have a private intention of this sort, but that it be *publicly* committed to such a target. First, a public commitment is likely to make it easier for the central bank's policy deliberations to remain focused on the right criterion — the one with the property that systematic conformity to it leads to an optimal equilibrium — rather than being tempted to "let bygones be bygones." And second, the benefits associated with commitment to a history-dependent policy depend entirely on this aspect of policy being *anticipated* by the private sector; otherwise, it *would* be rational to "let bygones be bygones." There is no point to a secret commitment to the future conduct of policy in accordance with a history-dependent rule, while the private sector continues to believe that the central bank will act in a purely forward-looking fashion; thus the target should be explained as clearly as possible to the public, and shown to be guiding the bank's decisions.

1.2 Communication About What?

Which specific types of communication by central banks are most important, in light of the objectives discussed above? It is possible to distinguish among at least four broad classes of issues, about which a central bank may consider revealing more or less to the public. The first is the central bank's interpretation of *economic conditions,* including (perhaps) the central bank's view of the outlook for the future, to the extent that this is shaped by factors other than the bank's intentions with regard to policy. Central banks typically have large staffs devoted to collecting and analyzing information about current conditions in the economy, as an input into policy deliberations; and the accuracy of private-sector understanding of the state of the economy might be improved if the central bank were to reveal more about what it believes it has learned. A second topic is the content of the *policy decisions* that are made in the central bank about current operating targets. For example, as noted in the introduction, the Fed did not publicly confirm the existence of an operating target for the federal funds rate prior to 1994, whereas current practice is to release a statement immediately following each meeting of the FOMC, which, among other things,

announces the operating target agreed upon at that meeting. A third possible kind of communication would be a description (which might be more or less explicit) of the *strategy* that guides the central bank's policy decisions in general. A fourth type of communication, much debated in the U.S. at present, makes statements about the *outlook for future policy*, in light of the current situation, without necessarily asserting that this illustrates a general rule that will always be followed.

These are all types of communication in which the public might be interested, and a general commitment to increased "transparency" might be taken to require greater explicitness about all of these matters. But the way in which " transparency" about one or another of these matters relates to the goal of more effective stabilization policy is somewhat different in each case. The first two types of communication are the ones that are least controversial among central bankers;[6] to the extent that there are doubts about the desirability of saying more about the central bank's analysis of current conditions, for example, this is largely connected to the way that the public may use this information to make inferences (rightly or wrongly) about the bank's intentions regarding future policy. And it is in any event the effect of central-bank talk on the public's expectations regarding future policy that is critical for the concerns introduced above. Hence it is communication about the way in which policy should be conducted in the future (the third and fourth types of communication listed above) about which I wish to speak here.

One might, first of all, make statements about the *targets* or objectives that future policy decisions will aim to achieve; ideally, one might imagine a full description of a *policy rule* to which the policy committee intends to conform. This is the ideal suggested by the theoretical literature, on the basis of the considerations summarized above. On the one hand, private-sector decisions depend, in principle, not just on near-term expectations, but on the expected state-contingent evolution of the economy far into the future, and not just on what is most likely to happen, but on how the economy will evolve under all possible future contingencies; and one could only hope to communicate about what should happen in *all* of the relevant future states through a discussion of the bank's general strategy. Moreover, an optimal policy requires that the central bank commit itself to behave in a different way than would correspond to discretionary optimization. It is difficult to imagine institutionalizing such conduct other than through a conscious commitment to a particular strategy inside the

[6]Note, however, some qualifications to this in section 1.3 below.

central bank itself; and if such a conscious intention exists, a public statement of the commitment is likely to help the policy committee to remember its intention.

But what does any of this have to do with communication policy? The public commitment of a central bank to particular targets or to a particular policy rule will not be matters for routine, ongoing communication with the public that requires institutionalization. It is true that from time to time it will be appropriate to change the targets — as, for example, in the case of the change in the U.K. announced in December 2003, from an RPIX target of 2.5 percent per annum to a CPI target of 2.0 percent — but announcements of this kind are not what is generally understood by "communication policy." Would communication policy be important, then, for a central bank that was actually able to commit itself to a sensible policy strategy?

There are two reasons why it surely would be. The first is the need for *verifiability* of the central bank's commitment. One might imagine that the central bank's seriousness about its declared targets could be ensured by checking whether they are met, without requiring the bank to say anything about *how* it ensures that they are met. For example, under a rumor that was widespread at one time, accountability was ensured in New Zealand by a "contract" with the Governor of the RBNZ according to which the Governor could be fired if realized inflation ever went outside a certain band. In practice, however, it makes more sense to monitor the existence of good-faith efforts to achieve the bank's targets than to suppose that one can demand that the targets will actually be fulfilled at all times; and this will require communication by the central bank about the rationale for its policy decisions. Moreover, to the extent that optimal target criteria involve the expected paths of variables that cannot yet be directly measured, as is typically the case, it is appropriate to check, not whether the *actually realized* values satisfy the target criterion, but whether it would have been reasonable for the central bank to *expect* them to satisfy the criterion at the time of its policy decision. This requires the central bank to discuss the projections on the basis of which the policy decision was made.

The second reason is that in practice, the strategy that a sensible central bank follows (and may wish to be understood to follow) will be too complex to explain through any one-time official statement of its "policy rule." On the one hand, the set of contingencies that may arise (and matter substantially for policy if they do) are extremely various. As a consequence, an explicit rule of conduct (one specific enough to indicate unambiguously the instrument setting appropriate to any given

circumstances) would either contain too many provisos to actually be written down, or would deal in a grossly inadequate way with the situations actually encountered with some frequency.[7] Moreover, as Bank of England Governor Mervyn King (2005b) has stressed in his recent Mais Lecture, the central bank's understanding of the monetary transmission mechanism will surely continue to evolve; but this means that an explicit rule that was judged to be optimal on the basis of the bank's preferred model of the economy in one year would surely no longer be judged optimal from the point of view of the bank's best understanding a few years later.[8]

These considerations are sometimes taken to imply that the very idea of advance commitment to a policy strategy is impractical, and that the only sound approach will be something close to pure discretion. I do not believe so; nor is that the point of Governor King's discussion. What they do imply, however, is that in practice, the kind of commitment that it is sensible for a central bank to make in advance, and the kind that it is sensible for it to try to explain to the public, is a commitment to a general *strategy*, with the implications of this strategy for the precise instrument settings that will be appropriate under particular circumstances left to be determined when it is known which circumstances have arisen. Similarly, the general strategy should be one to which the bank can expect to adhere even as its views about the details of the monetary transmission mechanism change, though its current best guess about those details will play an important role in deliberations about the particular actions that will best implement the general strategy.

The general strategy to which the bank commits itself nonetheless can and should be more specific than a mere promise to do "whatever best serves social welfare" in whatever circumstances have arisen, and it should require a different approach to policy than the one that would be chosen by a discretionary optimizer. For example, it should bring about a lower average rate of inflation than would result (according to most plausible economic models) from discretionary optimization; and it should require departures of the inflation rate from that long-run average to be less persistent on average than would most likely result from discretionary optimization as well.[9]

[7]The point here is that while any one "special" situation with which the rule does not deal might be highly unlikely *ex ante*, the central bank would likely face *some* such situation quite often, as the number of possible "special" situations is so large.

[8]Charles Goodhart has also stressed (in private communication) that a monetary policy committee's views will inevitably change over time with changes in the composition of the committee.

[9]It should also allow policy to be history-dependent in a way that discretionary optimization is

And in order for the benefits of these aspects of the bank's strategy to be obtained, the consequences of this commitment for the economy's likely future evolution must be made clear to the public, at least to the extent that this is possible given the uncertainty faced by the central bank itself.

Because the way in which the strategy will be implemented in practice cannot be reduced to an explicit instrument rule, ongoing communication on the part of the central bank can play an important role in clarifying the consequences of the general strategy. It can be very useful to *illustrate* the consequences of the bank's approach to policy deliberations by showing how the particular situations that have already arisen were analyzed; over time, the observation of a sufficient number of such cases should help the private sector to some degree of understanding of the central bank's "reaction function." (The chances of this occurring, of course, are vastly greater in the case that the bank does itself seek to base its decisions on a stable set of principles, despite the varying kinds of information that are considered on different occasions, and in the case that it explains its decisions each time by reference to those same principles.) But this approach to making public the systematic character of policy will depend on a commitment to frequent communication about ongoing policy deliberations within the bank. Ideally, such communication will be regular, detailed, and structured, as in the case of the Inflation Reports of the inflation-forecast targeting central banks (discussed further in section 3).

A somewhat different way in which central-bank talk can convey information about future policy is through direct statements about the current outlook for policy. Such statements — illustrated by the comments that have recently been included in the post-meeting statements of the FOMC in the U.S. (discussed further in section 2) — fall considerably short of stating a general rule for the conduct of policy, and are likely to refer only to future policy over a fairly short horizon. They can, however, be much more specific about matters such as how the policy instrument will be set than descriptions of the bank's general strategy are likely to be.

There are a number of reasons why this kind of communication can also be useful. First of all, it might be used to some extent as a *substitute* for communication about a general strategy, for those central banks that are reluctant to commit themselves to any target or strategy in general, but may nonetheless be willing to commit themselves occasionally to an *ad hoc* departure from fully discretionary policy. In fact,

not; this is discussed further below.

communication of this kind has been used most notably thus far by central banks such as the U.S. Federal Reserve and the Bank of Japan,[10] which have *not* been willing to explicit themselves to quantitative inflation targets.

But discussion of the outlook for policy is likely to play a useful role even in the case of a bank that is as explicit about its general commitments as it is likely to be possible for any bank to be. As just discussed, in practice it would not be possible for a bank to commit itself to an explicit instrument rule. There would thus be considerable room to give further information about the likely path of the policy instrument on particular occasions, that would neither contradict nor be made unnecessary by the bank's commitment to its general strategy. Such communication would help to flesh out the concrete implications of the general strategy, and increase the ability of the private sector to make correct inferences about the consequences of the bank's commitments for the future evolution of the economy. This kind of amplification of the general strategy is likely to be especially useful when unusual circumstances arise, so that the implications of the strategy for circumstances of that kind might not be at all apparent simply from observation of the bank's past behavior. The situation discussed in the next section — where the Federal Reserve found itself at least potentially constrained by the zero lower bound on nominal interest rates, though this constraint had been irrelevant for more than fifty years — provides a good example of such an occasion.

A further argument for the desirability of communication about the outlook for future policy — and one in no way tied to unusual circumstances — follows from the history-dependence of an optimal policy commitment. Optimal policy requires not only that the central bank commit itself to a particular rule of conduct, but that the rule be history-dependent: it must take account of past conditions, even some that no longer matter for an evaluation of what it would be possible to achieve from now on. Hence any institutionalization of an optimal rule must involve keeping some record of past conditions. It is furthermore worth noting that what matters is not what the past was actually like (as viewed from the future, when the history-dependent policy action is to be taken), but how matters *appeared* then, as this is what would determine the value at the earlier time of being able to shift expectations regarding

[10] I discuss recent policy signaling by the Fed in section 2. On recent policy signaling by the Bank of Japan, see Bernanke *et al.* (2004), Fujiki and Shiratsuka (2002), Iwamura *et al.* (2004) and Oda and Ueda (2005).

future policy.[11] Thus implementation of an optimal policy requires that a record be kept of how matters appeared to the policy committee in the past, and that those past views condition the later policy decision. And while history-dependent policy requires only that there exist an internal record, the *benefits* of history-dependence depend on its being understood by the public; this makes a public statement about the aspects of the current situation that should change future policy deliberations appropriate.

There are various ways in which the relevant aspects of past deliberations might be encoded, and in which those records might be used in subsequent deliberations. But one fairly straightforward one — which would make it especially easy for the public to understand the consequences for future policy — would be to indicate at the earlier date the future policy that should be expected to be implemented later, in the absence of developments unforeseen at that time. The policy committee would then be committed to actually implement the policy announced earlier, unless circumstances changed in ways not previously foreseen. Deciding policy in advance (to this extent) would be an obvious way of allowing the policy committee to internalize the effects of anticipations of its later policy, and making public the committee's forecast of future policy would be an obvious way of making clear the expectations regarding future policy that should follow from the intention to make policy history-dependent. Of course, in order to prevent such an advance commitment from implying a non-state-contingent (and hence suboptimal) rule of conduct, it would be important to specify the assumptions regarding economic developments under which the forecast about future policy had been made, so that the nature of the contingency of the commitment would be clear.

1.3 Can a Central Bank Talk Too Much?

There are nonetheless a number of questions that may be raised about the desirability of central-bank communication, especially in the case of communication about future policy intentions. One point of view — once fairly common among central bankers,

[11] In optimal policy calculations like the ones discussed in the next section, the history-dependence of optimal policy results from the presence of lagged Lagrange multipliers in the first-order conditions that characterize the optimal state-contingent evolution of the economy. The lagged values of these Lagrange multipliers depend on the decision problem faced by the central bank *at its last decision point*. For further discussion, see Woodford (2003, chap. 7) and Svensson and Woodford (2005).

though less common now — would question whether it is actually desirable to increase the degree of precision with which the markets are able to anticipate the actions of the central bank, arguing that market interventions by the central bank will be more *effective* to the extent that the bank is able to *surprise* the markets. The idea, essentially, is that unanticipated trading by the central bank should move market rates by more, owing to the imperfect liquidity of the markets. Instead, if traders are widely able to anticipate the central bank's trades in advance, a larger number of counter-parties should be available to trade with the bank, so that a smaller change in the market price will be required in order for the market to absorb a given change in the supply of a particular instrument.

But such an analysis assumes that the central bank better achieves its objectives by being able to move market yields more, even if it does so by exploiting temporary illiquidity of the markets. Yet the temporarily greater movement in market prices that is so obtained — if any greater movement *is* obtained[12] — occurs only because these prices are temporarily less well coupled to decisions being made outside the financial markets. Hence it is not at all obvious that any actual increase in the effect of the central bank's action upon the economy – upon the things that are actually relevant to the bank's stabilization goals – can be purchased in this way.[13]

Another ground for caution about the amount that the central bank should say about its view of the future is provided by the analysis of Morris and Shin (2002) of the possible disadvantages of public information provision. Morris and Shin consider a stylized game in which individual market participants each choose an action on the basis of their observation of both a public signal (common knowledge to all market participants) and a private signal. Both the public signal and the private signal are noisy measures of some payoff-relevant "fundamental" state variable; each market

[12]Demiralp and Jorda (2002) find that it has been possible for the Fed to move the funds rate with a *smaller* quantity of open-market operations since 1994 than before, and interpret this as an effect of the FOMC's greater transparency about its funds rate target since 1994. This would suggest that advance signaling of what the Fed wishes to achieve makes it *easier* for the Fed to move interest rates where it wishes them to be, contrary to the argument mentioned in the previous paragraph. There is a simple reason why this is likely to be the case, namely, intertemporal substitution in the demand for federal funds as a result of the fact that reserve requirements require only a certain average level of reserves over a two-week maintenance period.

[13]I develop this point in more detail in Woodford (2001), where a simple model of policy effectiveness with incomplete market participation is presented.

participant is assumed to care both about choosing an action that is appropriate given the fundamental state of the economy, and about choosing an action that is not too different from others' actions. (The latter aspect of the assumed payoffs in the game gives it some of the characteristics of Keynes' famous "beauty contest".) Morris and Shin show that while an increase in the precision of the private information available to market participants necessarily increases social welfare, an increase in the precision of the *public* signal has an ambiguous effect on welfare. On the one hand, it will increase the accuracy of each market participant's assessment of the current state of the economy, with the result that equilibrium actions are on average more appropriate to current fundamentals. But on the other hand, it will reduce the weight that each market participant puts on her private information in forming her estimate of current conditions and hence in choosing her action, and increase the weight placed on the public signal instead. This second effect makes the average action less appropriate to the current state, because the error in the public signal affects everyone's action (while the errors in participants' private signals instead cancel out, and have no effect on the average action). It is possible for the second effect to outweigh the first, so that welfare is reduced by an increase in the precision of the public signal — that is, by an increase in the amount of information conveyed by it.[14]

Morris and Shin stress that a leading application of their analysis should be to the question whether increased transparency on the part of a central bank is necessarily a good thing,[15] and their argument has received a great deal of attention in central banks and in the financial press,[16] often in the context of discussions of the desirability of the kind of signaling of future policy described in section 2 of this paper. However, the applicability of their analysis to this kind of central-bank communication is far from obvious.

[14]The decision to release more information is represented in the Morris-Shin model by the release of a signal that is a less noisy measure of the fundamental state. Their conclusion that under some circumstances it may be better for a central bank to say less is in fact a finding that under certain circumstances it would increase social welfare for the central bank to release estimates of the state of the economy *that contain more random noise*. Stating the conclusion this way would make it seem more paradoxical; but this is actually what their formal analysis implies.

[15]The application of the Morris-Shin insight to the issue of the desirable amount of central-bank communication is developed especially in Amato, Morris and Shin (2002) and Amato and Shin (2003).

[16]See, for example, the discussion in the *Economist* (2004), and by Kohn (2005) and Issing (2005).

It is important to recognize that while Morris and Shin show that a (small) increase in the precision of the information released by a public authority *can* be welfare-reducing under some circumstances, the conditions under which this result is obtained are quite special. First of all, the perverse outcome requires that the central bank's announcement not be too accurate as an indicator of the "fundamental" in question, while market participants' private information about that same state variable must be sufficiently precise. In fact, as Svensson (2005) points out, in the Morris-Shin model the precision of each participant's private information must be *at least* 8 times as great as the precision of the public signal in order for the perverse outcome to be possible.[17] And it is not obvious that private information should be so much superior to the information that would be revealed by a central bank that makes an effort to tell what it knows.

Of course, central banks are themselves less than omniscient, and one argument within central banks for limiting the amount that is said to the public is the straightforward observation that the central bank *does not know* the answers to all of the questions about which market participants would wish to be informed. But the Morris-Shin result requires not only that the central bank's assessment of the fundamental, should it choose to reveal it, would not be God's truth; it requires that it be much *less accurate* on average than the estimate that any market participant would make on her own, in the absence of comment by the central bank. Even in the case of an assessment of economic conditions that are largely outside the control of the central bank, it is not plausible that a central bank's guess should be this bad. For example, Romer and Romer (2000) find that Federal Reserve Board staff forecasts compare favorably with the accuracy of even the most sophisticated private forecasters' forecasts, and this is hardly surprising given the size of the Fed staff and its privileged access to certain kinds of information.

And the assumption about relative accuracy required for a perverse result is least plausible of all in the case of central-bank communication about *likely future policy*. If there is one issue about which a central bank should have better information than that of market participants, it is the bank's own deliberations about matters (such as the path of the federal funds rate) that are essentially under its direct control.[18]

[17]The minimum required ratio depends on the parameter r of the loss function (1.1) below, but it is always 8 or higher; see the Appendix for details.

[18]Governor Kohn (2005) suggests that the danger identified by Morris and Shin applies even more

Again, it is often objected that even if the federal funds rate is completely subject to decision by the FOMC, this does not mean that the FOMC already knows what it will decide about where it will want the funds rate to be next year; and it is suggested that an inaccurate forecast may be worse than none at all. But even if the Committee does not yet know the precise answer to questions such as when the current series of quarter-point increases in the funds rate target will end, it surely is in a better position than outsiders to make an informed guess. It is therefore implausible that a public revelation of its best guess about this matter could reduce welfare on the grounds proposed by Morris and Shin.

A second requirement for the perverse result in the Morris-Shin model is that the game played by market participants must have elements of a "beauty contest". Market participants must care, not just about acting in a way that conforms as much as possible with current fundamentals, but also about acting similarly to the way that others do; and they must care *sufficiently strongly* about conformity relative to their concern with fundamentals. Specifically, Morris and Shin assume a game in which each player i wishes to minimize the expected value of a loss function

$$L^i \equiv (1-r)(a_i - \theta)^2 + rE_j(a_j - a_i)^2, \tag{1.1}$$

where a_j is the action of a generic player j, θ is the unknown value of the "fundamental" state, E_j denotes an average over the continuum of players indexed by j, and $0 < r < 1$ is the relative weight on the conformity objective. The possibility of a perverse result requires not only the presence of the second term in the loss function, but that $r > 1/2$, so that the weight on the second term is greater than the weight on the first; if $r \leq 1/2$, then an increase in the precision of the public signal raises welfare regardless of what one may assume about the relative precisions of the public and private signals.

But again it is not obvious that one should assume that this is the relevant case where signals regarding the future path of interest rates are concerned. Is it really true that, holding constant a bond trader's estimate of the "fundamental" value of a

in the case of communication about the "policy inclination" than communication about the economic outlook, because markets are especially likely to pay great attention to what a central bank says about future policy. But in the model of Morris and Shin, market participants put greater weight on the public signal the greater the expected relative precision of that signal (*i.e.,* the greater the extent to which the public authority is believed to be relatively better informed), but the more this is true the stronger will be the relative strength of the *desirable* effect of increased transparency.

bond (based on both the central bank's hints about the future path of interest rates and his own information), the fact that other traders currently wish to buy the bond would make him wish to buy more of it? Might the second piece of information not instead make him think it is a good time to *sell* the bond that is overvalued on average, perhaps because average opinion has been influenced by views of the central bank that he disagrees with? If so, then the game among market participants would be characterized by "strategic substitutability" rather than "strategic complementarity," and, as Morris and Shin note in their article, the informational-externality argument would instead imply that market participants will put *less* weight on the public signal than would be socially optimal, rather than paying too much attention to the central bank's announcements.

Finally, the perverse result is possible in the model of Morris and Shin only because of a particular assumption about the proper measure of social welfare, that in fact is highly debatable. Recall that individual market participants are assumed to care about two distinct objectives — acting in a way that is appropriate given the fundamental state, and acting in the same way that others act — represented by the two terms in the loss function (1.1). But Morris and Shin rank alternative equilibria using a welfare criterion that reflects only one of these private objectives; they assume that public policy should seek to minimize a social loss function

$$L^{soc,MS} \equiv E_i[(a_i - \theta)^2], \tag{1.2}$$

the average squared distance of individual actions from the one that would be appropriate given the fundamental state θ. While individuals are assumed to dislike taking an action that differs from the actions taken by others, there is assumed to be no social welfare consequence of less coordination across the actions taken by different market participants.

It is not obvious, however, that this makes sense. The same factors that make individuals seek to avoid actions that are too far out of line with the actions of others may well imply that there are *social* losses from such lack of coordination. And the simplifying assumption made in the example of Morris and Shin is not innocuous. For the factor that they omit from their consideration of social welfare is one that necessarily favors greater precision of the public signal. An increase in the precision of the public signal will necessarily reduce the dispersion of individual market participants' actions, exactly because it leads them to put less weight on their

private information, which is the source of dispersion.

I show in the Appendix that if one were instead to rank outcomes on the basis of a social loss function proportional to $E_i[L^i]$ — that is, by the population average of the individual loss function (1.1) — then this alternative social loss function is *necessarily* reduced by increasing the precision of the public signal, even though the Morris-Shin loss function $L^{soc,MS}$ may be increased. As noted above, the only case in which (1.1) can be increased is when r is large; but this is exactly the case in which the goal of reducing the dispersion of opinion becomes the more important factor for social welfare, under the alternative proposed here. Since the objective proposed here seems the more reasonable one, I find little reason to be troubled by the Morris-Shin example, even when one grants the parametric assumptions required for their perverse case.[19]

Of course, central bankers may have other reasons to be concerned about saying too much about matters about which they are themselves uncertain. One of the reasons most often cited is a concern that members of the public could be harmed by reliance on bad information supplied by the central bank. But this would not be a concern if the central bank's audience could be assumed to consist of rational maximizers who optimally use the information available to them, as in the model of Morris and Shin.[20] And even granting that not all market participants can be relied upon to be quite this sophisticated in the way that they respond to news, it is not obvious that one should expect them to make fewer mistakes if left to puzzle things out for themselves. The fact that people are not ideal information processors means that a central bank should give thought to the question of what market participants most need to know and how best to express what it is trying to tell them; thus it needs to have a *communication strategy*, and not simply a concern for "transparency" in

[19]Roca (2005) obtains a similar result in the case of a model of price-setting under monopolistic competition of the kind discussed by Amato and Shin (2003), when a welfare objective is used that is based on the preferences of the households in the model. Hellwig (2004) similarly finds that transparency is welfare-increasing in an explicit model of complementarities in price-setting, while Angeletos and Pavan (2004) obtain a similar conclusion in a model of complementarities in investment.

[20]In the model of Morris and Shin, no *individual* is harmed by her observation of the public signal; if she were, she could choose to ignore it. Each market participant conditions her action on the public signal to precisely the extent that minimizes her expected losses, taking into account the likely error in the available sources of information.

the sense of letting anyone see whatever they wish. But it is not a reason for central banks not to try to increase the amount that is effectively communicated.

I do not wish to minimize the difficulties associated with effective communication with the public, especially about matters as subtle as the likely conduct of policy under future circumstances that are not yet known. In practice, communication strategies improve only through a process of trial and error, even when central banks give considerable attention to the problem of how to tell the public more; for market participants must learn to interpret what the central bank is saying, and the central bank must learn to anticipate how its statements will be interpreted. The remaining sections of this paper consider two recent case studies in which central banks have grappled with the question of how to talk about the outlook for future monetary policy.

2 Signaling Future Policy Near the Interest-Rate Lower Bound

A case in which the benefits of being able to steer expectations regarding the future conduct of policy, other than through current policy actions alone, are especially clear is in the case in which overnight interest rates are already as low as it is possible or desirable to make them, while underutilization of productive capacity and/or unduly low inflation continue to suggest a need for further monetary stimulus. Japan, where the overnight rate (call rate) has been at zero almost continuously since 1999, yet deflation continues, provides an obvious example of the possibility of such a situation. But the U.S. faced a similar situation, or at least the risk of one, in the spring and summer of 2003, as the federal funds rate operating target was reduced to only one percent, while the strength of the recovery remained doubtful and inflation remained lower than the Fed was entirely comfortable with.

Does monetary policy become impotent when the zero lower bound is reached, as classic analyses of the possibility of a "liquidity trap" in static models would suggest? Eggertsson and Woodford (2003) show that the answer is yes, in the context of an explicit general-equilibrium model of the transactions demand for money with sticky prices, if monetary policy is understood to consist solely of various ways in which the monetary base might be expanded through current open-market operations, without

any change in the way in which monetary policy is expected to be conducted in the future.[21] But they also show that changes in the expected conduct of monetary policy in the future — after real conditions change, so that the policy that would be preferred at the time, but for the need to fulfill earlier commitments, would be one in which nominal interest rates would be well above zero — can have a very substantial effect on inflation and real activity during the period in which the zero bound is a binding constraint. This indicates the possibility of substantial benefits from signaling that future policy will be conducted in a different way than might otherwise have been expected, simply as a result of the economy's having been temporarily constrained by the interest-rate lower bound.

2.1 An Optimal Policy Commitment when the Lower Bound Binds

It is worth recapitulating some of the details of the analysis of optimal policy by Eggertsson and Woodford (2003), as a basis for discussion of the recent use of communications policy in both the U.S. and Japan. The exposition is simplest if we proceed directly to a log-linear approximation to their intertemporal equilibrium model with Calvo-style staggered price-setting. In this approximation (which, except for the imposition of the zero bound, is identical to the one used in studies such as Clarida *et al.*, 1999), inflation π_t and the output gap x_t are determined by a pair of equations each period,[22]

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1}, \tag{2.1}$$

$$x_t = E_t x_{t+1} - \sigma[i_t - E_t \pi_{t+1} - r_t^n], \tag{2.2}$$

where κ, σ are positive coefficients, and $0 < \beta < 1$ is the utility discount factor, i_t is the riskless short-term (one-period) nominal interest rate, and r_t^n is the natural (real)

[21] This analysis extends the discussion of Krugman (1998) to include a more developed treatment of the dynamics of price adjustment, the connection between interest-rate policy and the generation of inflationary expectations, and the consequences of alternative forms of open-market operations.

[22] Equation (2.1) is here written without the "cost-push shock" term that plays a central role in the analysis of optimal policy in Clarida *et al.* The issue with which we are here concerned (the possible difficulties for policy created by the zero bound) is not one for which the existence of a "cost-push" term is important, whereas the existence of fluctuations in the natural rate of interest r_t^n is instead critical. The optimal policy rule derived by Eggertsson and Woodford (2003), however, is also optimal in the presence of "cost-push shocks" of the kind hypothesized by Clarida *et al.*

rate of interest, that evolves exogenously as a result of real disturbances. The interest rate i_t is assumed to be the instrument of monetary policy, and is here treated as under the direct control of the central bank. We may then suppress the equations of the model involving the demand for base money. However, it is important to note that the interest rate will satisfy

$$i_t \geq 0 \tag{2.3}$$

no matter how much base money is supplied; this lower bound is the constraint on policy with which we are here concerned.

Let us suppose that the objective of policy is to minimize a discounted loss function of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda x_t^2], \tag{2.4}$$

with some weight $\lambda > 0$. It follows from equations (2.1) – (2.2) that as long as $r_t^n \geq 0$ at all times, one possible rational-expectations equilibrium is one in which inflation and the output gap are both zero at all times, and in such a case, this is obviously the equilibrium that minimizes the loss function (2.4), and so is optimal. But it is possible for real disturbances to cause the natural rate of interest to be temporarily negative.[23] In such a case, the zero-inflation equilibrium is no longer a possibility. (Note that this equilibrium requires that $i_t = r_t^n$ at all times, so that (2.3) is satisfied only if $r_t^n \geq 0$ at all times.)

Given that the pursuit of zero inflation at all times would be optimal in the event that the lower bound on interest rates were never a problem, one might suppose that even given the possibility of an occasionally binding lower bound, it would be optimal to pursue zero inflation at all times, *if* the interest-rate lower bound allows it. But this is not true. As an illustration, consider the particular kind of real disturbance analyzed in the numerical example of Eggertsson and Woodford. The "normal" (long-run average) level for the natural rate of interest is $\bar{r} \equiv \beta^{-1} - 1 > 0$. However, at some date, an unexpected disturbance temporarily lowers the natural rate to a level $\underline{r} < 0$. There is then a probability $0 < p < 1$ each period that "fundamentals" revert to their normal state, so that $r_t^n = \bar{r}$ again, and in this case, the natural rate of interest is expected to equal $\bar{r}$ indefinitely (there are no further disturbances). With

[23] This is Krugman's (1998) analysis of the situation of the Japanese economy since the mid-1990s. See also Woodford (2003, chap. 4) for discussion of some of the kinds of real factors that can shift the natural rate of interest.

probability $1-p$, instead, the low-natural-rate state will continue in the following period, conditional on the natural rate remaining at $\underline{r}$ in the current period.

In this case, under the hypothesized policy, the central bank will achieve zero inflation from period T onward, where T is the random date at which fundamentals revert to their normal state. This will be associated with a constant output gap of zero, and a constant nominal interest rate equal to $\bar{r} > 0$. Prior to this date, inflation will equal the same rate $\underline{\pi}$ each period, the output gap will equal the same value $\underline{x}$, and the nominal interest rate will equal the same value $\underline{i}$,[24] where these constant values satisfy

$$\underline{\pi} = \kappa \underline{x} + \beta(1-p)\underline{\pi}, \tag{2.5}$$

$$\underline{x} = (1-p)\underline{x} + \sigma[\underline{r} + (1-p)\underline{\pi} - \underline{i}], \tag{2.6}$$

as a result of equations (2.1) – (2.2), together with the requirements that

$$\underline{\pi} \leq 0, \qquad \underline{i} \geq 0, \tag{2.7}$$

and that at least one of the inequalities in (2.7) must hold with equality. (The central bank achieves the zero inflation target at dates prior to T, unless the zero bound prevents the inflation rate from being raised to zero.)

An equilibrium of this form exists as long as

$$(1-p)\left(\beta + \frac{\kappa\sigma}{p}\right) < 1, \tag{2.8}$$

i.e., as long as the degree of persistence of the disturbance is not too great. One can easily show that under the assumption that $\underline{r} < 0$, it is the lower bound on interest rates that is binding in (2.7), and the solution is given by

$$\underline{\pi} = \frac{\underline{r}}{(\kappa\sigma)^{-1}p[1-\beta(1-p)] - (1-p)} < 0, \qquad \underline{x} = \frac{1-\beta(1-p)}{\kappa}\underline{\pi} < 0, \tag{2.9}$$

together with $\underline{i} = 0$. Thus deflation and output below the natural rate continue for as long as the disturbance to fundamentals does; even if prices are revised fairly often on average, the binding lower bound on interest rates can result in a slump that lasts for years. Furthermore, (2.9) implies that even a very mildly negative value for

[24] Here I restrict attention to the Markovian (minimum-state-variable) equilibrium consistent with the hypothesized policy. Note that when this equilibrium exists, it represents at least *one* possible outcome, and the fact that it may be very bad indicates the problem with this approach to policy.

the natural rate of interest can result in very severe deflation and contraction of real activity. Note that if the left-hand side of (2.8) is close enough to 1 (and there is no reason why it may not be), the rate of deflation and the size of the negative output gap in (2.9) become arbitrarily large, regardless of the degree to which $\underline{r}$ is less than zero.[25]

The contractionary effects may be quite large, due to a chain of circular causation. A real interest rate above the natural rate (owing to the zero bound) causes a negative output gap and deflation prior to period T. The anticipation of these effects, in the contingency that the natural rate continues to be negative in the following period, then depresses demand further and creates even stronger deflation in any period prior to T, owing to the effects of expectations $E_t\pi_{t+1} < 0, E_t x_{t+1} < 0$ in equations (2.1) – (2.2). The anticipation of these even stronger deflationary and contractionary effects causes still greater deflation and contraction, and so on, in a cumulative process that does not even converge unless (2.8) holds.

It is crucial in the above reasoning that the central bank is expected to target zero inflation again as soon as this becomes possible. While this would in fact make sense *ex post* — and so would be the outcome in a Markov equilibrium with discretionary optimization by the central bank — a better outcome is possible if the central bank commits to behave otherwise once fundamentals revert to their normal state. This can be seen if we repeat the above calculations, but assume that the central bank will bring about an inflation rate $\bar{\pi} > 0$ (and an associated output gap $\bar{x} > 0$) in period T. Under this variation on our assumptions, the solution for $\underline{\pi}$ in (2.9) generalizes to[26]

$$\underline{\pi} = \bar{\pi} + \frac{\underline{r} + \bar{\pi} + p\sigma^{-1}[\bar{x} - x^{ss}(\bar{\pi})]}{(\kappa\sigma)^{-1}p[1 - \beta(1-p)] - (1-p)}, \tag{2.10}$$

where $x^{ss}(\bar{\pi}) \equiv (1-\beta)\kappa^{-1}\bar{\pi}$ is the steady-state output gap associated with steady-state inflation rate $\bar{\pi}$.

In the case that $\bar{x}$ is increased along with $\bar{\pi}$ to the extent that it would in the case of a permanent commitment to the inflation target $\bar{\pi}$, the multiplier effect of an

[25] Of course, for a large enough rate of deflation and departure from the natural rate of output, the local approximations in (2.1) – (2.2) cease to be accurate; but this suffices to show that the departures from the zero-inflation steady state need not be *small*, for if they were small the local approximations would be valid and equations (2.9) would be approximately correct.

[26] Equation (2.10) applies as long as $\bar{\pi} \leq -\underline{r}$, so that the zero bound continues to bind when $r_t^n = \underline{r}$.

increase in the long-run inflation target $\bar{\pi}$ on the inflation rate $\bar{\pi}$ during the "liquidity trap" is given by

$$\frac{\partial \underline{\pi}}{\partial \bar{\pi}} = 1 + \mu,$$

where $\mu > 0$ is the (possibly very large) multiplier $-\partial \underline{\pi}/\partial \underline{r}$ implied by equation (2.9). There is a correspondingly large effect of a commitment to target an inflation rate $\bar{\pi} > 0$ on the value of $\underline{x}$ as well. Thus a commitment to a future inflationary policy can mitigate the effects of the zero lower bound, as argued by Krugman (1998). In the forward-looking model of inflation and output determination used here, these effects are quite large, owing to the same chain of circular causation as above, but now operating in the opposite direction (a "virtuous circle").

However, the optimal policy commitment (in order to minimize (2.4) is not a simple commitment to a higher long-run inflation target. The effects just discussed on inflation and output while the zero lower bound binds depend only on $\bar{\pi}$ and $\bar{x}$ being delivered in period T, the *first* period in which the natural rate of interest is again positive; there is no need to commit to *continued* inflation forever, and this will instead lead to unnecessary distortions in the long run. Eggertsson and Woodford (2003) show that the optimal policy involves a commitment to the creation of a modest inflationary boom in period T, and then stabilizing the price level shortly thereafter (*i.e.,* returning to a long-run inflation rate of zero), at a level slightly higher than the one that would have been reached in the absence of the disturbance.[27]

This is illustrated in Figure 1, which shows the paths of the nominal interest rate, the inflation rate, and the output gap under the optimal state-contingent policy commitment (for particular numerical parameter values discussed by Eggertsson and Woodford), in the case that T is exactly 15 quarters after the onset of the real disturbance. The figure also shows the paths of all three variables in the case of a commitment to zero inflation (or discretionary optimization). While the creation of the inflationary boom (by keeping interest rates low for five more quarters, rather than immediately raising them to the level that would be required to achieve price stability immediately at date T) results in mild distortions after date T, these are only temporary (as price stability and a zero output gap are achieved fairly soon),

[27] Jung *et al.* (2005) reach a similar conclusion in the case of different assumed dynamics for the natural rate of interest. Adam and Billi (2003) characterize optimal policy in the same model in the case of continuing stochastic fluctuations in the natural rate that cause the zero bound to bind periodically.

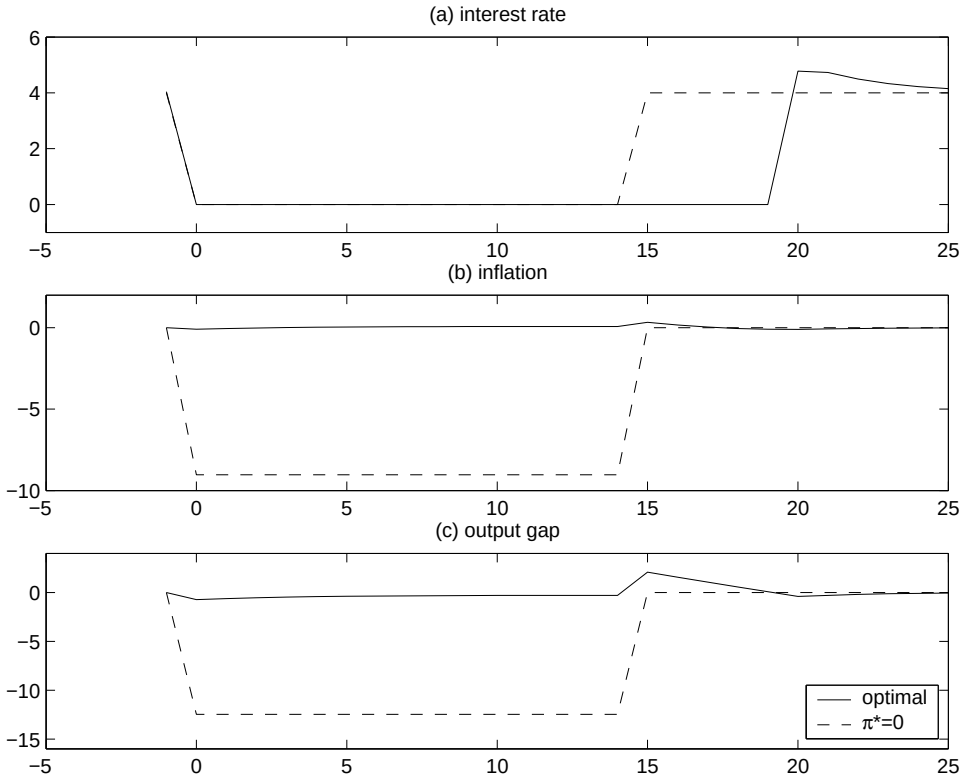


Figure 1: Comparison of state-contingent paths under two alternative monetary policies, in the case that the natural rate of interest is negative for 15 quarters. (Source: Eggertsson and Woodford, 2003)

and are quite mild relative to the size of the distortions prior to date T that are thereby avoided. The strong effect of the commitment to subsequent reflation of the economy occurs because of the chain of circular causation just discussed.[28]

This numerical example illustrates several points of more general importance. First, it shows how a credible commitment regarding the future conduct of policy can, at least in principle, greatly expand a central bank's ability to achieve its stabilization objectives. But in addition, it shows that the an optimal commitment requires not only that the central bank pledge to behave in a different way than would a dis-

[28]These effects are quite strong because it is assumed in the example that there is only a ten percent chance each quarter that fundamentals will revert to the normal state. Thus at any point in time while the natural rate of interest is negative, it is expected that this situation is likely to persist for two or more additional years.

cretionary optimizer; the conduct of policy must also be *history-dependent*. For the inflation rate that should be targeted once the natural rate of interest is positive is not the one that the central bank *always* targets in the case of fundamentals of that kind; rather, the central bank should temporarily behave differently because of what the economy's situation *has been* in the recent past, even though those circumstances no longer affect what it would be possible for policy to achieve from now on. As I have discussed above, this history-dependence of the optimal policy commitment strengthens the case for explicit discussion by the central bank of the way in which current conditions change the outlook for future policy. If it were desired simply to *always* target an inflation rate $\bar{\pi} > 0$, then it might not be necessary for the central bank to talk about this while in the liquidity trap; one might suppose that the central bank's long-run inflation target would already have been learned by the private sector from its previous behavior, and that people might confidently expect the central bank to return to the pursuit of this target once circumstances allowed it to be achieved, without any need for comment to that effect. But if it is desired that the public understand, while policy is constrained by the zero bound, that future policy will be *different* from what it usually is under similar circumstances, because of the current difficulties, then it is reasonable to suppose that the central bank may need to discuss this, rather than expecting this to be obvious from past experience. The case will be even stronger if the circumstances under which the zero bound becomes a constraint are fairly unusual.

It is also worth noting that the advantages of commitment to a history-dependent future policy do not depend on reaching the zero bound. It is simply important that there be *some* lower bound on the level of short-term nominal interest rates that the central bank is willing to target; none of the analysis just sketched depends on there being satiation in money balances when this bound is reached. In the analysis of Eggertsson and Woodford, there is no need for history-dependent policy unless there is some state in which the zero bound binds; but that is because they assume there is no *other* obstacle to lowering interest rates. If (as was arguably the case for the Fed in 2003) there is a positive level of interest rates i^l below which the central bank does not wish to go, a similar analysis applies in the case of this lower bound, except that now history-dependent policy becomes valuable if the natural rate of interest ever drops below i^l, which is even more likely to occur if i^l is positive.[29] Regardless of where the

[29]The analysis can similarly be generalized to the case of an objective function in which the

lower bound lies, a commitment to lower interest rates later can substitute, at least partially, for being able to lower interest rates immediately, so that history-dependent policy can relax the constraint implied by the lower bound.

2.2 Policy Signaling in Practice

The situation faced by the U.S. Federal Reserve in the summer of 2003 was arguably of the sort contemplated in the above analysis (though the model used in the calculations is obviously an extreme over-simplification). The federal funds rate operating target had been reduced to 1.0 percent by June of that year, and (at least according to speculation in financial markets and in the press) the FOMC may have been reluctant to move lower than that. Nonetheless, inflation remained low; according to the minutes of the August 2003 meeting of the FOMC, inflation was "already near the low end of what some members regarded as an acceptable range," and "a number of members expressed the view that some further disinflation was probable over the year ahead." While the committee was at least guardedly optimistic about real growth over the next year, it was believed that a substantial period of growth faster than the economy's potential growth rate would be needed to close "the economy's currently wide output gap."

Because of the risk of undesired further declines in inflation — that posed a particular risk insofar as once inflation *expectations* also fell, the level of real interest rates associated with the nominal interest rate floor would become an even higher one — the FOMC did not wish to tighten policy, despite the improving outlook for real activity. In this regard, it was not judged to be enough that they leave the funds rate target at one percent; for, as explained in section 1, it is primarily the *expected future path* of the funds rate (and other short rates tied fairly closely to it) that affects spending and pricing decisions, rather than the current level of the funds rate. And there was concern that the public's expectations regarding the future path of interest rates could move sharply upward as news about the real economy improved, because of the way in which the FOMC had typically responded to improvements in real activity in the past (as described, for example, by the "Taylor rule"). These

optimal inflation rate is assumed to be some positive inflation rate π^*, rather than zero. In this case, history-dependent policy is needed only if the natural rate of interest sometimes falls below the level $i^l - \pi^*$.

expectations, if allowed to respond in that apparently reasonable way, might slow the recovery of real activity and plunge the U.S. economy into deflation.

The minutes of the August meeting indicate the committee's concern with the recent evolution of market expectations, as indicated by long-term bond yields. The minutes discuss the "dramatic" increase in the ten-year Treasury yield in particular that had occurred in July (see Figure 2, below). "The increase appeared to be based on a number of factors, including investors' interpretation of the Chairman's congressional testimony, the release of Committee members' relatively bullish economic projections, and incoming news regarding the economy and corporate earnings that was seen as signaling a more likely upturn in economic growth," as a result of which the markets were evidently anticipating that increases in the funds rate might come as early as the fall. While the minutes do not clearly identify the reason for the FOMC's decision to introduce an explicit comment on the likely nature of future policy into its post-meeting statement on this occasion, it seems likely that their concern with movements in long-term bond yields on the basis of speculation about their future policy decisions was an important element in the decision to not leave the judgment of the market about this matter to guesswork.[30]

After reporting that the funds rate operating target would remain at 1.0 percent for another month, and assessing the "balance of risks" ("The Committee judges that, on balance, the risk of inflation becoming undesirably low is likely to be the predominant concern for the foreseeable future"), the statement included a final sentence of a new type: "In these circumstances, the Committee believes that policy accommodation can be maintained for a considerable period." While no outright commitment was made, the minutes indicate that the members in favor of this statement believed it likely "that the Committee would want to keep policy accommodative for a longer period than had been the practice in past periods of accelerating economic activity."

In fact, the committee acted as though they regarded themselves as committed not to raise rates, without some months of advance warning. The "considerable period" language was repeated in the statements released following each of the next three meetings as well (*i.e.,* through the end of 2003).[31] When the likelihood of

[30]Concern about the effectiveness of their communication strategy was also indicated by the fact that the committee scheduled a meeting for September, prior to the next policy decision, to reconsider "its practices regarding the communication of its policy decisions and its assessment of the risks to its objectives.".

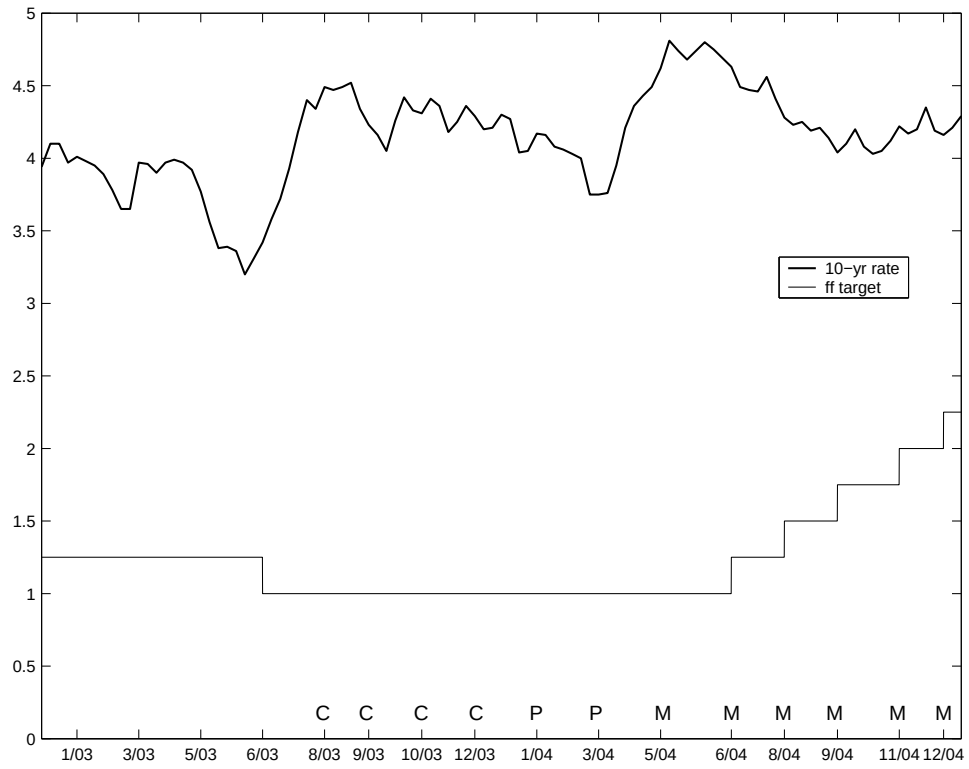


Figure 2: The federal funds rate target and the 10-year Treasury rate. Comments about likely future FOMC policy in post-meeting statements are indicated by letter codes: C = "considerable period", P = "patience", M = "measured pace."

interest-rate increases by the middle of 2004 became apparent, the fact that the funds rate would not remain at 1.0 percent indefinitely was indicated by dropping the "considerable period" language.[32] Instead, the final sentence of the statement released following the meeting at the end of January 2004 said that "the Committee believes that it can be patient in removing its policy accommodation." According to the minutes, "all the members agreed that a change in wording was desirable, not to signal a policy tightening move in the near term, but rather to increase the

[31]The dates on which statements have been released containing various types of language regarding likely future policy are indicated in Figure 2 by the codes 'C', 'P', and 'M'. The code 'C' indicates a reference to maintaining accommodation for a "considerable period."

[32]The commitment to keep rates low "for a considerable period" had already been qualified in the December 2003 statement, by the inclusion of a reference tying this policy to continued low inflation and resource "slack".

Committees flexibility to take such an action when it was deemed to be desirable and to underline that any such decision would be made on the basis of evolving economic conditions." This language was included again in the March 2004 statement, while the May 2004 statement instead indicated that "the Committee believes that policy accommodation can be removed at a pace that is likely to be measured." Even at this meeting, the funds rate target was not raised; but notice had been given that it *would* now be raised, albeit at a "measured pace." The funds rate target was indeed raised, beginning at the next meeting; it has now (at the time of writing) been raised by 2.25 percentage points, through a succession of quarter-point increases at nine successive meetings. Throughout this period of steady increases, the post-meeting statements have continued to include the reference to expecting to remove policy accommodation at a "measured pace."

What has the new policy of commenting on the likelihood of future changes in the funds rate target achieved? The "considerable period" language seems to have been intended to influence market expectations in a way that would stimulate additional spending — higher spending, that is, than would have occurred if expectations had been allowed to change in the direction that it was feared they would in the absence of such assurances from the FOMC. As in the scenario described by Eggertsson and Woodford, a statement that interest rates would be kept low for a longer period of time was able to substitute for an immediate cut in rates.[33] The signal furthermore seemed to be effective. Governor Ben Bernanke, speaking the following year, argued that "the language of the statement in August 2003 and subsequent meetings persuaded the markets that an autumn tightening was not in the cards, and market expectations adjusted accordingly. Crucially, this change in expectations resulted in lower interest rates at all maturities, a development that helped support the expansion in the latter part of last year" (Bernanke, 2004). The decline in the 10-year bond rate following the introduction of this commitment is shown in Figure 2. While bond yields move for a

[33]There is no indication in the minutes that the FOMC did not believe that interest rates could ever be cut below one percent. But for whatever reason, no further cut in the funds rate target was made, despite the desire to head off further disinflation, and to signal to the markets that policy would not be tightened as much as was widely believed. In discussing policy in the summer of 2003, Governor Donald Kohn speaks of a funds rate target of zero as having "uncertain consequences," and says that signaling about future policy instead "seemed to be the less-risky way" to stimulate demand (Kohn, 2005). This suggests that there was indeed a reluctance to cut the funds rate target, at least given the existence of an alternative lever of policy.

variety of reasons, this decline even as fears of deflation dissipated and the outlook for the real economy continued to improve suggests that the new communication policy had its intended effect on expectations.[34]

The eventual transition to a higher level for the funds rate has also been managed in a way that has involved few surprises for the markets, and that has not resulted in a bond market rout, sending long rates to levels greater than those consistent with the intended medium-term level of interest rates. As it became clear in the spring of 2004 that the commitment to maintain rates at their existing (historically low) level was about to end, with no indication from the Fed as to how dramatic a change in policy might be coming, long bond yields again rose fairly sharply. But the introduction of the commitment to a "measured pace" in May allowed long rates to decline again, as shown in Figure 2. Since then, it has been possible to increase the funds rate target by a total of two percentage points, while the 10-year bond rate remains (at the time of writing) at about the same level as it was late in 2003, and below the level that it had reached in the summer of 2003, prior to the introduction of the "considerable period" language. This sort of outcome is not especially paradoxical when increases in the funds rate target occur within the context of a funds rate path that was already fairly predictable, so that the increases themselves cause no change in market expectations about the likely level of interest rates over the next several years (and may actually contribute to *lower* expectations regarding nominal interest rates some years out, insofar as they confirm that the central bank will be vigilant to contain inflation). Bernanke (2004) suggests that the new communications policy has also likely contributed to a recent decline in "overall financial market volatility ... by reducing the uncertainty surrounding the future course of policy."

The degree to which recent decisions about the funds rate operating target have failed to surprise financial markets is shown in Figure 3. In the figure, bars indicate the size of the changes in the FOMC's funds rate operating target that occurred at various dates in the period 2001-04. Associated with each date at which there was a change (or at which there was a meeting at which the FOMC chose not to change the target, so that the announcement of no change was itself news) is an asterisk indicating the size (in basis points) of the component of this change that was

[34]Discussing the same period, Governor Kohn (2005) states "I would judge the outcome to have been successful. We did influence rates to better reflect the actual path of policy," with a good outcome for the economy as well.

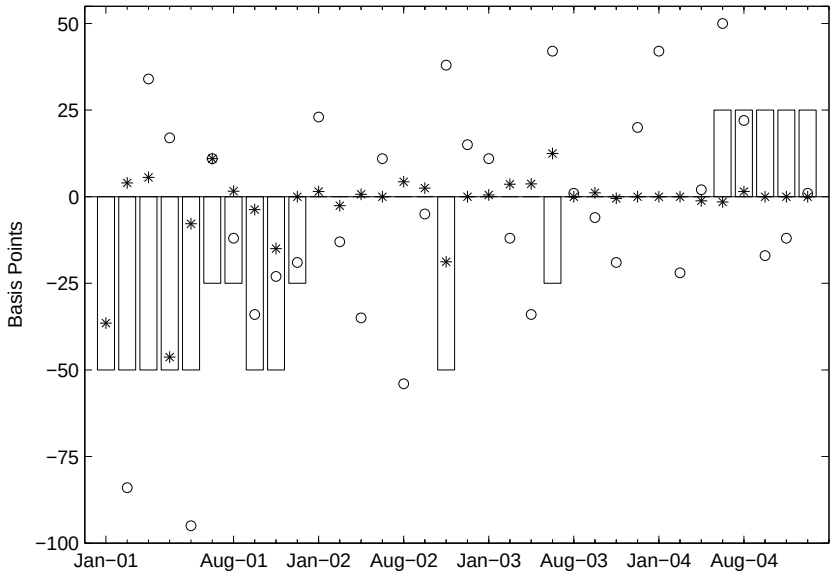


Figure 3: Federal funds rate target changes (indicated by bars), the surprise component of each change (indicated by asterisks), and change in the "slope factor" (indicated by circles). Sources: target changes: Federal Reserve Board; policy shocks: Gurkaynak *et al.* (2005); slope factor: Gurkaynak (2005).

not already anticipated by financial markets just before the announcement.[35] One observes that beginning with the August 2003 meeting, the surprise components of federal funds rate target announcements have been extremely small. Surprises were typically larger in earlier periods, even in periods when the target was left unchanged for several meetings in a row, and especially at times when a loosening or tightening cycle began, as in early 2001. Recently, instead, the surprises have been negligible, even when the Fed moved from its constant one percent target to a period of steady tightening.

At the same time, this does not mean that FOMC announcements have failed to affect the markets. In a recent paper, Gurkaynak (2005) investigates the effects of the

[35]The surprise components are measured by observing the change in federal funds futures prices for the following month that occurs over a one-hour time window around the announcement. The data on policy surprises are taken from Table 2 of the data appendix to Gurkaynak *et al.* (2005), available on the *IJCB* website.

release of FOMC statements on market expectations (as indicated by the fed funds futures market) regarding what the funds rate will be between the current meeting and the next, what it will be after the next meeting, and what it will be after the meeting after that. (The first of these changes in expectations is the "policy surprise" plotted in Figure 3.) Gurkaynak finds that while there has been almost no change in expectations regarding the current funds rate, there have still been changes in expectations regarding what the funds rate will be further in the future.

In particular, what he calls "slope surprises" (indicated by the circles in Figure 3[36]) — changes in expectations regarding the funds rate two meetings in the future, over and above what one would typically expect given the news about the current funds rate target and the one expected to be chosen at the next meeting — have continued to be about two-thirds as volatile since the fall of 2003 as they were previously, over the period (since early 1998) for which they can be measured.[37] This indicates that under the new regime, FOMC statements still change expectations regarding the *future path* of the funds rate — which, as argued above, is essential if they are to effect the economy — but they now achieve this *without* a need for surprise changes in the current funds rate target.[38] Insofar as the avoidance of unexpected movements in short-term rates is desirable, to the extent that it is possible without compromising other stabilization objectives, this can be judged an improvement in the skill with which monetary policy is conducted.[39]

[36]The values plotted for the slope surprises in Figure 3 are ten times as large as the definition used by Gurkaynak, to make the variations more visible in the figure.

[37]Over the period between January 1998 and June 2003, the root-mean-square size of the slope surprise (under the normalization used in Figure 3) was 0.53, while in the period between January 2001 and June 2003 (shown in the figure), it was closer to 0.54. Between August 2003 and December 2004 (the last FOMC statement in Gurkaynak's sample), the r.m.s. size has been 0.33 — a reduction in volatility, but only a modest one.

[38]In a related analysis, Gurkaynak *et al.* (2005) decompose changes in the term structure of interest rates upon the release of an FOMC statement into two (orthogonalized) principal components. They interpret the factor that is constructed to be uncorrelated with "policy surprises" as measure of changes in the expected future path of interest rates independent of any change in the current policy rate; as their figure indicates, the volatility of this "path factor" has also remained as large since August 2003 as it had been before.

[39]The observation that funds rate decisions are no longer surprises on the day of the meeting under the new regime contrasts with the findings of Ehrmann and Fratzscher (2005) about the effects of the Fed's change in disclosure policy in 1999. They find that funds rate decisions were *not* any more predictable on the day of the meeting under the regime where the FOMC released a

The primary difficulty associated with the new policy of signaling funds rate target changes in advance has been some degree of discomfort, among at least some members of the FOMC, with the degree to which statements of this kind constrain the policy decisions that can be made at later meetings.[40] Of course, one of the *advantages* of such statements that I have suggested above is precisely that they can constrain a policy committee to not behave in the way that would otherwise seem appropriate *ex post*; such constraint can be desirable, because of the benefits that flow from being anticipated to conduct policy in a history-dependent fashion. For example, in the case of the optimal state-contingent policy depicted in Figure 1, it would be important that the central bank feel itself constrained not to raise interest rates in quarter 15, despite the fact that it has learned that (relatively unexpectedly) the natural rate of interest has returned to its normal, positive level, and despite the fact that failing to do so means creating a mildly inflationary boom. Policymakers who reason like discretionary optimizers will indeed be uncomfortable with such behavior, and they would be right to anticipate — when considering the advisability of such a commitment at the earlier time — that they will subsequently find the commitment an annoying constraint. But a policymaker who thinks more deeply should realize that it is nonetheless desirable to constrain oneself, at least in the case that the constraint can be expected to be understood by the public.

But in all likelihood, the discomfort is not solely due to a failure to understand the logic of Kydland and Prescott (1977), but also to frustration with the crudeness of the kinds of commitments that can be made using such simple statements as the ones just paraphrased. The signals that have been given thus far through the post-

"balance of risks" statement than they had been previously, though they do find that the increased communication after 1999 *did* have the effect of increasing the degree to which funds rate decisions could already be forecasted by the markets immediately following the previous FOMC meeting (so that less learning occurred during intermeeting periods under the new regime), and also that FOMC meetings contributed less to financial market volatility under the new regime. Thus while the post-1999 "balance of risk" statements did increase the forecastability of policy (albeit not at such short horizons as a day in advance), the new policy signaling regime appears to be even more effective at this.

[40]For example, the minutes for the May 2004 meeting, at which the "measured pace" language was introduced, reports that "a number of policymakers were concerned that such an assertion could unduly constrain future adjustments to the stance of policy should the evidence emerging in coming months suggest that an appreciable firming would be appropriate," though all members eventually endorsed the language adopted.

meeting statements all attempt to say something about the likely path of the funds rate for the next several months; they refer (in a way that is open to interpretation) to rates of change and periods of time, but, except for the constant qualification that the statement is only an indication of "likely" policy, they do not speak of the way in which future policy should be *contingent* on circumstances that are not already evident. If the statements are interpreted as *commitments* to particular non-state-contingent paths for the funds rate — albeit commitments that specify the path only for a fairly short distance ahead — then they are likely to constrain policy in ways that are not fully ideal.

For while an optimal policy commitment will generally imply that policy should be *history-dependent,* as stressed above, it will also generally imply that policy should be *state-contingent* as well. For example, in the policy problem considered by Eggertsson and Woodford (2003), the optimal policy involves keeping the nominal interest rate at its lower bound for several quarters even after the natural rate returns to its normal level; but the *exact number* of quarters for which this should be done depends on the evolution of the exogenous disturbance, and is not known until date T is reached. Moreover, the fact that, in the particular numerical exercise considered in Figure 1, it is possible to make a definite commitment once period T is reached depends on the fact that there is assumed to be no further uncertainty about the evolution of fundamentals after date T. In the case of a more general disturbance process $\{r_t^n\}$, the number of periods for which interest rates should be kept low will also depend on the path of the natural rate *after* it has again become positive; under an optimal policy, the central bank would not generally know whether it was yet time to raise rates until the time to do so was reached.

The kind of commitment that needs to be communicated, in order to allow a closer approximation to fully optimal policy, is one that would indicate the way in which future policy should depend on future economic developments. This might seem so complex as to not be usefully explained to the public; but my own view is that even a very general indication of the kind of factors that should be crucial for future policy decisions would greatly help to clarify the public's view of the likely state-contingent evolution of interest rates and of the economy. In the context of the simple policy problem considered above, Eggertsson and Woodford show that the optimal time and degree to which interest rates should eventually be raised can be explained in terms

of a fairly simple formula. Under the policy rule that they propose,[41] the central bank should set its policy rate so as to achieve a particular pre-announced target level for an output-gap-adjusted price-level target,[42] to the extent that this is consistent with the interest-rate lower bound. The interest rate will thus be kept at the lower bound *as long as it continues to be impossible to reach (or exceed) this target,* even with interest rates at the lower bound; rates should be raised above the lower bound once the target level for the output-gap-adjusted price level is reached, and not sooner. A commitment of this kind would *imply* that the lower-bound policy should be expected to continue for a considerable period, in the case that the output-gap adjusted price level is currently well below the target; and a policy commitment that emphasized this target (and hence the size of the gap that would remain to be closed at any given time) would even allow the private sector some basis for judging the likely length of such a period. But it would also automatically imply that increases in interest rates would likely be appropriate soon, as the gap with respect to the target shrank, and would furthermore provide guidance as to how policy should be expected to be conducted thereafter, all without any need for a change in the language of the commitment. And it would be a form of commitment that would make evident the state-contingency of the implied path of interest rates.

The FOMC has made some effort to communicate the contingent nature of its statements about the outlook for future policy.[43] In December 2003, the statement that rates were likely to remain low for a "considerable period" was explicitly linked to the observation of "inflation quite low and resource use slack." According to the minutes of the January 2004 meeting, this language had been intended to "underscor[e] the notion that a move away from the current degree of policy accommodation would depend on economic conditions rather than simply on the passage of time."

[41] Other approaches to the implementation of the optimal state-contingent policy in a similar model are discussed by Svensson (2004) and by Sugo and Teranishi (2005).

[42] This variable is equal to the log price level plus a positive multiple of the output gap. Thus the price level target is automatically increased slightly in the event of a negative output gap.

[43] Bernanke *et al.* (2004) note that during the period of the FOMC's commitment to keep rates low for a "considerable period," each FOMC statement also discussed labor market conditions, which might have been interpreted as indicating that a change in the policy would be conditional upon labor market developments. As evidence in support of the conditionality have been understood by market participants, they find that Treasury yields became more responsive to news in the monthly payroll number in the employment report during this period.

These minutes also indicated that the language adopted in the statement released after the January meeting, referring to "patience" rather than to a "considerable period," was intended to indicate more clearly that an increase in interest rates later in the year was possible, "and to underline that any such decision would be made on the basis of evolving economic conditions." While the reference to "patience" itself may have been a rather cryptic way of indicating that policy would be state-contingent, the eventual publication of the minutes would have made this clearer; and the current policy of expedited release of the minutes makes the possibility of clarifying the state-contingent character of future policy through the minutes, rather than through the post-meeting statement itself, a realistic possibility.[44]

It would likely be desirable for the FOMC to experiment further with clarifications of this kind. As the period in which regular quarter-point interest-rate increases at each meeting were fully predictable comes to an end, it will become more important for the FOMC to attempt to communicate about the state-contingent character of policy, if it is to comment on future policy at all. Recent experience suggests that comments on future policy can help, both in reducing the number of policy surprises and in keeping bond-market expectations in line with the FOMC's own outlook for rates over the medium term, and these indications of success in steering expectations should increase the effectiveness of policy. But a continuation of this success under more normal circumstances will require the development of more flexible ways of speaking about the likely character of future policy.

3 Assumptions About Future Policy in Inflation-Forecast Targeting

Probably the most important advances in communications policy over the past ten to fifteen years have been made by the inflation-targeting central banks, among which banks such as the Sveriges Riksbank, the Bank of England, and the Reserve Bank of New Zealand have been especially important innovators in the development of new methods of communication with the markets and the general public. The Inflation

[44]Governor Kohn (2005) lists as an important reason for the earlier release of the minutes the desire to "help spell out the linkage the Committee may see between any policy inclination and its economic outlook" and "convey the conditionality of Committee thinking."

Reports of these banks provide good practical examples of communication with the public about the central bank's policy commitments. These reports do not pretend to give a blow-by-blow account of the deliberations by which the central bank reached the position that it has determined to announce; but they do explain the *analysis* behind the decision that has been reached. This analysis provides information about the bank's systematic approach to policy by illustrating its application to the concrete circumstances that have arisen since the last report; and it provides information about how conditions are likely to develop in the future through explicit discussion of the bank's own projections. Because the analysis is made public, it can be expected to shape future deliberations; the bank knows that it should be expected to explain why views expressed in the past are not later being followed. Thus a commitment to transparency of this sort helps to make policy more fully rule-based, as well as increasing the public's understanding of the rule.

The periodic publication of Inflation Reports is a key element in the kind of policy regime that Svensson (1999) calls "inflation-forecast targeting." Under this approach, the central bank does not only announce a quantitative target (its inflation target) that defines the goal of policy. It is also committed to a particular kind of decision procedure, under which projections are made of the future evolution of inflation and other variables under a particular assumed stance of policy, and the assumed policy is to be implemented (until the matter is reconsidered in the next decision cycle) only if the projections satisfy a certain target criterion. For example, in the case of the Bank of England, the criterion given primary emphasis (at least in the public justification of the policy decisions that are taken) is one that requires projected CPI inflation to equal 2.0 percent at a horizon eight quarters in the future.[45] And finally, it explains its policy decisions to the public in terms of their conformity with the target criterion. In the case of the central banks just mentioned, this means the publication of Inflation Reports several times per year, that give prominent attention to the projections that justify the current stance of policy.

While the development of inflation-forecast targeting represents a substantial advance, both in the commitment of central banks to the conduct of policy in accordance with an explicit rule and in the transparency of communication with the public about

[45]Before 2004, the criterion required the projection of a different inflation measure, RPIX inflation, to equal 2.5 percent at the 8-quarter horizon (Vickers, 1998; Goodhart, 2001). The change in target criterion is discussed in Bank of England (2004a).

policy deliberations, the precise techniques that are used continue to be refined. One of the most debated aspects of current practice has been the question of what kind of assumption to make about the *future* conduct of policy when preparing the projections that will be used to judge the appropriateness of *current* policy.

3.1 Constant-Interest-Rate Projections

Computation of projected paths for variables such as inflation and output some years into the future requires that one make assumptions about the future conduct of monetary policy — at least over the horizon of the projection, and, in the case of a model of the transmission mechanism that incorporates forward-looking behavior by the private sector, even farther. This is a particularly delicate aspect of such exercises, because of the implied need for the central bank to take a stand on the question of how it is likely to conduct policy in the future, and even more critically, because of concern that public discussion of the assumptions being made could be interpreted as a statement of the bank's intentions. Statements of intentions regarding future policy have often been regarded as problematic, for reasons of the kind already discussed above.

A common way of seeking to avoid any statement about future policy has been to base policy deliberations — or at least, the way these deliberations are presented in the banks' Inflation Reports — on projections of the future evolution of inflation and other variables under an assumption that the interest-rate target (repo rate) will remain *constant* over the horizon of the projection, at the level that is chosen currently. (This was, for example, the assumption made in the base-case projections that were presented in the overview section of the Bank of England *Inflation Reports,* prior to August 2004; and it is still the assumption made in the base-case projections in the *Inflation Reports* of the Sveriges Riksbank.) The policy decision is then a search for an interest rate with the property that constant-interest-rate projections assuming that particular rate satisfy a particular target criterion (for example, projected CPI inflation of 2.0 percent two years in the future).[46] The Inflation Reports then justify the interest-rate decision that has most recently been made by presenting these

[46]Former Bank of England MPC member Charles Goodhart (2001) describes himself as having tried to set interest rates in this way, and says "This was, I thought, what the exercise was supposed to be" (p. 177). Jansson and Vredin (2003) describe the similar procedure used by the Sveriges Riksbank.

projections and noting that they conform fairly well to the criterion in question.

This way of avoiding any need for the central bank to show its cards with regard to future policy has the advantage of being simple to explain to the public — as long as the public is not sophisticated enough to ask what it really means — but has a number of unappealing implications.[47] First of all, many optimizing models of the monetary transmission mechanism have the property first demonstrated by Sargent and Wallace (1975) for a rational-expectations IS-LM framework, namely, that the equilibrium path of the price level (and hence of the inflation rate) is *indeterminate* under the assumption of a fixed nominal interest rate (or indeed, any exogenously specified interest-rate process).[48] If such a model were to be used for the central bank's projection exercise, the staff would be unable to compute predicted paths for inflation or other variables under the hypothesis of any constant level of nominal interest rate, and so unable to assert that one particular level would imply satisfaction of the target criterion.[49]

Alternatively, many backward-looking models (including optimizing models in which expectations are assumed to be based on extrapolation from past time series) have the property discussed by Friedman (1968), namely, that maintaining a constant nominal interest rate indefinitely will lead to explosive inflation dynamics, through a Wicksellian "cumulative process."[50] Goodhart (2001) suggests that the Bank of England's model has this latter property, and that as a result, "the rate of change of most variables visible at the two-year horizon in the Bank's forecast generally (though not invariably) tends to persist, and on occasion to accelerate, in the third and subsequent years" (p. 171).[51] An example of this property is shown in the

[47]Goodhart (2001) reviews what he calls "the prima facie case against" this approach before offering his defense of it. Other critical discussions include Leitemo (2003), Svensson (2003), and Honkapohja and Mitra (2003).

[48]See Woodford (2003, chap. 4) for further discussion.

[49]Leitemo (2003) discusses possible interpretations of the constant-interest-rate projection exercise that would allow it to yield a policy recommendation even in the case of a forward-looking model of the transmission mechanism; but these do not eliminate the other unappealing features of such a procedure.

[50]See Bullard and Mitra (2002) and Preston (2005) for analyses of forward-looking models with least-squares learning by the private sector.

[51]Goodhart (2005) indicates that "in medium run simulations at the Bank of England running much beyond [a] two-year horizon, the constant two-year rate assumption had to be linked into a Taylor-type reaction function to prevent nonsensical trends from developing as the horizon extended

box on pp. 42-43 of the August 2004 *Inflation Report* (Bank of England, 2004b), where the constant-interest-rate inflation projection from the February 2004 report is extended another year into the future. While the projection showed CPI inflation rising to about 2.0 percent, under the most likely scenario, by early 2006, it showed inflation continuing to rise, to about 2.4 percent, by early 2007, with no indication of convergence even at that level. In this case, it is possible to ask which constant interest rate would imply satisfaction of the target criterion at a certain finite horizon, but only at the expense of making it clear that hitting the target at (say) the 8-quarter horizon does *not* also imply expecting to hit it in subsequent quarters. Hence it cannot be the case that one expects to be content to maintain the constant-interest-rate policy indefinitely, even in the absence of any developments that cannot already be foreseen.

Moreover, if one's model currently implies that inflation will depart significantly from the target rate at the three-year horizon if interest rates are maintained at their current level for that long, then it also implies that one should expect that a year from now — barring unforeseen developments — if interest rates have been maintained at their current level, it will then be forecasted that inflation will depart from the target at the *two-year* horizon if interest rates are not changed. For example, in the case just mentioned, the projection in February 2004, based on an assumption that the repo rate would remain at 4.0 percent over the following three years, implied showed CPI inflation accelerating to about 2.4 percent by early 2007. But this projection would then imply that under the most likely scenario, keeping the repo rate at 4.0 percent throughout 2004 would be expected to result in the Bank's projecting in February 2005 that CPI inflation should reach 2.4 percent in only two years, at which point (if not sooner), the exercise should require the repo rate to be raised. Thus the projection would imply that one should *not* expect the repo rate to remain at its current level for an entire year, even in the absence of any "news". It should have been expected to be raised fairly soon, as indeed it was (by 75 basis points over the next six months).

---

beyond two years (p. 7)." Goodhart argues that this does not present any problem for an exercise in which the constant-rate assumption is maintained only over a two-year horizon. But if one admits the acceptability of imposing a reaction function as the policy assumption beyond the two-year horizon, it is unclear why one should not be willing to impose a reaction function for earlier dates as well, at least to the extent that the MPC is not willing to take a stand on a particular direction of likely near-term deviation from the 'typical' reaction function.

The publication of constant-interest-rate projections — and the public justification of policy decisions by reference to them — is in no way intended to suggest that the central bank intends to maintain interest rates constant over the period of the projection. (Indeed, the most important justification for the use of constant-interest-rate projections seems to be a desire *not* to express any intention regarding future policy.) Nor can it be defended as representing the central bank's own best current estimate of the future path of interest rates; after all, the implication of the projections explained above would be evident above all within the central bank itself.[52]

But this implies that the targeting procedure is based on forecasts that are not actually believed, even in the central bank itself. Such a procedure has the paradoxical implication that the central bank may choose a policy under which it does *not* truly expect the target criterion to be satisfied, though it may believe that it would be under the counterfactual hypothesis of the constant interest rate.[53] Such a state of affairs can hardly be defended as conducive to transparency in the conduct of monetary policy. If policy is genuinely based on constant-interest-rate conditional projections, then one's policy decisions are not aimed at ensuring satisfaction of the

[52]Even before the Bank of England ceased to use the constant-interest-rate projections as the base case in its *Inflation Report,* it was fairly clear that these projections did not represent the Bank's own forecast of how the economy was most likely to evolve. This was conceded at least implicitly in the Bank's published discussions of the accuracy of its projections, for example in the *Inflation Reports* of August 2001 and August 2002. In these discussions, the Bank gave exclusive attention to the projections that it also published in which an interest-rate path was assumed corresponding to current market expectations, rather than to the projections conditional on the constant interest-rate path. If the Bank regarded the constant-interest-rate assumption as the best available forecast of its behavior, it would want to test the accuracy of the projections made under that assumption, rather than under contrary assumptions that might be made by traders in financial markets.

[53]In the case of the February 2004 projections of the Bank of England already discussed above, the constant-interest-rate projection used as the main basis for policy deliberations at the time indicated inflation near 2.0 percent at the 8-quarter horizon, but surging above two percent over the next year. As shown in the Bank's subsequent discussion of its decision to de-emphasize the constant-interest-rate projections (Bank of England, 2004b, pp. 42-43), a projection based on market anticipations regarding the future path of interest rates (according to which interest rates would soon be raised above the 4.0 percent decided upon in February) instead implied that inflation should remain well below two percent over the entire two-year period, though this projection would imply an inflation rate nearing two percent if extended three years into the future. If the latter projection were really the one regarded as more realistic by the MPC in February, this would imply that they were not really basing their decision on projected inflation at the 8-quarter horizon.

target criterion that is announced to the public; and the projections published by the central bank are not accurate forecasts that should better help the private sector to correctly anticipate the economy's evolution. On the other hand, if the central bank genuinely does expect the target criterion to be satisfied, then policy is not actually determined in the way that the official rhetoric implies that it is; and if the forecasts are unbiased, then they are not the kind of forecasts that they are officially described as being.

3.2 Projections Based on Market Expectations

The Bank of England has evidently accepted the force of at least some of the criticisms that have been raised of the use of constant-interest-rate projections, and since the August 2004 *Inflation Report* it has ceased to emphasize those projections in its justification of current policy.[54] It now focuses its evaluation of policy on a set of projections that are conditional on the path of short-term interest rates implied by the term structure of yields on longer-term Treasury securities.[55] This allows the projections to be based on a more realistic assumption regarding future interest rates, while still allowing the MPC to remain silent as to whether the interest-rate assumptions used in the projection exercise agree with their own.[56] The way in which

[54]The Minutes of the monthly Monetary Policy Committee meetings in 2004 indicate considerable emphasis on the projections based on market expectations, as opposed to the constant-interest-rate projections, in the actual policy decisions from at least March 2004 onward. The August *Inflation Report* was the first one to recognize this change in the MPC's thinking by presenting the market-expectations projections as the base case.

[55]The Sveriges Riksbank has also begun, in its 2005:Q1 *Inflation Report*, to give more prominence to projections conditional on an interest-rate path inferred from the term structure. However, the Riksbank continues to base its "main scenario" on the assumption that the repo rate will remain unchanged. The alternative scenario is also presented, late in the report, to "provide a broader base for discussions of monetary policy" (Sveriges Riksbank, 2005, p. 5), but is not referred to in the main discussion of inflation risks under current policy. The *Inflation Reports* of Norges Bank currently give greatest prominence to a baseline scenario in which the interest-rate assumption is described as "based on market interest rate expectations" (see, *e.g.*, Norges Bank, 2005, p.2), but which seems actually to represent a view of the Bank itself as to the most likely path of interest rates, as discussed further below.

[56]The explanation of the interest-rate assumption includes the following qualification: "It is important to stress that the market rate path does not represent the MPC's forecast of official interest rates.... Financial markets may judge economic prospects differently from the MPC" (Bank of

these projections are used in the MPC's decision procedure has not been explicitly discussed. Apparently, the idea is that if the projections based on market expectations satisfy the target criterion, then the MPC should set the repo rate at the level *expected by the markets.* Thus in August 2004 the rate was raised from 4.50 to 4.75 percent, to conform to market expectations of a repo rate averaging 4.6 in 2004:Q3 and rising to an average of 4.9 in 2004:Q4. In November 2004, it was left unchanged, because market expectations at this time forecasted only an average of 4.7 in both 2004:Q4 and 2005:Q1; and in February 2005, it was again left unchanged, because this continued to be what the markets expected.[57] If the projections based on market expectations were to fail to satisfy the target criterion, then presumably policy should deviate from the market expectation, though it is not clear by how much; this situation appears not yet to have arisen.

While this alternative avoids some of the problems associated with constant-interest-rate projections, it does not avoid all of them, and introduces some new problems of its own.[58] The problems of inconsistency that arise when policy deliberations (or at least the public justification of these deliberations) are based on an interest-rate assumption different from what the MPC believes is most likely to occur remain, as long as the assumed path does not actually represent the MPC's forecast. The new approach makes the contradiction less glaring, since it is no longer obvious in which particular way the MPC should be expecting something different from what is assumed. But if the MPC *does* disagree with the assumed interest-rate path, the procedure is still incoherent and/or misleading; and it is hard to see how continued coyness of the MPC as to whether it agrees or not with its stated assumptions can serve the goal of clarifying the way in which policy is conducted.

The problem that the interest-rate assumption would not allow an equilibrium path to be computed, in the case of a standard forward-looking model, or would imply unstable dynamics, in the case of many backward-looking models, is also not avoided by the new proposal. These problems arise whenever *a path for the nominal interest rate is assumed,* that is independent of the evolution of the endogenous variables in the model; they have nothing to do with the *constancy* of the assumed

England, 2004b, p. 41).

[57] See Table 6.A, "Market expectations of the Bank's official interest rate," in each of these issues of the *Inflation Report.*

[58] See also Goodhart (2005) for criticism of the new procedure.

path.[59] When a particular path for interest rates is inferred from the term structure and plugged into a model of the monetary transmission mechanism as if it were given exogenously, the same problems will arise as in the case of a constant-interest-rate assumption. (Even if nominal interest rates are assumed to rise over the coming year, the fact that they are assumed to rise by an amount independent of any increase in inflation gives rise to unstable dynamics: higher inflation will lower the real interest rate, stimulating higher inflation, and so on.) The only solution for this problem is to include in one's model a realistic representation of *endogenous* variation in short-term nominal interest rates.

At the same time, a decision procedure using projections based on market expectations introduces some new problems. Most importantly, it runs the risk of making policy too sensitive to market expectations regarding policy, in a way that fails to provide any anchor for those expectations or any stable course for policy. Consider an extreme version of the approach to policy sketched above, in which the central bank simply infers market expectations regarding the path of its policy rate, and sets the rate in accordance with market expectations. As Blinder (1998) points out, "following the markets" in this way would be quite dangerous, for there would then be no reason for the markets to expect one kind of policy rather than any other; arbitrary notions could easily become self-fulfilling, and the sources of significant instability and/or inefficiency.

Of course, the approach to policy that has apparently been adopted by the Bank of England is more sophisticated than that; there is no commitment to follow the markets' lead, unless the Bank's projections imply that doing so is consistent with a projection of inflation near two percent two years in the future. But it is not clear

[59] Here I am interpreting a projection based on "market expectations" as one that computes a particular forward path for interest rates from the term structure, and then substitutes that interest-rate path into the structural model as the specification of monetary policy. One could imagine carrying out such an exercise in other ways — in the next section I discuss an alternative — and it is not clear from published discussions exactly what kind of computations are undertaken by the Bank of England staff. But the interpretation discussed here is not only the simplest version of such an exercise, it is the only version that would allow a central bank to avoid making any assumptions of its own about how it will conduct policy in the future. Thus even if it is not the kind of projection exercise currently used at the Bank of England, it is a type of exercise that might well be attempted by a central bank seeking to follow their lead in using projections based on market expectations.

to what extent such a qualification should succeed in preventing fluctuations due to self-fulfilling expectations. This is because, even when arbitrary fluctuations of that kind occur, inflation may still be (correctly) predicted at any point in time to revert back to its target level within a few quarters.

As an example, consider the simple model of the monetary transmission mechanism discussed above, according to which inflation and output are determined each period by equations (2.1) and (2.2). Let s_t be an exogenous random process that has no relation to economic "fundamentals"[60] — what is sometimes called a "sunspot" variable — and suppose that its dynamics can be described by a first-order autoregressive equation,

$$s_t = \rho s_{t-1} + \epsilon_t, \tag{3.1}$$

where $0 < \rho < 1$ is the coefficient of autocorrelation and $\{\epsilon_t\}$ is a bounded white-noise disturbance. I wish to consider the possibility of an equilibrium in which inflation and output fluctuate in response to variation in the sunspot variable, simply because people have come (for arbitrary reasons) to expect this, and have observed that the variable does indeed help them to forecast the economy's future evolution. Consider, for example, the possibility of an equilibrium in which

$$\pi_t = \phi s_t,$$

for some coefficient $\phi \neq 0$. Equation (2.1) will be satisfied as long as the equilibrium fluctuations in the output gap are given by

$$x_t = \frac{1 - \beta\rho}{\kappa} \phi s_t.$$

Equation (2.2) will also be satisfied as long as the nominal interest rate satisfies

$$i_t = r_t^n + \left[\rho - (1-\rho)\left(\frac{1-\beta\rho}{\kappa}\right)\sigma^{-1}\right]\phi s_t. \tag{3.2}$$

Now suppose that in each period t, the exogenous state of the world (including the current values of r_t^n and s_t) is first revealed; then futures markets are open, in which traders bet on the value of that period's interest rate i_t; and finally, the central bank chooses its operating target for i_t, after learning the market forecast $i_{t|t}$

[60] In particular, it is unrelated to the real factors that cause variation in either the natural rate of output or the natural rate of interest.

implied by the futures prices, and period t inflation and output are then determined in accordance with equations (2.1) – (2.2). In the equilibrium just conjectured, the market expectation should be given by

$$i_{t|t} = r_t^n + \left[\rho - (1-\rho)\left(\frac{1-\beta\rho}{\kappa}\right)\sigma^{-1}\right]\phi s_t.$$

If the central bank observes this, and sets $i_t = i_{t|t}$, then (3.2) will be satisfied, and the equations given above for inflation and output will represent a rational-expectations equilibrium.

But should the central bank be willing to follow the markets and set $i_t = i_{t|t}$? In the equilibrium just described, a correct forecast of the future path of inflation will in any period be given by

$$E_t\pi_{t+j} = \rho^j \phi s_t,$$

for arbitrary $j \geq 0$.[61] If the central bank forecasts in this way, conditional on policy that coincides with market expectations, it will conclude that the inflation rate should converge back to its long-run target value (here assumed to equal zero[62]) as the horizon j is extended farther into the future; and this will be true no matter how large ϕ is, and hence no matter how large the short-run fluctuations in inflation and output due to self-fulfilling expectations may be. Furthermore, if ρ is not close to 1, the convergence will be predicted to be nearly complete after only a few quarters.[63]

[61] Substitution of the interest rate path $\{i_{t+j|t}\}$ expected by the markets into the model consisting of equations (2.1) – (2.2) would not yield this as the only possible equilibrium, because of the indeterminacy problem already referred to in the case of any exogenously specified interest-rate path. But it would be *one* possible projection consistent with both the model and with market expectations regarding the path of interest rates; and since it is the *correct* forecast, in the conjectured equilibrium, I shall suppose that the central bank is sophisticated enough to produce this forecast.

[62] For simplicity, I here assume that the target inflation rate is zero. This allows me to avoid discussing the question whether it makes sense to suppose, as in the Calvo pricing model used here, that prices should remain fixed in nominal terms between the occasions on which they are re-optimized, even when the central bank's target rate of inflation is positive and inflation is positive on average. The point made here about the possibility of sunspot equilibria in which inflation reverts quickly to the target level would be equally valid if the target rate of inflation were assumed to be positive.

[63] For example, in the case that $\rho = 0.7$ and periods represent quarters, one could have an equilibrium in which actual inflation varies over a range as wide as 6 percentage points (3 points above and below the target), solely as a result of self-fulfilling expectations, but in which projected inflation 8 quarters in the future would never be more than 10 basis points away from the target.

It is thus not obvious that a central bank that pays attention only to the projected inflation rate two or more years in the future would find any reason not to follow the markets at all times, even if this policy allowed large transitory fluctuations in inflation (and associated large swings in output relative to potential) to occur, due solely to self-fulfilling expectations.[64]

One might argue that in the case just described, the possibility of fluctuations in response to the sunspot disturbance would be eliminated as long as the central bank were to establish a reputation for responding to even small departures of the projected inflation rate two years in the future from the target value. For example, if the central bank were to commit itself to a rule of the form

$$i_t = i_{t|t} + \psi \hat{E}_t \pi_{t+8}, \tag{3.3}$$

for some coefficient $\psi > 0$, where $\hat{E}_t \pi_{t+8}$ represents the central bank's projection of inflation 8 periods in the future, conditional on its conducting policy in the way anticipated in the futures markets, then the only possible rational-expectations equilibrium would have to be one in which $E_t \pi_{t+8} = 0$ at all times. This would preclude any equilibrium of the kind conjectured above, except the one with $\phi = 0$. However, this result depends on assuming that the central bank would respond systematically (albeit to only a small extent) to even small departures from satisfaction of its target criterion, which seems unlikely given that the policy rate is ordinarily moved only in discrete steps. More to the point, even rule (3.3) would allow the existence of sunspot fluctuations of arbitrary amplitude, in the case of a sunspot variable s_t that evolved as an MA(7) process (or any lower-order MA process), so that $E_t s_{t+8} = 0$ at all times.

Another problem with the current procedure of the Bank of England is that it is unclear how the MPC is intended to determine the correct current repo rate in the event that the interest-rate path expected by the markets is judged to imply projections inconsistent with the Bank's target criterion. Would an attempt be made to determine the current repo rate that would lead to an acceptable projection, under the assumption that *the path of the repo rate after the current month would follow the path anticipated by the markets?* This would typically require an extreme adjustment

[64]The problem just described arises because of the use of a target criterion that involves only the economy's projected state 8 quarters and more in the future; it could be eliminated by the use of a nearer-term target criterion. This would be desirable on other grounds as well, as discussed in Woodford (2004).

of the current repo rate, as a change in the repo rate for only one month would have to change the path of inflation over the following two years by enough to get the projected inflation rate two years in the future on track. A more sensible approach would surely involve adjusting the entire path of interest rates to one that the MPC would view as more sound, rather than acting as if the committee expected itself to behave in the future in the way currently anticipated by the markets, even though it was planning to depart substantially from the markets' expectation in the short run. But in this case, projections would have to be produced on the basis of an assumption about future policy other than the one corresponding to market expectations. The idea that the MPC would be able to avoid taking a stand (at least in its internal deliberations) on a reasonable future path of interest rates, by insisting on using the markets' forecast in its projections, is not tenable.

3.3 Projections Based on a Model of Central-Bank Behavior

The kind of forecast-targeting procedure recommended by Svensson and Woodford (2005) as a way of implementing optimal monetary policy is of a different sort than either of the two approaches just discussed. In this procedure, one projects the economy's future evolution under alternative contemplated policy decisions, assuming that in future decision cycles the central bank will *again act to ensure satisfaction of the target criterion.* This amounts to asking what action is needed in order to project that that the criterion should be satisfied in the current period, taking as given that it is expected to be satisfied in later periods (as a result of the policy actions to be taken in those periods). Such a calculation yields a determinate outcome as long as there is a determinate rational-expectations equilibrium implied by the target criterion; this is always the case if the target criterion is selected according to the method of Giannoni and Woodford (2002).

Thus I would argue that policy should be based on a projection exercise that includes a model of the central bank's own future behavior — one that is furthermore consistent with the procedure that it actually follows in making its policy decisions. This is the kind of projection exercise used as the basis for policy decisions at some central banks, notably the Reserve Bank of New Zealand, which also publishes some information about the non-constant interest-rate path implicit in its projections, along with its projections for inflation and other variables.[65]

Charles Goodhart (2001, 2005) objects that such a procedure is impractical, on the ground that it would be much more difficult for a monetary policy committee to reach agreement on an entire future path for interest rates, rather than allowing them to decide only about the current interest rate each time they meet. But the procedure described by Svensson and Woodford does not involve a multi-dimensional decision problem in each decision cycle. As with the constant-interest-rate projection method, one makes a decision for the current period only, on the basis of projections of the future that (necessarily) incorporate a hypothesis about future policy; the hypothesis about future policy is simply a more realistic one than the notion that interest rates will not change, regardless of how inflation and output evolve. And there is no greater need for agreement among the members of the policy committee about that particular aspect of the model specification than about the other assumptions involved in making projections for the future.[66] The fact that the *Inflation Reports* of Norges Bank have begun to include projected paths for money-market interest rates that (at least for the final two years of the projection period) represent the judgment of the Bank rather than market expectations also suggests that agreement on a path is not necessarily such a complex multidimensional decision.

Goodhart also argues that revealing a projected non-constant path for interest rates is problematic, because "any indication that the MPC is formally indicating a future specific change in rates ... would be taken to indicate some degree of com-

[65]The models used at some other banks, such as the Bank of Canada, similarly include equations intended to represent future policy; but these banks do not publish their projections. Norges Bank is another inflation-forecast-targeting central bank that has recently begun to publish projections based on a non-constant interest-rate path that appears to represent the Bank's own view of the most likely scenario, though it is unclear how this path is determined. The interest-rate assumption is described as "based on market interest rate expectations," and the interest-rate path assumed in the baseline scenario typically coincides with the path implied by forward rates for the first several quarters, though it may deviate from market expectations in later quarters if the Bank disagrees with the apparent expectations of the markets (Norges Bank, 2005, chart 3.2).

[66]The members of the monetary policy committee *might* wish, under certain circumstances, to modify the model's default rule for the future conduct of policy because they do not regard it as correctly representing their intentions regarding near-term future policy. But this would only occur to the extent that it was in fact possible for them to agree that current conditions would justify a departure from typical policy, as in the case of the recent response to the perceived threat of deflation in the U.S., discussed above.

mitment" (2001, p. 175).[67] This is clearly a delicate issue in the proper explanation to the public of how the central bank's projections are to be interpreted. Yet the danger is not as great in this case as in the case of the signals regarding future policy that have recently been included in the post-meeting statements of the FOMC in the U.S. For as has just been noted, the MPC would not have to decide on a forward path for interest rates, that is then fed into the bank's model in order to generate the projections; instead, the model could incorporate an equation representing typical policy. The resulting projections would include a path for the policy rate, and to the extent to which the MPC announced that this projection represented its best judgment about how the economy was likely to evolve, it would be endorsing a forecast of its future decisions. But it would be fairly easy in this case for the discussion of the projections to include a disclaimer stating that the interest-rate projection is not intended to pre-judge the policy decisions that the MPC will actually make at later dates.

And the experience of New Zealand suggests that it is possible to reveal interest-rate projections to the public without being understood to have made an advance commitment about the path of the official cash rate. Archer (2004) discusses the RBNZ's experience with the publication of a forward path for interest rates. While he mentions the possibility of the projected interest rate path's being misinterpreted "as a policy plan" (p.9), this is cited as a concern expressed to him by other central bankers, rather than as a problem that has arisen in practice in New Zealand. The main practical problem that he cites with regard to the publication of an interest-rate path is that it has not always been possible to produce a model-based projection that coincided with the policymaker's assessment of the likely future path of interest rates, in which case the quantitative projection and the discussion of future policy in the Monetary Policy Statement do not agree with one another. According to Archer, "it is noteworthy that on all such occasions, market analysts and position takers expressed frustration at the inconsistency" between the two contrasting messages regarding future policy (p. 10). The problem in such cases has been a failure of policymaker confidence in the assumptions reflected in the model-based projection. It is not clear why this problem should be greater in the case of projections regarding the path of interest rates than with other variables, though central banks may well be especially careful about what they say about this particular aspect of the future

[67] See also Goodhart (2005, p. 6).

outlook.

Moreover, a "fan chart" for the path of interest rates ought to make it clear that the bank is not committing itself to a definite path; rather, the expected evolution will depend on a variety of contingencies that can at best be assigned probabilities. The practical possibility of communicating about this kind of uncertainty is illustrated by the way that the Bank of England currently discusses the interest-rate assumptions used in its projections "based on market interest-rate expectations." When explaining the particular interest-rate path that is assumed in the projections, the *Inflation Report* always takes care to note that "there is a high probability that official interest rates will not follow this path .... Given the great uncertainties, the economy will almost certainly not evolve in the way either the MPC or the markets expect," and a fan chart is presented indicating the degree of market uncertainty about interest rates at progressively longer horizons, based on the prices of options on futures contracts.[68] The fan chart would seem an effective device for communication about the uncertainty around the published interest-rate path. In the case of an interest-rate path implied by a quantitative model augmented by a policy equation, it would be possible to generate the probability distributions that are plotted in the fan chart in the same way as the fan charts for inflation and output are currently generated.

Once one admits that the model used in one's projections needs an equation representing monetary policy, and not simply an interest-rate path as the specification of current and future policy, it is nonetheless possible to carry out exercises under which the projections are made to be consistent with a particular interest-rate path specified from outside the model. One might add an interest-rate equation to the model representing the central bank's reaction function (say, some version of a "Taylor rule"), and then allow the intercept term in this equation to specified arbitrarily at each future date. If one wants to produce projections consistent with a particular constant-interest-rate path, for example, one could search for the sequence of (non-constant) intercept terms that will result in a projected path for the (endogenous) short-term nominal interest rate with exactly this feature; one could use the same approach to produce projections consistent with market expectations inferred from the term structure of interest rates. An advantage of this roundabout approach over a simple postulation of a nominal interest-rate path is that it will allow one to obtain a

[68] See, for example, Chart 6.1 in Bank of England (2004b, p. 41); the quotation in the text is from the same page.

determinate rational-expectations equilibrium in the case of a forward-looking model of the monetary transmission mechanism, and some central banks have begun to use methods of this kind as a consequence of adopting forward-looking models.[69] While this would avoid one of the problems with approaches that use an assumed interest-rate path discussed above, it solves *only* the problem of determinacy of equilibrium; such procedures would still be vulnerable to the other criticisms offered in sections 3.1 and 3.2.

Thus there seems no coherent alternative to an approach to inflation-forecast targeting in which the central bank's projections are produced on the basis of an assumption that the MPC is willing to make about the way in which its own future policy should endogenously respond to alternative possible future conditions. This need not involve an attempt to determine in advance the actual future path of interest rates (except, of course, in a "central case" scenario that is understood to be unlikely to actually occur), and the central bank might or might not wish to reveal much about the interest-rate path implicit in its projections for other variables. My own view is that communication about this path can help to facilitate the transmission mechanism of policy. For example, Archer (2004) concludes that the slope of the path announced by the RBNZ has effected the slope of the market yield curve in New Zealand, and an ability to influence the expectations reflected in the yield is exactly what a central bank should wish, as argued in section 1. But even if a central bank were to wish to communicate less about this aspect of its deliberations — and it must be admitted that even the RBNZ limits the amount of detail with which it describes the interest-rate path implied by its projections — it should not allow such concerns to prevent it from considering the appropriate character of future policy in its forecast-targeting exercise. And to the extent that it is concerned at all about transparency and accountability, it should not pretend to the public that its deliberations make no assumptions about future policy.

4 Conclusion

The increased willingness of the FOMC under the Chairmanship of Alan Greenspan to speak openly about both current policy decisions and the Committee's view of

[69] In the case of a purely backward-looking model of the transmission mechanism, the alternative procedure would lead to an identical outcome and would only be more cumbersome computationally.

likely future policy has greatly increased the ability of markets to anticipate Fed policy. There is every reason to believe that this has been beneficial, not only from the point of view of reducing the uncertainty with which traders and other economic decisionmakers must contend, but also from that of enhancing the accuracy with which the FOMC is able to achieve the effects on the economy that it desires, by keeping the expectations of market participants more closely synchronized with its own. It is very much to be hoped that the FOMC will continue to build on this legacy under its next Chairman.

Despite the worldwide movement toward greater transparency and increased communication on the part of central banks — at least as notable in the case of the inflation-targeting central banks as it has been in the U.S. — both of the case studies just discussed indicate that there continues to be a good deal of uncertainty in central banks about the degree to which it is desirable to publicly indicate the bank's own view of the likely future path of interest rates. Nonetheless, it seems likely that in coming years central banks will develop even more effective ways of communicating about this issue.

In the case of the inflation-forecast targeting central banks, the logic of the approach to the conduct of monetary policy — and in particular, the approach to communication with the public about monetary policy decisions — that they have already adopted (and found to be largely successful) will, in my view, almost inevitably force these banks to base the projections that are at the heart of their public explanation of their policy decisions on an explicit model of their own likely future approach to policy. Their existing commitment to transparency will also require them to publicly explain the assumptions about future policy that are implicit in such projections, at least to some extent, and with suitable caveats about the fact that such policy assumptions do not represent a commitment to actually set policy in the precise way indicated by the projection. Insofar as the FOMC increases the amount that it makes public about the projections upon which its own deliberations are made — which would certainly be a logical development of the current trend toward greater transparency, even if the Fed does not adopt inflation targeting — this could also become an issue that will have to be confronted at the Fed under the next Chairman.

One cannot be quite as certain about the future role of explicit indications about the expected path of interest rates of the kind with which the FOMC has experimented over the past two years. Thus far, I believe that the new approach has been a

successful one. On the other hand, one might argue that this kind of explicit comment on the intended path of interest rates (as opposed to more general discussion of the considerations that will be taken into account in future policy decisions) is most useful under relatively special circumstances such as the ones faced by the Fed in the summer of 2003; other ways of guiding expectations about future policy might prove more suitable under more ordinary conditions. At the very least, an extension of the current policy will require the development of a more extensive lexicon of possible statements about expected future policy. But while Fed communication policy will surely develop further, the challenge of efficient communication with the markets about the future conduct of policy is one that should continue to engage the FOMC in coming years, and the relative success of recent innovations should encourage further efforts to increase the degree to which the FOMC is able to share its view of the future evolution of policy with the public.

A The Model of Morris and Shin (2002)

Morris and Shin consider a stylized game in which each of a continuum of market participants, indexed by i, chooses an action a_i (a real number) after observing a public signal y and a private signal x_i. As explained in the text, the payoff of each agent i depends on her own action a_i, the average action $\bar{a}$ chosen by market participants as a whole, and an unknown "fundamental" state θ (another real number). The public signal is common knowledge, and given by

$$y = \theta + \eta,$$

where the random noise term η is independently of the value of θ, according to a normal distribution with mean zero and variance $1/\alpha$. Each of the market participants' private signal is given by

$$x_i = \theta + \epsilon_i,$$

where the random noise term ϵ_i is distributed independently of θ, η, and each of the other ϵ_j, according to a normal distribution with mean zero and variance $1/\beta$. Thus α and β measure the precision of the public and private signals respectively.

Morris and Shin assume that the payoff of each market participant is inversely related to a loss function of the form

$$L^{i.MS} \equiv (1-r)(a_i - \theta)^2 + r[E_j(a_j - a_i)^2 - E_k E_j (a_j - a_k)^2], \tag{A.1}$$

where $0 < r < 1$ and, as in the text, I use the notation $E_j(\cdot)$ for the average of some quantity over the continuum of agents indexed by j.[70] This differs from the specification (1.1) given in the text by the presence of the second term inside the square brackets. The additional term is additive and independent of the action of agent i, so it has no effect on equilibrium behavior. It matters only for the evaluation of welfare, and is introduced by Morris and Shin to justify their assumption that the dispersion of actions across market participants is of no consequence for welfare, even though it is a concern that is given considerable weight in the decisions of individuals.

Under these preferences (whether written as in (1.1) or as in (A.1)), the optimal action of agent i will be given by

$$a_i^* = E[r\bar{a} + (1-r)\theta | y, x_i],$$

[70] Morris and Shin write integrals for this operator.

i's estimate of a weighted average of the fundamental state θ and the average action of others $\bar{a}$. (Here $E[\cdot|\cdot]$ is the conditional expectation operator, averaging over possible realizations of θ and the random signals.) Thus a higher value of r implies a greater tendency to choose an action similar to what one guesses that others are doing (greater strategic complementarity).

Under the assumption of prior beliefs about θ given by an improper uniform prior on the real line, an agent's optimal estimate of the fundamental state after observing both private and public signals will be

$$E[\theta|y, x_i] = \frac{\alpha}{\alpha+\beta}y + \frac{\beta}{\alpha+\beta}x_i,$$

a weighted average of the two signals with weights that depend on the relative precisions of the signals. One can then show that the equilibrium action of each market participant will be given by

$$a_i^* = \frac{\alpha y + \beta(1-r)x_i}{\alpha + \beta(1-r)}. \tag{A.2}$$

Expression (A.2) shows how the weight that market participants put on their private information decreases as a result of an increase in the precision α of the public signal.

Morris and Shin measure social welfare by the negative of the social loss function

$$L^{soc,MS} = \frac{1}{1-r}E_i L^{i,MS},$$

which is easily seen to correspond to the expression (1.2) in the text. It then follows from the description (A.2) of equilibrium behavior that

$$\Lambda^{MS} \equiv E[L^{soc,MS}|\theta] = \frac{\alpha + \beta(1-r)^2}{[\alpha + \beta(1-r)]^2}. \tag{A.3}$$

The key result of Morris and Shin is that the function on the right-hand side of (A.3) is not necessarily decreasing in α, for given values of β and r. In particular, one can show that $\partial\Lambda^{MS}/\partial\alpha > 0$ if and only if

$$\frac{\alpha}{\beta} < (1-r)(2r-1). \tag{A.4}$$

Since the right-hand side of (A.4) is positive for all $1/2 < r < 1$, there exist parameter configurations that satisfy (A.4). Hence it is possible to construct examples in which

a (small enough) increase in the precision of the public signal would increase expected social losses, under the Morris-Shin measure. As noted in the text, however, condition (A.4) is fairly stringent. Svensson (2005) points out that the right-hand side of (A.4) is no greater than $1/8$ for any value of r; hence the perverse case cannot arise unless $\alpha < \beta/8$. Moreover, the right-hand side is only positive if $r > 1/2$, so the perverse case does not arise, even when α is much smaller than β, unless $r > 1/2$, so that the weight on the conformity-with-others objective in market participants' preferences is greater than the weight on the conformity-with-fundamentals objective.

The possibility of a perverse case also depends on the (completely unmotivated) inclusion of the final term in the individual loss function (A.1). Suppose that one instead represents individual preferences by the loss function (1.1) proposed in the text, and uses as a social loss function the corresponding population average,

$$L^{soc} = \frac{1}{1-r} E_i L^i.$$

Equilibrium actions are again given by (A.2), but the implied expected level of social losses is now

$$\Lambda \equiv E[L^{soc}|\theta] = \frac{\alpha + \beta(1-r^2)}{[\alpha + \beta(1-r)]^2}. \tag{A.5}$$

The right-hand side of (A.5) is easily seen to be a globally decreasing function of α, for any values of β and r. Thus in this case, the release of a more precise estimate by the public authority will necessarily improve social welfare, regardless of how imprecise the estimate may be that it is possible for the authority to release, and of how great may be the concern of market participants to behave as others do.

References

Adam, Klaus, and Roberto M. Billi, "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," unpublished, University of Frankfurt, August 2003.

Amato, Jeffery D., Stephen Morris, and Hyun Song Shin, "Communication and Monetary Policy," *Oxford Review of Economic Policy* 18: 495-503 (2002).

Amato, Jeffery D., and Hyun Song Shin, "Public and Private Information in Monetary Policy Models," BIS working paper no. 138, July 2003.

Angeletos, George-Marios, and Alessandro Pavan, "Transparency of Information and Coordination in Economies with Investment Complementarities," *American Economic Review* 94(2): 91-98 (2004).

Archer, David, "Communication with the Public," unpublished, Reserve Bank of New Zealand, May 2004.

Bank of England, *Inflation Report,* February 2004a.

—- —-, *Inflation Report,* August 2004b.

Bernanke, Ben S., "Central Bank Talk and Monetary Policy," remarks at the Japan Society corporate luncheon, New York, October 7, 2004.

Bernanke, Ben S., Vincent R. Reinhart, and Brian P. Sack, "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment," *Brookings Papers on Economic Activity* 2004:2, pp. 1-78.

Blinder, Alan, Charles Goodhart, Philipp Hildebrand, David Lipton, and Charles Wyplosz, *How Do Central Banks Talk?* Geneva Report on the World Economy no. 3, International Center for Monetary and Banking Studies, 2001.

Bullard, James, and Kaushik Mitra, "Learning about Monetary Policy Rules," *Journal of Monetary Economics* 49: 1105-1129 (2002).

Clarida, Richard, Jordi Gali and Mark Gertler (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37: 1661-1707.

Demiralp, Selva, and Oscar Jorda, "The Announcement Effect: Evidence from

Open Market Desk Data," Federal Reserve Bank of New York *Economic Policy Review* 8: 29-48 (2002).

Economist, "It's Not Always Good to Talk," July 22, 2004, page 71.

Eggertsson, Gauti B., and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity* 2003:1, pp. 139-233.

Ehrmann, Michael, and Marcel Fratzscher, "Transparency, Disclosure and the Federal Reserve," ECB working paper no. 457, March 2005.

Eusepi, Stefano, "Central Bank Transparency under Model Uncertainty," unpublished, Federal Reserve Bank of New York, January 2005.

Fujiki, Hiroshi, and Shigenori Shiratsuka, "Policy Duration Effect under the Zero Interest Rate Policy in 1999-2000: Evidence from Japan's Money Market Data," Bank of Japan *Monetary and Economic Studies* 20(1): 1-31 (2002).

Giannoni, Marc P. and Michael Woodford, "Optimal Interest-Rate Rules: I. General Theory," NBER working paper no. 9419, December 2002.

Goodhart, Charles A.E., "Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates," Federal Reserve Bank of St. Louis *Review,* July/August 2001, pp. 165-181.

—- —-, "The Interest Rate Conditioning Assumption," unpublished, London School of Economics, May 2005.

Gurkaynak, Refet S., "Using Federal Funds Futures Contracts for Monetary Policy Analysis," Federal Reserve Board, Finance and Economics Discussion Series working paper no. 2005-29, June 2005.

—- —-, Brian Sack, and Eric Swanson, "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking* 1: 55-93 (2005).

Hellwig, Christian, "Heterogeneous Information and the Benefits of Transparency," unpublished, UCLA, December 2004.

Honkapohja, Seppo, and Kaushik Mitra, "Problems in Inflation Targeting Based on Constant Interest Rate Projections," unpublished, University of Helsinki, August 2003.

Issing, Otmar, "Communication, Transparency, Accountability: Monetary Policy in the Twenty-First Century," Federal Reserve Bank of St. Louis *Review* 87(2, pt. 1), pp. 65-83 (2005).

Iwamura, Mitsuri, Takeshi Kudo, and Tsutomu Watanabe, "Monetary and Fiscal Policy in a Liquidity Trap: The Japanese Experience 1999-2004," unpublished, Waseda University, August 2004.

Jansson, Per, and Anders Vredin, "Forecast-Based Monetary Policy: The Case of Sweden," *International Finance* 6: 349-380 (2003).

Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, "Optimal Monetary Policy at the Zero-Interest-Rate Bound," *Journal of Money, Credit and Banking,* forthcoming 2005.

King, Mervyn, "What Has Inflation Targeting Achieved?" in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate,* Chicago: University of Chicago Press, 2005a.

—- —-, "Monetary Policy: Practice Ahead of Theory," 2005 Mais Lecture, delivered at Cass Business School, London, May 17. 2005b.

Kohn, Donald L., "Central Bank Communication," remarks at the annual meeting of the American Economics Association, Philadelphia, January 9, 2005.

Krugman, Paul (1998), "It's Baaack! Japan's Slump and the return of the Liquidity Trap," Brookings Papers on Economic Activity 1998:2, pp. 137-187.

Kydland, Finn E., and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85: 473-491 (1977).

Lange, Joe, Brian Sack, and William Whitesell, "Anticipations of Monetary Policy in Financial Markets," Federal Reserve Board, FEDS paper no. 2001-24, April 2001.

Leitemo, Kai, "Targeting Inflation by Constant-Interest-Rate Forecasts," *Journal of Money, Credit and Banking* 35: 609-26 (2003).

Morris, Stephen, and Hyun Song Shin, "The Social Value of Public Information," *American Economic Review* 92: 1521-34 (2002).

Norges Bank, *Inflation Report,* June 2005.

Oda, Nobuyuki, and Kazuo Ueda, "The Effects of the Bank of Japan's Zero Interest Rate Commitment and Quantitative Monetary Easing on the Yield Curve: A Macro-Finance Approach," Bank of Japan Working Paper no. 05-E-6, April 2005.

Orphanides, Athanasios, and John C. Williams, "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate,* Chicago: University of Chicago Press, 2005.

Poole, William, "Remarks: Panel on 'After Greenspan: Whither Fed Policy?'" remarks delivered at the Western Economics Association International conference, San Francisco, California, July 6, 2005.

Preston, Bruce, "Learning about Monetary Policy Rules when Long-Horizon Forecasts Matter," *International Journal of Central Banking,* forthcoming 2005.

Roca, Mauro, "Transparency and Monetary Policy with Imperfect Common Knowledge," unpublished, Columbia University, February 2005.

Romer, Christina D., and David H. Romer, "Federal Reserve Private Information and the Behavior of Interest Rates," *American Economic Review* 90: 429-457 (2000).

Sargent, Thomas J., and Neil Wallace, "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule," *Journal of Political Economy* 83: 241-254 (1975).

Sugo, Tomohiro, and Yuki Teranishi, "How Should Optimal Monetary Policy Rules Be Designed in a Low Inflation Era?" unpublished, Bank of Japan, March 2005.

Svensson, Lars E.O., "Inflation Targeting as a Monetary Policy Rule," *Journal of Monetary Economics* 43: 607-654 (1999).

—- —-,"The Inflation Forecast and the Loss Function," in P. Mizen, ed., *Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart,* vol. 1, Edward Elgar, 2003.

—- —-, "The Magic of the Exchange Rate: Optimal Escape from a Liquidity Trap in Small and Large Open Economies," unpublished, Princeton University, July 2004.

—- —-, "Social Value of Public Information: Morris and Shin (2002) is Actually Pro Transparency, Not Con," unpublished, Princeton University, February 2005.

—- —- and Michael Woodford, "Implementing Optimal Policy through Inflation-Forecast Targeting," in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate,* Chicago: Univ. of Chicago Press, 2005.

Sveriges Riksbank, *Inflation Report,* First Quarter, 2005.

Vickers, John, "Inflation Targeting in Practice: The U.K. Experience," *Bank of England Quarterly Bulletin,* November 1998.

Woodford, Michael, "Monetary Policy in the Information Economy," in *Economic Policy for the Information Economy*, Kansas City: Federal Reserve Bank of Kansas City, 2001.

—- —-, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.

—- —-, "Inflation Targeting and Optimal Monetary Policy," Federal Reserve Bank of St. Louis *Economic Review,* July/August 2004, pp.15-41.

NBER WORKING PAPER SERIES

CENTRAL BANK COMMUNICATION AND MONETARY POLICY: A SURVEY OF THEORY AND EVIDENCE

Alan S. Blinder
Michael Ehrmann
Marcel Fratzscher
Jakob De Haan
David-Jan Jansen

Working Paper 13932
http://www.nber.org/papers/w13932

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2008

Blinder: Princeton University. Ehrmann: European Central Bank. Fratzscher: European Central Bank. De Haan: University of Groningen and CESifo. Jansen: De Nederlandsche Bank. The authors are grateful to Sylvester Eijffinger, Gabriel Fagan, Andreas Fischer, Otmar Issing, Frederic Mishkin, Glenn Rudebusch, Pierre Siklos, Eric Swanson, Charles Wyplosz, and the editor and two anonymous referees of this Journal for valuable comments on earlier drafts. Views expressed in this article do not necessarily coincide with those of the European Central Bank, de Nederlandsche Bank, the Eurosystem. or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Alan S. Blinder, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence
Alan S. Blinder, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan , and David-Jan Jansen
NBER Working Paper No. 13932
April 2008
JEL No. E52,E58

ABSTRACT

Over the last two decades, communication has become an increasingly important aspect of monetary policy. These real-world developments have spawned a huge new scholarly literature on central bank communication -- mostly empirical, and almost all of it written in this decade. We survey this ever-growing literature. The evidence suggests that communication can be an important and powerful part of the central bank's toolkit since it has the ability to move financial markets, to enhance the predictability of monetary policy decisions, and potentially to help achieve central banks' macroeconomic objectives. However, the large variation in communication strategies across central banks suggests that a consensus has yet to emerge on what constitutes an optimal communication strategy.

Alan S. Blinder
Department of Economics
Princeton University
Princeton, NJ 08544-1021
and NBER
blinder@princeton.edu

Michael Ehrmann
European Central Bank
Postfach 16 03 19
D-60066 Frankfurt am Main
GERMANY
michael.ehrmann@ecb.int

Marcel Fratzscher
European Central Bank
Kaiserstrasse 29
D-60311 Frankfurt/Main GERMANY
marcel.fratzscher@ecb.int

Jakob De Haan
Faculty of Economics and Business
University of Groningen
PO Box 800
9700 AV Groningen, Netherlands
and CESifo
jakob.de.haan@rug.nl

David-Jan Jansen
De Nederlandsche Bank
Westeinde 1
PO Box 98
1000 AB Amsterdam
The Netherlands
d.jansen@dnb.nl

1. A Revolution in Thinking and Practice

Prior to the 1990s, central banks were shrouded in mystery—and believed they should be. Conventional wisdom in central banking circles held that monetary policymakers should say as little as possible, and say it cryptically. In 1981, Karl Brunner (1981, p. 5) wrote, with evident sarcasm:

> Central Banking… thrives on a pervasive impression that [it]… is an esoteric art. Access to this art and its proper execution is confined to the initiated elite. The esoteric nature of the art is moreover revealed by an inherent impossibility to articulate its insights in explicit and intelligible words and sentences.

Fifteen years later, in his 1996 Robbins lectures at the London School of Economics, one of the authors of this paper (Alan Blinder (1998), pp. 70-72) expressed a view of what central bank communications *should be*—one that had been lurking around in the underbrush but was far from mainstream at the time:[1]

> Greater openness might actually improve the efficiency of monetary policy… [because] expectations about future central bank behavior provide the essential link between short rates and long rates. A more open central bank… naturally conditions expectations by providing the markets with more information about its own view of the fundamental factors guiding monetary policy…, thereby creating a virtuous circle. By making itself more predictable to the markets, the central bank makes market reactions to monetary policy more predictable to itself. And that makes it possible to do a better job of managing the economy.

Five years later, Michael Woodford (2001, pp. 307 and 312) told an audience of central bankers assembled at the Federal Reserve's 2001 Jackson Hole conference that:

> successful monetary policy is not so much a matter of effective control of overnight interest rates… as of affecting… the evolution of market *expectations*... [Therefore,] transparency is valuable for the effective conduct of monetary policy… this view has become increasingly widespread among central bankers over the past decade.

Notice the progression here: from Brunner's 1981 lament about central bankers' refusal to communicate, to Blinder's 1996 argument that more communication would

[1] For example, the basic idea was stated in Marvin Goodfriend (1991). We thank Michael Woodford for this reference.

enhance the effectiveness of monetary policy, to Woodford's 2001 claims that the essence of monetary policy is the art of managing expectations *and that this was already received wisdom*. Woodford probably exaggerated that last point. But the view that monetary policy is, at least in part, about managing expectations is by now standard fare both in academia and in central banking circles. It is no exaggeration to call this a revolution in thinking.

These new ideas have made a mark on central bank practice as well. At the Federal Reserve, for example, then-Chairman Alan Greenspan, who once prided himself on "mumbling with great incoherence," was by 2003 explicitly managing expectations by telling everyone that the Fed would keep the federal funds rate low "for a considerable period." This guidance was only the latest step in what was, by then, a long march toward greater transparency that began in February 1994 when the Federal Open Market Committee (FOMC) first started announcing its decisions on the federal funds rate target. In May 1999, the FOMC began publishing an assessment of its "bias" with respect to future changes in monetary policy in its statements. It also began issuing fuller statements, even when it was not changing rates. About three years later, it began announcing FOMC votes—with names attached—immediately after each meeting. Starting in February 2005, the FOMC expedited the release of its minutes to make them available before the subsequent FOMC meeting. And most recently, starting in November 2007, the Fed has increased the frequency and expanded the content and horizon of its publicly-released forecasts.

Other central banks have also become remarkably more transparent in the last 10-15 years and are placing much greater weight on their communications. In fact, the Fed is more of a laggard than a leader in this regard. The Reserve Bank of New Zealand and the Bank of

England were early and enthusiastic converts to greater transparency, and Norges Bank (the central bank of Norway) and Sveriges Riksbank (the central bank of Sweden) may now be in the vanguard. Arguably, the European Central Bank (ECB) has been more transparent than the Fed ever since it opened its doors in 1998. More extensive central bank communication is truly a worldwide phenomenon.

One important driver of increased transparency is the notion that more independent central banks should be more accountable—that they have a duty to explain both their actions and the thinking that underlies those actions. But the intellectual arguments just mentioned also played a role. As it became increasingly clear that managing expectations is a useful part of monetary policy, communication policy rose in stature from a nuisance to a key instrument in the central banker's toolkit. In this survey, we concentrate on how central bank communication can be used to manage expectations both by what might be called "creating news" and "reducing noise."

These real-world developments have spawned a huge new scholarly literature on central bank communication—almost all of it written in this decade. While this new literature includes some theoretical contributions, most of it is empirical; and this survey reflects that weighting. Studies of how central bank communications *create news* focus on how, e.g., the central bank's pronouncements influence expectations and therefore move asset prices. In extreme circumstances, communication, used to anchor and guide market expectations, may even become the main tool of monetary policy. Studies of *reducing noise* focus, e.g., on how central bank talk increases the predictability of central bank actions, which should in turn reduce volatility in financial markets. As William Poole (2001, p. 9) put it: "The presumption must be that market participants make more efficient

decisions… when markets can correctly predict central bank actions." In both cases, the central bank's presumed objective is to raise the signal-to-noise ratio, and one major concern of this essay is how successful that effort has been.

That said, communication is no panacea. As with all human endeavors, there are pitfalls and occasional errors. One famous example came in October 2000 when then-ECB President Wim Duisenberg hinted to an interviewer that there would be no further central bank intervention to support the euro. Those words led to an immediate depreciation of the euro and to heavy criticism of Duisenberg. Similarly, when a supposedly off-the-record remark made in April 2006 by Fed Chairman Ben Bernanke, stating that his recent Congressional testimony had been misinterpreted, was reported, markets reacted strongly—as investors concluded that Bernanke was "reversing himself" and saying that interest rates could easily go up.

What constitutes "optimal" communication strategy is by no means clear. And these two examples illustrate that more talk is not always better.[2] The key empirical question is whether communication contributes to the effectiveness of monetary policy by creating genuine news (e.g., by moving short-term interest rates in a desired way) or by reducing noise (e.g., by lowering market uncertainty). There are two main strands in the literature. The first line of research focuses on the impacts of central bank communications on financial markets. The basic idea is that, if communications steer expectations successfully, asset prices should react and policy decisions should become more predictable. Both appear to have happened. The second line of research seeks to relate differences in communication strategies across central banks or across time to differences in economic performance. For

[2] However, in the Bernanke case, the Fed *was* going to raise rates further. So disabusing markets of the false notion that the tightening cycle was finished probably did manage expectations in a constructive way.

example, does announcing a numerical inflation target help anchor the public's long-run inflation expectations? The answer seems to be a qualified yes.

This article reviews the impressive number of mostly empirical studies of central bank communication that have been written in the last several years, mostly focusing on the experience of advanced economies. We take stock of what we now know about how central bank communication can contribute to the effectiveness of monetary policy, and we identify places where additional research is needed. Section 2 discusses in more detail why central bank communication matters. Section 3 examines the practices of three major central banks: the Federal Reserve, the ECB, and the Bank of England. Section 4 reviews the first strand of empirical research mentioned above, and Section 5 discusses the second. Finally, Section 6 provides our answers to the question of how central bank communication can contribute to the effectiveness of monetary policy and identifies avenues for future research.

2. Why does central bank communication matter? Theory

Central bank communication can be defined as the provision of information by the central bank to the public regarding such matters as the objectives of monetary policy, the monetary policy strategy, the economic outlook, and the outlook for future policy decisions.

Nowadays, it is widely accepted that the ability of a central bank to affect the economy depends critically on its ability to influence market expectations about the *future path* of overnight interest rates, and not merely on their current level. The reason is simple. Few, if any, economic decisions hinge on the *overnight* bank rate. According to standard theories of the term structure, interest rates on longer-term instruments should reflect the expected sequence of future overnight rates. So, for example, the n-day rate should be, approximately:

(1) $R_t = \alpha_n + (1/n)(r_t + r^e_{t+1} + r^e_{t+2} + \ldots r^e_{t+n-1}) + \varepsilon_{1t}$,

where r_t is the current overnight rate, r^e_{t+1} is today's expectation of tomorrow's overnight rate (and so on for t+2, t+3,…), α_n is a term premium, and the error term indicates that the term premium might be stochastic.[3] Equation (1) makes it clear that intermediate and long-term rates should depend mostly on the public's *expectations* of *future* central bank policy. Today's overnight interest rate barely matters. A particularly extreme case arises when interest rates get close to their zero lower bound. As long as the *current* overnight rate is stuck at or near zero, central bank communication about expected future rates becomes the essence of monetary policy (Ben S. Bernanke, Vincent Reinhart and Brian Sack 2004; Gauti Eggertsson and Woodford 2003).

Let us now embed this idea in a simple macroeconomic framework designed to illustrate the role of central bank communications, henceforth denoted by the vector s_t (for "signals").[4] Imagine that r in (1) is the short rate and R is the long rate. Then aggregate demand depends on r, R, expected inflation (π^e_t), and a host of other factors which need not be listed explicitly:

(2) $y_t = D(r_t - \pi^e_t, R_t - \pi^e_t, \ldots) + \varepsilon_{2t}$,

The aggregate supply relation could (but need not) be something like the New Keynesian Phillips curve:

(3) $\pi_t = \beta E(\pi_{t+1}) + \gamma(y_t - y^*_t) + \varepsilon_{3t}$,

[3] The time subscript can be thought of as indexing days, months, quarters, etc. The same interpretation holds. The weaknesses of the expectations theory of the term structure are well known. We use it here only for illustrative purposes.

[4] We deliberately keep this model simple for expositional purposes. It could be expanded in several directions. For example, in a New Keynesian setting, expected output would appear on the right-hand side of equation (2)—which would open up another channel by which central bank communications could matter. One could also add a more complex financial sector and/or more complex interactions between the real and financial sectors. None of this is necessary for current purposes.

where π_t is inflation and y_t and y^*_t are, respectively, actual and potential real output. The model could be closed by appending a central bank reaction function (e.g., a "Taylor rule"):

(4) $r_t = G(y_t - y^*_t, \pi_t, \pi^*_t, \dots) + \varepsilon_{4t}$.

where π^* denotes the central bank's inflation target.

Now imagine that the economic environment is *stationary* (that is, equations (1)-(3) do not change over time), that the central bank is *credibly committed* to an unchanging policy rule (4), and that expectations are *rational.* In that unrealistic case, central bank communication has no independent role to play. Any systematic pattern in the way monetary policy is conducted would be correctly inferred from the central bank's observed behavior (Woodford, 2005). In particular, when it comes to predicting future short-term rates, it would suffice to interpret incoming economic data in the light of the central bank's (known) policy rule. Any explicit central bank communication would be redundant. Under Jon Faust and Lars Svensson's (2001, p. 373) definition of central bank transparency—that is, how easily the public can deduce central-bank goals and intentions from observable data—the central bank would be fully transparent without uttering a word.

This extreme case points to four features that have the potential to make central bank communication matter: *nonstationarity* (whether of the economy or the policy rule), the *learning* that is a natural concomitant of such an environment, and either *non-rational expectations* or *asymmetric information* between the public and the central bank. If one or more of these conditions hold, central bank communication can matter.

Needless to say, these four conditions are the norm, not the exception. The real world is constantly changing, as Alan Greenspan never tired of emphasizing.[5] So learning, including learning both by and about the central bank, never ends. Furthermore, it is

[5] See Blinder and Ricardo Reis (2005), especially pages 15-24.

virtually inevitable that the central bank will know more about its own thinking than the public does. In addition, contrary to the impression given by simple Taylor rules, monetary policy decisions depend on much more than current inflation and output gaps (Svensson, 2003). It is also extremely unlikely that the central bank would stick to an unchanged policy rule for long. For example, Bernanke (2004) noted that "specifying a complete and explicit policy rule, from which the central bank would never deviate under any circumstances, is impractical. The problem is that the number of contingencies to which policy might respond is effectively infinite (and, indeed, many are unforeseeable)." Likewise, President Jean-Claude Trichet has repeatedly emphasized that the ECB takes its decisions one step at a time, rather than following a rule.

Under conditions like that, as Bank of England Governor Mervyn King (2005, p.13) has observed, "Rational optimising behaviour is … too demanding, and actual decisions may reflect the use of heuristics." Since central bank communication undoubtedly plays a role in shaping beliefs about those heuristics, it also plays a potentially important role in anchoring expectations.[6] Similarly, Bernanke (2004) used the recent academic literature on adaptive learning to explain why communication affects monetary policy effectiveness. When the public does not know, but instead must estimate, the central bank's reaction function, there is no guarantee that the economy will converge to the rational expectations equilibrium because the public's learning process affects the economy's behavior. The feedback effect of learning on the economy can lead to unstable or indeterminate outcomes—which effective communication by the central bank can help to avoid (see, for example, Stefano Eusepi and Bruce Preston 2007).

[6] For example, King (2005, p. 12) suggests that, under inflation targeting, a good heuristic would be "expect inflation to be equal to target."

In addition, the central bank may have, or may be believed to have, superior information on the economic outlook. Central banks usually devote many more resources than private sector forecasters to forecasting and even to estimating the underlying unobservable state of the economy. Various studies find that financial markets react to information on the outlook that central banks provide (e.g., Malin Andersson, Hans Dillén and Peter Sellin 2006). Apparently, investors update their own views in response to the information conveyed by the central bank. Donald L. Kohn and Sack (2004) argue that private agents may attach special credence to the economic pronouncements of their central bank, especially if the bank has established its *bona fides* as an effective forecaster. They point out that the Federal Reserve has been broadly correct on the direction of the economy and prices over the past two decades, on occasion spotting trends and developments before they were evident to market participants. In a well-known paper, Christina Romer and David Romer (2000) provide statistical evidence that Federal Reserve staff forecasts of inflation were far more accurate than private sector forecasts over a period of several decades.

Central bank communication and learning are inextricably tied, despite a dearth of scholarly attention to that obvious point. There are exceptions, however. In Athanasios Orphanides and John C. Williams (2004), the public is assumed to *know* the form of the equation describing inflation dynamics but to employ standard statistical methods to *learn about* its parameters—which depend on the unobserved objectives and preferences of the central bank. The learning process leads to different behavior than in the rational expectations equilibrium. For example, while people are learning, an increase in inflation may lead the public to revise its estimate of long-run average inflation upward, which, in

turn, raises actual inflation.[7] As Bernanke (2004) pointed out, such a situation opens up a clear opportunity for the central bank to improve economic performance by providing information about its long-run inflation objective. As is true in many contexts, an information problem can be cured by providing more information.

We capture these ideas within our simple framework by replacing the assumption of rational (really, "model-consistent") expectations by an explicit equation for interest rate expectations such as:[8]

$$(5) \quad r^{e}_{t+j} = H_j(y_t, R_t, r_t, \ldots, s_t) + \varepsilon_{5t} ,$$

where s_t is a vector of central bank signals, which might range from crystal clear (e.g., announcing a numerical inflation target) to cryptic (e.g., some of the Fed's words.) Some of these communications, such as the inflation target, might be long-term and durable while others, such as the daily reactions to data releases, might be high-frequency and fleeting—a distinction that will assume some importance in our review of the empirical evidence. There is no need to specify the details of equation (5), which can stand for a variety of possibilities for learning.

In this schema, the total effect of any central bank action operates through at least three distinct channels:

- the direct effect of the overnight rate on aggregate demand—D_r in equation (2)—which is probably quite small;
- the direct effects of central bank signals on expected future short rates: H_s in equation (5), including any learning that might take place;

[7] Some other examples are Glenn Rudebusch and Williams (2008) and Michele Berardi and John Duffy (2007).
[8] We focus on interest rate expectations for simplicity. Expectations of inflation or even of output may be equally important.

- the effect of changes in the short rate on expectations of the entire sequence of *future* short rates, via equations (1) and (5), and their consequent feedback onto long rates, R_t, and therefore onto demand (D_R). This channel will undoubtedly be influenced by the central bank's signals, s_t.

It should be clear from this trichotomy that any account of monetary policy that ignores central bank communication is seriously deficient. Indeed, if the first channel is as unimportant as we suggest, then the communication channels constitute most of the story—which is what Woodford meant and is why many economists these days characterize the job of monetary policy as one of *managing expectations.*

This modern view of monetary policy leads directly to several empirical questions that form the central concerns of this survey. First, what does the vector s_t look like in practice—and why might it vary across central banks? Second, what evidence is there that central bank communications influence expectations *directly*, as posited by equation (5)? Third, how do particular elements of the vector s_t affect measurable variables like interest rates, stock prices, and exchange rates? Fourth, the framework suggests that skillful communications can (a) raise the signal-to-noise ratio, (b) reduce financial market volatility, and (c) lead to better monetary policy outcomes (e.g., lower variances of inflation and output). Is there evidence that it does?

Is there a downside to communication?

All that said, poorly designed or poorly executed communications clearly can do more harm than good; and it is not obvious that a central bank is always better off by saying more. In practice, central banks do limit their communications. In most cases, internal deliberations are kept secret. Only a few central banks project the future path of their policy

rate. (More on this later). And most observe a blackout or "purdah" period before each policy meeting, and in some instances also before important testimonies or reports. The widespread existence of such practices illustrates the conviction of most central bankers that communication can, under certain circumstances, be undesirable and detrimental. In fact, communication during the purdah period has been shown to lead to excessive market volatility (Ehrmann and Fratzscher 2008).

The theoretical literature has not generated clear conclusions regarding the optimal level of transparency (Petra Geraats 2002, Carin van der Cruijsen and Sylvester Eijffinger 2007). The models differ with respect to both which aspects of central bank transparency they consider and their assumptions about how communications influence the monetary transmission mechanism. Looking at real-world central bank behavior, the range of views on what constitutes the "optimal" degree (and types) of communication has clearly evolved over time—mainly in the direction of greater openness. Are there valid—*and empirically relevant*—arguments for limiting communication on monetary policy?[9]

One possible argument dates back to the seminal paper by Alex Cukierman and Allan Meltzer (1986).[10] Their case for obfuscation rested on two assumptions: that only unanticipated money matters, and that the central bank's preferences are not precisely known by the public. Under these assumptions, some degree of opacity enhances the effectiveness of monetary policy because a fully-transparent central bank cannot create surprises. However, two decades later, Pierre Gosselin, Aileen Lotz, and Charles Wyplosz (2007) pointed out that both the view that only unanticipated money matters and the idea

[9] We mention here, but do not discuss further, a few obvious ones: the need to preserve confidentiality, the fact that financial stability sometimes limits central bank talk, and the obvious point that no central bank can divulge what it does not know

[10] For related work see Michelle Garfinkel and Seonghwan Oh (1995), Faust and Svensson (2001), and Henrik Jensen (2002).

that the central bank conceals its preferences in order to pursue its own agenda seem increasingly anachronistic.

Anne Sibert (2006) has recently raised doubts about the Cukierman and Meltzer argument. Her two-period model of a non-transparent central bank focuses on the role of private information. The central bank's welfare is increasing in unexpected inflation (because it increases output) and decreasing in actual inflation. As is typical in models with such objective functions, an unobserved shock that is realized after the public's expectations are formed but before monetary policy is made offers the central bank an opportunity to exploit a short-run Phillips curve tradeoff. Nonetheless, one of Sibert's main conclusions is that both the central bank and society are always better off with increased transparency – mainly because it reduces the inflation bias.

Surely there are limits to how much information can be digested effectively (Daniel Kahneman, 2003). So a central bank should perhaps be wary of communicating about issues on which it receives noisy signals itself—such as the evolution of the economy (as opposed to, say, its upcoming interest rate decisions). This point has been emphasized in the literature on coordination games initiated by Stephen Morris and Hyun Song Shin (2002). Jeffery D. Amato, Morris, and Shin (2002) argue that central bank communication has a dual function: On the one hand, it provides signals about the private information of central banks, and on the other hand, it serves as a coordination device for the beliefs of financial market agents. They argue that central bank communication might be welfare-reducing if agents give too much weight to central bank communication as a focal point, and too little to their own information. The central bank might even coordinate the actions of markets *away from* fundamentals.

But is this likely? Svensson (2006a) shows that the validity of the argument requires that central bank communication has a much lower signal-to-noise ratio than that of private information. He argues that this assumption hardly ever holds in reality. Moreover, Woodford (2005) notes that the Morris-Shin problem is even less likely to arise if the coordination of private agents' actions is a welfare objective *per se*. And Gosselin *et al*. (2006) point out that it is unrealistic to think that a central bank can withhold information as Amato *et al.* suggest. For example, policymakers tacitly reveal some of what they know merely by setting the interest rate. Furthermore, if we focus on providing information about future monetary policy—as opposed to, say, forecasting the stock market or the exchange rate—there is an even simpler and more compelling objection to the Morris-Shin reasoning. Who, after all, knows more about the central bank's intentions than the central bank itself? Thus *honest* central bank talk is almost certain to coordinate beliefs in the *right* direction.

Finally, we should mention the "cacophony problem," pointed out by Blinder (2004, Chapter 2). When monetary policy decisions are taken *and subsequently explained* by a committee rather than by a single individual, there is a danger that too many disparate voices might confuse rather than enlighten the public—especially if the messages appear to conflict. If done poorly, uncoordinated group communication might actually lower, rather than raise, the signal-to-noise ratio. But the appropriate remedy for this problem, should it exist, is clarity, not silence.

Communication is not precommitment

Over the years, many central bankers and economists have at times confused *communication* with *commitment*—or worried out loud that the public might confuse the two. For example, it has been agued that words uttered today might restrict the freedom to

maneuver tomorrow. For example, then-Chairman Paul Volcker defended the Fed's refusal to announce its decisions immediately in 1984 as follows:

> One danger in immediate release of the directive is that certain assumptions might be made that we are committed to certain operations that are, in fact, dependent on future events, and these interpretations and expectations would tend to diminish our needed operational flexibility.[11]

In a similar vein, Alan Greenspan opposed immediate disclosure of the FOMC's decisions in 1989 because "a public announcement requirement also could impede timely and appropriate adjustments to policy."[12] (Yet less than five years later, he *voluntarily* did precisely that.)

From today's standpoint, the objections of Volcker and Greenspan to this minimalist disclosure proposal sound quaint—almost scholastic. While there are cases in which saying something *does* constrain future behavior—as in "giving a verbal commitment" —most central bank communication is not, or need not be, of this nature. In particular, the mere conveyance of information—such as about the policy decision, the inflation target, the forecast, etc.—does not commit the bank to any future action or inaction (although it might hint at such). Even the famous published "forward tracks" of the Reserve Bank of New Zealand (discussed later), which are *conditional forecasts* of its own future behavior, are conditioned on many future variables. That said, the conditional character of such forecasts may be difficult to convey (Otmar Issing 2005).

Of course, there may be cases in which a central bank *wants to* use words to commit itself in some way. For example, Bernanke, Thomas Laubach, Frederic Mishkin, and Adam Posen (1999) argued in favor of inflation targeting on precisely these grounds—as a way to

[11] Quoted in Goodfriend (1986), pp. 76-77. Goodfriend's paper was an early, and at the time highly controversial, critique of the Federal Reserve's secrecy—written by a Fed employee.
[12] Quoted in Blinder (1998), pp. 74-75.

constrain central bank discretion. But that is the exception, not the rule. For the most part, the sorts of communications that we deal with in this paper generally *do not* imply any form of commitment. Since there is already a huge and well-known theoretical literature on the role of commitment in monetary policy, we will not deal with that subject further.[13]

In sum, there are many theoretical reasons why central bank communication should be expected to matter, and many of them imply that skillful communication can improve macroeconomic outcomes. As against this, the arguments against greater transparency seem to be thin gruel: the profession no longer believes that only unanticipated money matters; the Morris-Shin coordination "problem" seems more likely to be an advantage of central bank communication than a disadvantage; and communication need not imply (unwanted) commitment. We turn now from theory to practice.

3. Central bank communication in practice

Many central banks with similar monetary policy objectives nonetheless follow fundamentally different communication policies; and these policies have evolved over time. In our framework, this means that the vector of communication signals, s_t, takes different forms in different times and places. In this section, we illustrate the diversity in current communication practices by examining the different types of signals that central banks send, concentrating mainly on three major central banks: the Federal Reserve System, the Bank of England, and the European Central Bank.[14] We first split the vector of central bank signals, s_t, by *content* (Section 3.1), and then by *sender* (Section 3.2).

3.1 What to communicate

[13] Among the many sources that could be cited, see Richard Clarida, Jordi Gali, and Mark Gertler (1999) or Woodford (2003).

[14] The diversity in communication practices across central banks is also illustrated in Issing (2005), particularly Table A2.

Central banks communicate about at least four different aspects of monetary policy: their overall objectives and strategy, the motives behind a particular policy decision, the economic outlook, and future monetary policy decisions. Central banks' objectives and strategies tend to be more stable, so the corresponding signals show less variability over time than signals about the other three items.

Objectives and strategy

An independent central bank should be given a clearly-defined mandate by its government. Generally, this is done by enunciating central bank objectives, sometimes in quantitative terms. Some central banks that are not given quantitative objectives by their governments have nonetheless decided (or been directed) to provide their own quantification, for at least two reasons. First, numerical targets facilitate accountability, enabling the performance of the central bank to be assessed against its mandated yardstick (Jakob De Haan and Sylvester Eijffinger, 2000). Second, a quantitative objective (or objectives) helps to anchor the expectations of economic agents. In terms of our simple modeling framework, agents' expectation formation in (5) is facilitated by knowing the targets y_t^* and π_t^* that enter the policy rule (4). In turn, well-anchored inflation expectations help to stabilize actual inflation by removing an important source of shocks. However, few if any central banks actually communicate a precise policy rule.[15] Instead, private agents learn about the "rule" both by watching what the central bank *does* and listening to what it *says*.

These accountability and anchoring arguments figure prominently in the debate over inflation targeting (IT) because better and more open communication is often taken to be a

[15] Even formulating an objective function may be a daunting task for a central bank. Some of the difficulties in doing so are described by Mishkin (2004) and Cukierman (2008).

defining virtue of IT. While the Bank of England, for example, sets interest rates independently, its inflation target comes from the Chancellor. The ECB, in contrast, was not given a quantitative objective by the Maastricht Treaty, but provided one for itself as an important part of its monetary policy strategy. Yet a third approach is followed by the Federal Reserve, which has *two* legislated objectives, namely price stability and full employment, neither of which is quantitative as yet. This diversity of practices among otherwise similar central banks is striking, and we will later investigate the extent to which these differences bear on economic outcomes.

Policy decisions

Most central banks nowadays inform the public about their monetary policy decisions on the day they are taken. However, this was not always so. Prominently, the Federal Reserve only began announcing changes in its target federal funds rate on the day of FOMC meetings in February 1994. Before that, markets had to infer the intended funds rate from the type and size of open-market operations--until the decision was published after the subsequent FOMC meeting. Prompt and clear announcement of monetary policy decisions clearly creates news, but it also reduces noise by eliminating any guessing on the part of the public. So this type of central bank communication evidently raises the signal-to-noise ratio. As we will see in the next section, it also leads to improvements in the efficiency of monetary policy.

Practices differ enormously regarding what central banks should or should not say in the statement that accompanies the decision and, presumably, explains it. In particular, central banks apparently disagree over how much should be disclosed about the decision-making process itself, e.g., through the release of minutes and voting records. The ECB

does not publish minutes, and insists that it makes monetary policy decisions by unanimity. The Fed and the Bank of England (BoE) do release minutes (and both recently expedited the release), along with recorded votes. This information is particularly important for the BoE, whose Monetary Policy Committee (MPC) members are individually accountable, and therefore need to have their votes recorded and scrutinized. Interestingly, dissents on the British MPC are much more frequent than they are on the FOMC, where decisions are typically unanimous and dissent connotes fundamental disagreement.[16] Instead of releasing minutes, some central banks (such as the ECB) hold press conferences immediately following their policy decisions. Press conferences may provide less detail than minutes, but they are more timely and more flexible, as they allow the media to ask questions. We will return to this issue later.

The economic outlook

Another important aspect of a central bank's communication strategy is the extent and content of any forward-looking information it provides. This information set includes the central bank's assessment (forecast) of future inflation and economic activity, and its own inclinations regarding future monetary policy decisions. Central banks differ sharply in whether and how they communicate such information.

Inflation-targeting central banks typically provide their assessment of expected future inflation in periodic reports. In that context, the Bank of England's display of probability distributions through "fan charts" has many imitators. However, central banks that are not inflation targeters also often release (some aspects of) their inflation forecasts. In the case of the ECB, this is done through the staff projections (now published four

[16] On this point, see Blinder (2007), Henry Chappell, Rob Roy McGregor and Todd Vermilyea (2004), Ellen Meade and Nathan Sheets (2005), and Laurence H. Meyer (1998).

times a year), which serve as an input to the Governing Council's discussions, but need not be endorsed by it—a very different role from inflation forecasts in an IT strategy. The Federal Reserve keeps its staff projections secret; but it now publishes FOMC forecasts of inflation four times a year. The November 2007 changes in its communication practices increased both the frequency and length of its publicly-released forecasts (see Bernanke 2007). Although these changes did not include the adoption of an explicit inflation target, the new three-year-ahead forecast effectively reveals the inflation rate that policymakers believe is consistent with the Fed's mandate to achieve "stable prices."

Until recently, the diversity across central banks was even wider when it came to the outlook for economic activity. However, the Federal Reserve has now joined the the Bank of England and the ECB in providing more frequent official forecasts of output measures. A number of central banks even publish estimates of the output *gap*. Given the difficulties in measuring and forecasting potential output,[17] the latter option is practiced by only a few central banks (including those of New Zealand, Norway, the Czech Republic, Sweden, and Hungary).

The path of future policy rates

When it comes to likely future policy decisions, many central banks provide some sort of forward guidance, albeit in very different ways. Some, such as the ECB, use indirect signals, often in the form of code words like "vigilance" (David-Jan Jansen and De Haan, 2007). Other central banks are more explicit. The FOMC, for instance, sometimes (but not always) issues a statement with a forward-looking assessment of future monetary policy. These statements, which began in earnest in May 1999, have evolved over time. They were originally phrased in terms of the policy "bias," then in terms of the "balance of risks" for

[17] See, e.g., Orphanides (2001).

the "foreseeable future," and so on. At times, especially during the 2003–2005 period, the FOMC has been quite direct about its expected future path of interest rates.[18]

A few central banks even provide *quantitative* guidance by publishing the numerical path of future policy rates that underlies their macroeconomic forecasts. Sweden and Iceland recently joined a small group that includes New Zealand and Norway in doing so. Some observers view the central bank's forecasting its own future behavior as the last frontier of transparency, and none of the three major central banks on which we have focused have yet been willing to go there. The issue remains highly controversial.[19]

Both Mishkin (2004) and Charles Goodhart (2001) argue against announcing the path of the policy rate on the grounds that it may complicate the committee's decision-making process. It may also complicate communication with the public, which may not understand the *conditional* nature of the projection. In practice, the main concern holding back many central bankers is that such communications might be mistaken for *commitments*. If the projected developments do not materialize, the discrepancy between actual and previously-projected policy might damage the central bank's credibility (Issing 2005). In addition, while forward guidance by the central bank is intended to correct faulty expectations, and thereby reduce misallocations of resources, inaccurate forecasts might actually *induce* such misallocations, e.g., if agents make economic decisions (such as taking on a mortgage) based on the central bank's communication.

To guard against these potential pitfalls, all central banks that provide forward guidance on interest rates emphasize that any forward-looking assessment is conditional on current information—and therefore subject to change. For example, the Riksbank regularly

[18] For a detailed description of the evolution of the FOMC's forward-looking language, see Rudebusch and Williams (2008).
[19] The case in favor is made by Svensson (2006b) and Woodford (2005).

stresses the conditionality of its projected repo rate path by repeating the mantra: "It is a forecast, not a promise."

3.2 How to communicate

Just as the content of signals differs markedly across central banks, so does the choice of communication tools. Central banks can choose from a large menu of communication instruments, and each central bank uses its own mixture.[20] This subsection provides a brief overview of one particularly important aspect of instrument selection, namely, the choice of sender (e.g., whether a signal is sent by the committee or by an individual committee member), which in turn may influence the precision of the signal. When signals are sent by or on behalf of the monetary policy committee, the appropriate content, timing, and channels must all be chosen. Communication by individuals raises further issues—such as whether one member (e.g., the chairman or governor) should serve as spokesperson for the committee, reflecting a more collegial approach to communication, or each member should present his or her own views, representing an individualistic communication strategy.

Communication by committees

The most natural occasion for communication by an MPC as a whole arises on meeting days, when decisions are announced. The timing of this communication and the amount of detail provided differ substantially across central banks. The Federal Reserve provides a short press release containing the decision, a concise (and typically stylized) explanation of its underlying reasoning, and (at times) some forward guidance. The Bank of England's press statement announces the decision, but normally provides an explanation

[20] See Blinder *et al.* (2001) for a detailed, though by now somewhat dated, account and explanation of the various instruments used by central banks.

only when interest rates are changed or when its decision was largely unexpected. Somewhat later, but prior to the subsequent meeting, both central banks provide detailed accounts and explanations of the decisions in the minutes.[21]

By contrast, the ECB not only releases a press statement with the policy decision, but also holds a press conference on the day of Governing Council meetings, including a question and answer session.[22] Compared to the approach of the other two central banks, there are four main differences. First, while providing background information on the rationale for the decision, the ECB press conference is generally less detailed than the minutes of the Bank of England or the Federal Reserve. In particular, it does not provide any information on voting.[23] Second, however, the press conference avoids the substantial time delay of the minutes. Some observers have argued that there is a trade-off between the timeliness of this type of communication and its accuracy (Goodhart 2005), as minutes usually undergo a detailed sanitizing process to put the best possible face on the committee's views and intentions, which delays their release. Third, the Q&A session allows the press to ask follow-up questions and thus can help clarify ambiguities (Ehrmann and Fratzscher 2007a). Fourth, press conferences are typically televised live, which gives the central bank an opportunity to reach out to the broader public.

Another natural communication opportunity for a committee inheres in its legal reporting requirements. For example, each of the three central banks is obliged to provide an annual report and to testify before its legislature. For the ECB and the Federal Reserve,

[21] In addition to the minutes, the Federal Reserve eventually also releases the transcripts of FOMC meetings, albeit only after a five-year lag.
[22] The central banks of the Czech Republic, Japan, New Zealand, Norway, Poland, Sweden, and Switzerland also hold regular press conferences.
[23] For a stimulating debate on the ECB's decisions not to release either minutes or individual voting records, see Willem Buiter (1999) and Issing (1999). One important argument in favor of publishing minutes is that they provide some information about the internal deliberations.

these hearings provide the committee's views, whereas Bank of England testimonies relate more to members' personal views.

Among the most important reporting vehicles are regular publications such as the ECB's *Monthly Bulletin,* which is published one week after each monetary policy meeting and contains both the assessment of economic developments and information on the analytical framework—e.g., models, methods and indicators—used in its decision-making process. For the Bank of England, the most closely-watched reports are its quarterly *Inflation Report*, which sets out the detailed economic analysis that underlies the MPC's decisions and presents the Bank's assessment of the prospects for inflation over the following two years, and the *Quarterly Bulletin,* with its commentary on market developments and monetary policy operations. The publication of the *Inflation Report* is also accompanied by an hour-long press conference. The Federal Reserve's closest counterpart is its semiannual *Monetary Policy Report to the Congress*, presented with the chairman's testimony to Congress.

Communication by individual committee members

Most central banks these days make decisions by committee, reflecting an apparent consensus that doing so leads to superior policy (Blinder, 2004, Chapter 2). But committees come in a wide variety of shapes and sizes. Blinder (2004) distinguishes among three types of committees—individualistic, genuinely collegial, and autocratically collegial—and characterizes the Bank of England's MPC as individualistic, the ECB's Governing Council as genuinely collegial, and the Federal Reserve's FOMC under Alan Greenspan as autocratically collegial.[24] He emphasizes that these distinct types of committees need different communication strategies. In the individualistic case, the diversity of views on the

[24] The FOMC is clearly becoming more genuinely collegial under Ben Bernanke.

committee should be apparent, as a way to help markets understand the degree of uncertainty surrounding monetary policy making. But in the collegial case, a similar diversity of views, if made public, might undermine clarity and common understanding. Therefore, communication should mainly convey the committee's views.

Since the importance of individual views in the communication strategy of a particular MPC will reflect the structure and functioning of the committee, it will vary both across banks and across time. Despite its collegial structure, the Federal Reserve pursues a somewhat individualistic communication strategy, which at times reveals highly diverse opinions across FOMC members. This diversity stands in sharp contrast to the ECB, which has followed a far more collegial communication strategy, often displaying a much higher degree of consistency among the statements of individual committee members (Ehrmann and Fratzscher 2007b).

One difference between communications by individual members and by entire committees is the greater flexibility in timing of the former. Communications by committees are generally pre-scheduled, and thus somewhat inflexible in timing. But changes in the circumstances relevant to monetary policy do not always coincide with meeting dates or testimonies. Furthermore, the central bank might want to provide more guidance to financial markets and the public in times of great uncertainty (Jansen and De Haan 2005). Occasional speeches and interviews by individual committee members between meetings offer a way to communicate changes in views rapidly, if so desired. But the large variation across central banks in the intensity of inter-meeting communication suggests that they differ greatly in how much importance they attach to timeliness.

This section has shown that central bank practices differ enormously, both across central banks and across time. However, there are clear trends toward more timely and more open communication. We have also highlighted the huge variety of signals that are subsumed under the symbol s_t in our schema. This huge heterogeneity in central bank practices raises an obvious question: Are there better and worse ways of communicating? This is clearly an empirical question—and a large one. The rest of this survey will suggest partial answers based on recent empirical evidence.

4. The impact of central bank communication on financial markets

In Section 2, we noted that central bank communications (s_t) influence expectations of future short-term rates (r^e_{t+j}), which, in turn, influence long-term rates and other financial-market prices (R_t). These prices, in turn, influence such macro variables as inflation and output (y_t and π_t). But there are at least two crucial differences between the earlier and later links in this causal chain. Both relate to timing. First, while central bank communications, just like monetary policy, affect financial markets very quickly, interest rates and asset prices affect the economy only gradually—with the proverbial long and variable lags. Second, many factors other than monetary policy influence macro variables such as y_t and π_t. But, at least over the narrow time windows used in many empirical studies, it is arguable that financial market variables are reacting only, or at least mostly, to central bank signals. The upshot is that it is a much easier econometric task to estimate the effects of central bank signals using high-frequency data from financial markets (the subject of this section) than using low-frequency data on macroeconomic performance (the subject of the next section). So we begin with financial market reactions—where there is also, not by coincidence, much more literature to survey.

In Section 4.1, we explain some of the methodological approaches that have been used to identify and measure communication signals. Section 4.2 then moves on to what seems to be the logically first substantive question: Do central bank signals successfully steer market expectations about future monetary policy and therefore render policy more predictable? The evidence suggests that they do, though imperfectly. Section 4.3 takes the next step and inquires about the impact of central bank communications on interest rates and other asset prices—as well as on their volatility. Some of these impacts appear to be sizable, others less so. Finally, in Section 4.4, we turn to the possibility that some communications wind up being *mis*communications, so that, e.g., relatively noisy central bank signals might actually increase, rather than decrease, uncertainty.

4.1 Identifying and measuring communication events

As we have noted, central bank communications differ in many ways and are sometimes difficult to measure. So we focus first on relatively well-defined, high-frequency signals, such as announcements and speeches. It has become standard practice to identify regular pre-scheduled communications (such as announcements of policy decisions) from central bank websites and irregular statements (such as speeches and interviews) through financial newswire reports. However, it is not always straightforward to determine exactly when a communication "event" took place. For example, when a late-Thursday media report on a Tuesday interview with a policymaker causes financial markets to react on Friday morning, was the communication event on Tuesday, Thursday, or Friday?

A second challenge is how to extract the intention or objective behind a policy statement—which is, after all, essential if we want to know whether central bank communications succeed. Different approaches to this issue have emerged. One line of

research does not even attempt to divine directional intent. Kohn and Sack (2004), Ellis Connolly and Marion Kohler (2004), and Rachel Reeves and Michael Sawicki (2007) study instead the effects of central bank communication events on the *volatility* of financial variables. The basic idea is that, if communications affect the returns on financial assets, the volatility of these returns should be higher on days of central bank communications, *ceteris paribus*, because the signals contain news.

Focusing on volatility makes it unnecessary to assign a direction to each statement, which is both a strength and a weakness. The researcher just studies whether central bank communications create news, not whether they move markets in the "right" direction. But there are other problems. One is that many factors, some of them unobservable, affect asset prices. So a rise in observed volatility may reflect the reaction of financial markets to shocks other than central bank communication. A second problem is that communication may be endogenous. A central bank may choose to communicate at a particular time because of a sudden change in the economic outlook or some other news. In this case, asset prices will probably be more volatile on communication days, but not necessarily because of the statement (Reeves and Sawicki 2007). Such endogeneity is less of a problem when the dates of major communications—such as policy decisions and testimonies—are known in advance. But it may be especially problematic for speeches or interviews of committee members, which are flexible in both timing and content (Ehrmann and Fratzscher 2007c).

Kohn and Sack (2004) show that the volatility of various asset prices reacts significantly to statements by the FOMC and its members. They argue that the increased volatility is evidence that central bank communication conveys relevant information to market participants. While this is an important first finding, it leaves several open questions.

Their approach says nothing about whether the policymaker *intended* to move asset prices in the way they did. It also does not tell us whether the rise in volatility reflects an *increase* in market uncertainty or a *decrease* in uncertainty that is outweighed by the immediate effect of the news. In other words, the Kohn and Sack approach establishes that central bank communication creates news, but it cannot determine whether it reduces noise.

Accordingly, a number of subsequent studies have attempted to *quantify* communication in order to assess both the direction and magnitude of its effects on asset prices—and thus to determine to what extent communication has its intended effects.[25] To do so, all statements must be classified according to their content and/or likely intention, and then coded on a numerical scale. Among the types of communications that have been studied in the literature are: statements on monetary policy and communication about the exchange rate (see Jansen and De Haan, 2005, for speeches and interviews by members of the ECB Governing Council), policy inclinations and the economic outlook (see Ehrmann and Fratzscher, 2007b, for communication by committee members at the Federal Reserve, the Bank of England, and the ECB), price stability, the real economy, and monetary indicators (see Helge Berger, Jan-Egbert Sturm and De Haan, 2006, for the introductory statement of the ECB's press conference and Stefan Gerlach, 2007, for the editorial in the ECB's Monthly Bulletin). All of these studies assign negative (positive) values to statements that are perceived as dovish (hawkish), and zero to those that appear to be neutral. Whereas some researchers restrict the coding to directional indications (e.g., Jansen

[25] This is similar to the so-called "narrative approach" to modeling monetary policymaking as, for instance, applied by Romer and Romer (1989) to identify times when the Fed tightened. Similarly, John Boschen and Leonard Mills (1995) generated a discrete measure of the Fed's policy stance taking five different values {-2, -1, 0, 1, 2}, where -2 indicates a very tight policy stance, while 2 indicates a very loose policy stance. Like the Romer and Romer dates, their indicator is based on FOMC minutes. But it is a more informative measure of monetary policy, since it differentiates according to size as well as direction.

and De Haan 2005, Ehrmann and Fratzscher 2007b), others assign a finer grid that is at least suggestive of magnitude (Carlo Rosa and Giovanni Verga 2007, Marie Musard-Gies 2006), e.g. by coding statements on a scale from -2 to +2.

The strength of the coding approach is that it helps us understand whether communication succeeds or fails. But this comes at a price. First, the required classification is necessarily subjective, and there may be misclassifications. This risk can be reduced through content analysis (Ole Holsti 1969), such as when several researchers independently classify statements (as in Berger *et al*., 2006); but it can never be eliminated.

Second, when statements are identified through newswire reports (as in Jansen and De Haan 2005, and Ehrmann and Fratzscher 2007b), communication is not measured at the source, but rather via the media intermediary. But these intermediaries could be selective or misleading in their reporting, or could release the story days after the interview. That said, if researchers are interested in testing market responses to communication, it may make sense to focus on statements that actually reach market participants, and on the content as conveyed by the media.

Third, subjective communication indicators are constructed *ex post* and may not accurately reflect how financial markets understood the signals at the time. The latter is likely to be affected by, e.g., expectations about monetary policy at the time of the statement. As we know, efficient financial markets should react only to the *unexpected* component of a statement.

One approach to these issues, first suggested by Refet Gürkaynak, Sack, and Eric Swanson (2005) and Gürkaynak (2005), and also implemented in Claus Brand, Daniel Buncic and Jarkko Turunen (2006), employs indirect measures derived from financial

market reactions. Gürkaynak *et al*. (2005) use a principal components method that identifies two common factors that describe asset price movements around FOMC announcements. The factor that is orthogonal to the surprise in the federal funds target rate is defined to be the communication effect of the FOMC statement. Brand *et al.* (2006) make use of the timing details of the ECB's meeting day communications. The ECB announces its decisions at 1:45 pm, without any explanatory statements, and then explains the decisions in detail in the press conference 45 minutes later. Because of that delay, the market reaction to the release of the decision can be distinguished from the market reaction to the forward-looking communication by using very high-frequency data. Doing so leads to results that are similar to those obtained with the principal components approach of Gürkaynak *et al*. (2005). Rosa (2008) approximates the unexpected components of both FOMC and ECB statements by estimating forecasting regressions that explain the content of each statement by the preceding statement and the slope of the yield curve just prior to the release of the statement. Interestingly, he finds stronger market reactions to unexpected Fed talk than to unexpected ECB talk.

4.2 Does central bank communication enhance the predictability of monetary policy?

Equation (5) in our conceptual schema posited effects of central bank signals on *expectations* of future monetary policy decisions. It seems that successful central bank communication efforts should make policy more predictable and market expectations of future short rates more accurate. Is it true?

Presumably, central banks are most concerned with *long-term* predictability, which comes when the public develops a genuine understanding of the way the central bank thinks and operates. This is what King (2000) meant by saying provocatively that a central bank

should be "boring." The idea is that the monetary policy reaction function should be so well understood that all the relevant news for financial markets would stem from developments in the economy, not from the actions or words of the central bank. In a similar vein, Blinder (2004, p. 25) suggested that "perhaps the best a central bank can do is to 'teach' the markets its way of thinking."

Unfortunately, long-term predictability is inherently difficult to gauge, so empirical studies typically resort to predictability at short horizons, which is more measurable.[26] Often, financial market prices are used for this purpose. And indeed, research has shown convincingly that the predictability of the interest rate decisions of the major central banks has improved remarkably in recent years—so much so that financial markets' expectations nowadays are generally well aligned with actual decisions (Kerstin Bernoth and Jürgen von Hagen, 2004).

The case of the Federal Reserve, with its continuously increasing level of transparency since the 1980s, has been studied most extensively. Poole and Robert Rasche (2003) provide evidence that the surprise component of policy decisions decreased considerably after February 1994. Joe Lange, Sack, and William Whitesell (2003) show that the ability of Treasury bill yields to predict changes in the federal funds rate some months in advance has increased since the late 1980s. Similarly, Swanson (2006) finds that U.S. financial markets and private sector forecasters have become both better able to forecast the federal funds rate at horizons out to several months and less uncertain about their interest rate forecasts *ex ante*, as measured both by interest rate options and by the cross-sectional variance of interest rate forecasts. Importantly, private sector forecasts of macroeconomic

[26] However, as more central banks begin to announce projected interest-rate paths, it will become possible to study the extent to which these forecasts are incorporated in yield curves.

variables have *not* shown similar improvement, which points toward a specific effect of communication about monetary policy rather than a general decline in macroeconomic volatility (the so-called Great Moderation).

Each of these authors argues that the Federal Reserve's practice of making same-day announcements of monetary policy decisions since February 1994 was an important factor in reducing uncertainty. Furthermore, Selma Demiralp and Oscar Jorda (2002) provide compelling evidence that announcing changes in the intended federal funds rate in real time enabled the Fed to move the actual funds rate with smaller open market operations—which suggests that greater transparency also helped make policy implementation more efficient.

Other types of central bank communication should also enhance predictability. Regardless of whether policymakers have superior information about the economy, communicating their macro forecasts, and possibly also offering indications about the future path of interest rates, should help financial markets predict future central bank actions better. A study by Andrea Fracasso, Hans Genberg, and Wyplosz (2003) sheds light on this issue. They develop a number of indicators of the *quality* of inflation reports for 19 countries and find that higher quality reports are associated with smaller policy surprises. In particular, three subjective indicators—how convincing the report is judged to be, how well it reflects the expertise of the staff, and the quality of the writing style—increase the predictability of interest-rate decisions. These results suggest that clearer communication, which creates a higher signal-to-noise ratio, reduces private agents' uncertainty. While its substantive concern (financial market volatility) is quite different, Jansen (2008) supports this suggestion. Using objective measures of readability to measure the clarity of the

Federal Reserve chairman's semi-annual Humphrey-Hawkins testimonies, he finds that greater clarity often reduces the volatility of interest rates.

Particularly strong effects can be expected if a central bank communicates its intentions about upcoming interest rate decisions. Unfortunately, we know of no research that analyses the effects of the publication of *quantitative* forward guidance, as practiced for years by the Reserve Bank of New Zealand (RNBZ) and, more recently, by Norges Bank and Sveriges Riksbank. New Zealand has the longest history of doing so by far, and David Archer (2004) provides some evidence that market reactions to the publication of the RBNZ's forward interest rate tracks are limited. He suggests that this might reflect the market's understanding of the conditionality of such publications: Market participants might adjust rates to match the Bank's forward path only if they share the underlying assumptions about the evolution of the economy. Beyond this, however, our knowledge of the effects of central bank interest rate projections on the market's understanding of monetary policy is minimal—which is inevitable given the short time span and the small number of central banks that have revealed such information. As more experience is accumulated, e.g., in Norway and Sweden, this will be a high priority area for future research. Already, Norwegian data show that markets do not always uncritically adopt the central bank's projection.

Another type of forward-looking information has been the subject of far more scholarly attention: the *qualitative* guidance provided by the Federal Reserve's (and other central banks') policy "bias" or "balance of risks." As we will see in the next subsection, these statements strongly affect financial markets. Surprisingly, however, their predictive power for future monetary policy is mixed.

In interpreting this evidence on the Fed, it is important to distinguish between bias statements made prior to and after May 1999. Until that date, the FOMC's policy directives were not public information. They were internal declarations of intent, focused narrowly on the inter-meeting period and released only after the *next* FOMC meeting. By contrast, the balance-of-risks assessments since May 1999 are external statements intended to provide markets with information about possible future policy decisions.

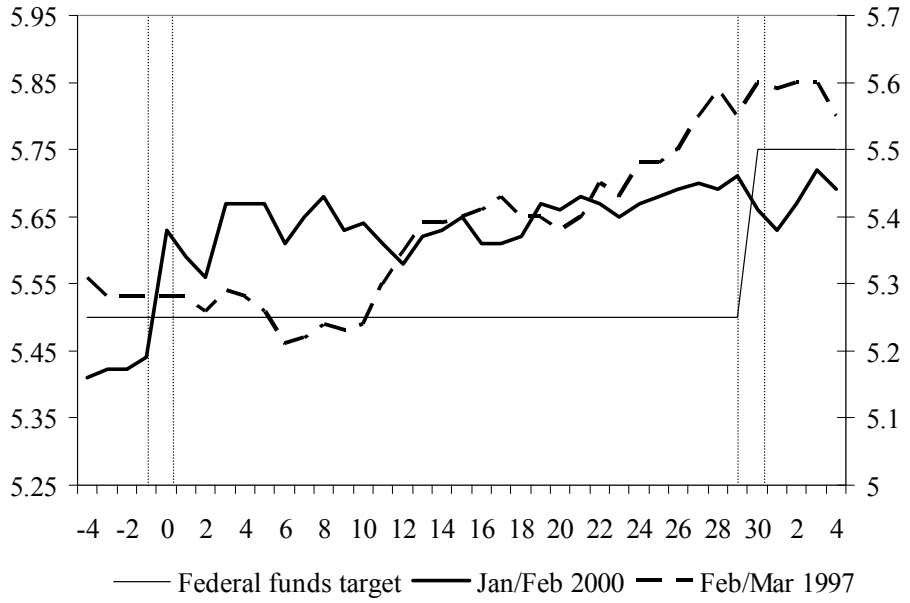
Consistent with their internal use, the pre-1999 bias statements have been found to be a statistically significant indicator of the likelihood and direction of changes in the fed funds target during the subsequent inter-meeting period (John Lapp and Douglas Pearce 2000), but not thereafter (Daniel Thornton and David Wheelock 2000). By contrast, Ehrmann and Fratzscher (2007d) focus on the period since May 1999 and do find consistency with subsequent interest rate moves. Similarly, Michael Pakko (2005) finds that the information content of these statements is a statistically significant variable predicting changes in the funds rate target in the context of an estimated Taylor rule.

Strikingly, however, Ehrmann and Fratzscher (2007d) find that policy decisions after May 1999 have *not* been anticipated better by financial markets than they were before. Chart 1 shows one case in point by displaying the adjustment of the three-month T-bill rate around two pairs of FOMC meetings: one prior to 1999, and one after. In both cases, there was no change at the first of the two meetings (marked as day 0 on the horizontal axis) and then a 25 basis points increase at the second (marked as day 30). Furthermore, in both cases, the FOMC had actually adopted a tightening bias at the meeting on day 0.

In both instances, the decision to raise the funds rate was well predicted by the market: T-bill rates had already increased substantially by the time the FOMC acted. But the

timing was very different in the two cases. When the tightening bias was released on day 0 (solid line), markets did most of the adjusting to the coming interest rate hike within one day. By contrast, when the bias was not released (dashed line), interest rates adjusted much later. Interestingly, T-bill rates started anticipating the Fed's next rate move on the day of Chairman Greenspan's testimony before the US Senate (day 11 in the chart).

Chart 1: Adjustment of market interest rates under alternative disclosure regimes: March 25, 1997 versus February 2, 2000



Source: Ehrmann and Fratzscher (2007d).

Note: Three-month money market rates for the 1997 tightening episode are shown on the right-hand scale whereas those for the 2000 episode are depicted on the left-hand scale. Both tightening days are scaled so as to be shown on day 30 on the horizontal axis. Day 0 refers to the corresponding previous FOMC meetings.

While the chart shows just two exceptionally clear cases, Ehrmann and Fratzscher's statistical tests confirm this as a general tendency. Markets anticipated the Fed's decisions rather well both before and after release of the "balance of risks," but the mechanisms were different. Nowadays, they extract information from the statements, whereas before they relied more on other types of Federal Reserve communication between meetings, such as speeches and testimonies by FOMC members.

A number of recent papers apply Pakko's (2005) approach to the case of the ECB, examining to what extent its communications add information to Taylor-type rules. There is general agreement that models based exclusively on communication variables perform worse than those that include relevant macroeconomic variables, like inflation and output gaps. However, whether communication can improve the fit of standard Taylor rules is in more dispute. Jansen and De Haan (2006a) find that the ECB's communications do not add informational value, but Rosa and Verga (2007) and Friedrich Heinemann and Karin Ullrich (2007) find that the press conferences do.

4.3 Do financial markets respond to (which form of) central bank communication?

The evidence obtained from various event studies demonstrates that central bank talk can move financial markets substantially. An early contribution to this literature is Graeme Guthrie and Julian Wright (2000), who, without distinguishing among different types of communication, provide compelling evidence that communication by the RBNZ has sizable and long-lasting effects on interest rates across the maturity spectrum.

Communication by the Federal Reserve also affects financial markets. As one indication of magnitudes, Table 1 reproduces the results of the seminal study by Kohn and Sack (2004). It is apparent that both FOMC statements and Greenspan testimonies move markets. The statements affect interest rates over short-to-medium-term horizons, while testimonies by then-Chairman Greenspan affect the entire yield curve. In contrast, Kohn and Sack find no significant effects of other Greenspan speeches on financial markets. Reeves and Sawicki (2007), following the same approach, find similar evidence for financial market effects of Bank of England communications. In particular, the release of the MPC

minutes and the Bank's *Inflation Reports* seem to move financial markets significantly.[27] But just like Kohn and Sack, Reeves and Sawicki find that speeches do not.

Table 1

Effects of Federal Reserve communication on the unconditional variance of asset prices

	Increase in unconditional variance due to		
	FOMC statements	*Greenspan testimony*	*Greenspan speeches*
Federal funds futures			
3 months ahead	24.1 ***	10.0 **	1.0
Treasury yields			
2-year	37.5 **	41.4 ***	4.3
10-year	16.4	37.1 ***	3.9
Treasury forward rates			
0 to 1 year ahead	28.9 **	21.8 **	2.1
1 to 2 years ahead	49.7 **	69.3 ***	6.2
2 to 3 years ahead	43.7	57.8 ***	4.1
3 to 4 years ahead	28.7	45.2 **	1.8

Source: Kohn and Sack (2004).
Numbers report the increase in variance (corrected for the effect of macroeconomic and monetary policy surprises) relative to that observed over the week preceding the communication events. Based on daily data.
***/**/* indicate significance at the 99%/95%/90% level.
Sample period: 1989-2003.

However, other studies have found significant effects of speeches on financial markets, perhaps due to different procedures for selecting communication events. Table 2, taken from Ehrmann and Fratzscher (2007b), displays the estimated effects of interviews and speeches by central bank committee members for the Federal Reserve, the Bank of England, and the ECB. It shows that speeches and interviews affect interest rates through most of the maturity spectrum, ranging from very short maturities up to five years. Since these results are based on coded communication events as explained above, they offer directional information, too. Ehrmann and Fratzscher (2007b) find that statements generally

[27] Details of timing may matter here. The *Inflation Reports* are found to move markets only when the frequency of the data is increased from daily (as in Kohn and Sack) to intraday. This is presumably because, with relatively low frequency data, the effects of other news gets mixed up with the communication effects. Using daily data, Connolly and Kohler (2004) and Jaghit Chadha and Charles Nolan (2001) also fail to identify effects of the *Inflation Report* on financial markets.

move financial markets in the intended direction: Statements suggesting tightening lead to higher rates, while statements suggesting easing lead to lower rates. They identify the largest effects for the ECB; statements containing some reference to monetary policy inclinations move interest rates an average of 1.5-2.5 basis points.

Musard-Gies (2006) finds similar results for the ECB, although the short end of the yield curve reacts more sharply to statements than the long end. In stark contrast, statements pertaining to the economic outlook are found to affect mainly US markets, especially at the medium-to-long end of the yield curve, but not (or much less in) markets in the UK or the euro area. Marek Rozkrut, Krzysztof Rybiński, Lucyna Sztaba, and Radosław Szwaja (2007) repeat this analysis for the Czech Republic, Hungary, and Poland. They find that speeches about monetary policy move markets, although to different degrees in each country, but speeches about the economic outlook do not. Finally, Andersson *et al.* (2006) find that speeches by the Riksbank's Executive Board members are important market movers.

Table 2

Effects of interviews and speeches by central bank committee members on interest rates

	Federal Reserve	Bank of England	ECB
3-month interest rates	0.97 ***	0.46 ***	2.05 ***
6-month interest rates	0.46	0.80 ***	1.57 ***
1-year interest rates	0.88 **	0.95 **	2.47 ***
2-year interest rates	1.01 *	0.13	2.48 ***
5-year interest rates	1.14 *	0.15	1.96 ***
10-year interest rates	0.76	-0.21	0.61
20-year interest rates	0.64	0.02	0.44

Source: Ehrmann and Fratzscher (2007b).
Note: The table shows the response of interest rates in basis points to speeches and interviews given by committee members of the different central banks (FOMC members, MPC members and Governing Council members, respectively), containing some non-neutral reference to monetary policy inclinations. Statements that are perceived as suggesting a policy tightening inclination coded as +1, statements that are perceived as suggesting a policy easing inclination as -1. Based on daily data.

*/**/*** denote significance at the 90%/95%/99% level.
Sample period: 1999 (1997 for Bank of England) to 2004.

Overall, the results from several countries are qualitatively consistent, even though heterogeneity in both communication practices and monetary policy strategies makes direct comparisons difficult. Central bank statements and speeches do get the attention of financial market participants, which finds its way into market prices. The quantitative magnitudes naturally differ across central banks. For example, the comprehensive study of different communication tools in six central banks by Connolly and Kohler (2004) finds that monetary policy reports in Australia, Canada, New Zealand, and the US provide information that significantly affects markets' expectations, and thus interest rate futures.[28] Parliamentary hearings affect futures rates in Australia, New Zealand, the UK (albeit only marginally), and the US, but not in Canada or the euro area. However, where they have effects, the impacts of hearings on interest rate expectations are the largest among the various communication tools. There is also some evidence suggesting that the financial markets respond more strongly when communication is from the head of the central bank (Andersson *et al*., 2006, Ehrmann and Fratzscher 2007b).

Central banks explain their policy decisions in different ways. As noted above, Kohn and Sack (2004) find that FOMC statements are important market movers—a result corroborated by Gürkaynak *et al.* (2005). In the principal components approach of the latter (described above), the "path factor," which is related to the forward-looking communication in the statement, affects interest rates across the yield curve--and appears to be the dominant factor at the long end. Similarly, Brand *et al.* (2006) find that the effects of the ECB's communications are particularly sizable at the long end of the yield curve.

[28] For the US, this effect also incorporates any effects of the Congressional testimony that coincides with publication of the report, which might explain the relatively large effect that they find.

The fact that the ECB reveals its decisions 45 minutes before it provides explanations in press conferences offers an interesting opportunity to separate the effects of the two different types of communication events. Ehrmann and Fratzscher (2007a) find that press conferences have, on average, larger effects on asset prices than do announcements of policy decisions. Interestingly, these larger effects on interest rates are accompanied by smaller effects on volatility, indicating that the signal to noise ratio of the press conference communication is high. Why? Their results suggest that the Q&A session gives journalists the opportunity to digest the information provided, to compare it to their own prior information, and to ask clarifying questions where necessary. In fact, there is evidence for such a clarification role. When macroeconomic uncertainty—proxied by the diversity of expectations about data releases—is higher, the market's response to the release of the monetary policy decision is more muted, suggesting that market participants wait for the clarification provided in the press conference before they react.

Central banks often provide substantially more detailed explanations of monetary policy decisions in the minutes of policy meetings. For these to have meaningful news content for financial markets, however, they must be released before the subsequent meeting of the committee. In recent years, both the Federal Reserve and the Bank of England have shortened the lag in releasing their minutes, moving it from after to before the subsequent meeting. As shown by Reeves and Sawicki (2007) and Reinhart and Sack (2006), there are discernable financial market reactions only with more timely release.

4.4 Communication about exchange rates

Central bank communication is not confined to monetary policy proper. While it is beyond the scope of this survey, some central banks make occasional statements on

exchange rates, and this special type of central bank communication has spawned a rapidly-growing literature of its own. We cannot summarize the literature on "oral interventions" here, but a few findings are worth noting because they dovetail nicely with those on interest rates:

- Given the small scale of actual forex interventions relative to today's massive and liquid markets[29]—whether sterilized or unsterilized, words may be more effective than actions, perhaps because words serve as the focal point that coordinates the actions of market participants (Lucio Sarno and Mark Taylor (2001)).[30]
- Communication can complement intervention as well as substitute for it (Fratzscher 2004, forthcoming).
- Central bank talk about exchange rates, just like central bank talk about interest rates, may have to substitute for interest rate cuts when the nominal interest rate reaches its zero lower bound (Lukas Burkhard and Andreas Fischer 2007).

In summary, while the evidence is mixed, on balance it seems fair to say that, if the central bank speaks clearly, its statements can (and apparently do) affect a variety of asset prices in non-negligible ways. While financial market reactions depend on the details of the communications, especially their timeliness and news content, central bank talk apparently does contain news that is quickly incorporated into asset prices.

[29] See Kathryn Dominguez and Jeffrey Frankel (1993).
[30] For some empirical evidence, see Michel Beine, Gust Janssen, and Christelle Lecourt (2006), Fratzscher (2006), Jansen and De Haan (2007), Dominguez and Freyan Panthaki (2007), and Siklos and Bohl (forthcoming).

4.5 Uncertainty in central bank communication

That said, we hasten to highlight an issue that has received scant attention in the literature on central bank communication, yet might be highly relevant to its success: namely, the role of uncertainty in central bank communication.

Uncertainty can arise from various sources. First, inconsistent signals can arise when different members of a monetary policy committee say different things—whether intentionally (e.g., by conducting a debate in public) or unintentionally (e.g., via uncoordinated communication). As Blinder (2007, p. 114) notes, "A central bank that speaks with a cacophony of voices may, in effect, have no voice at all." On the other hand, Bernanke (2004) argues that "the willingness of FOMC members to present their individual perspectives in speeches and other public forums provides the public with useful information about the diversity of views and the balance of opinion on the Committee." Both views have validity. Whether communicating individual committee members' views to the public enlightens or confuses is ultimately an empirical issue. And whether doing so is advisable or inadvisable depends, *inter alia*, on whether the committee has group or individual accountability.

As indicated by Bernanke's comment, FOMC members do indeed communicate with disparate voices (Ehrmann and Fratzscher, 2007b), whereas the ECB generally speaks more with a single voice.[31] However, this has not always been so. Jansen and De Haan (2006b) show that communication about monetary policy inclinations by individual

[31] Issing (1999, pp. 508 and 510) has justified this on the basis of the special circumstances in which the ECB operates: "What matters for transparency is therefore clarity as well as openness. For a new and supranational institution like the ECB, it is particularly important that it sends clear and coherent messages to the markets and the wider public. … Speaking with one voice – or at least speaking 'one language' – is of particular importance for transparency and clarity in the case of the Eurosystem."

members of the Governing Council was relatively high in the initial years of the ECB, but then declined over time.

To what extent do consistency and other characteristics of central bank communication practices affect the predictability of monetary policy? Ehrmann and Fratzscher (2007e) find that more dispersed communication on Federal Reserve monetary policy is associated with less predictable policy decisions at short- and medium-term horizons, and that the magnitude of this effect is sizeable. In another recent study, Rozkrut *et al*. (2007) corroborate this result for the central bank of Hungary; but for the Czech National Bank, more dispersed communication actually enhances the predictability of policy decisions.

There is also evidence that the voting records of the Bank of England's MPC members help predict future policy changes (Petra Gerlach-Kristen 2004). Casting a minority vote appears to be a bigger step, and therefore one that carries more information, than merely expressing a personal dissenting view in public.

Markets, of course, adapt to a central bank's communication style. When central banks give out relatively dispersed communications, financial markets attempt to identify pivotal committee members and attach more weight to their statements. For example, Andersson *et al*. (2006) find evidence that markets react more strongly to statements by the Riksbank's governor. Ehrmann and Fratzscher (2007b) show the same for the Fed chairman. But they find that, in the case of the ECB's more collegial approach to communication, markets react more equally to statements by all Governing Council members.

Central bank communication is also a two-way street: It must have both a transmitter and a receiver, and either could be the source of uncertainty or confusion. Moreover, on the receiving end, the same message might be interpreted differently by different listeners who may have different expectations or believe in different models. One good example is provided by Fracasso *et al*. (2003), who use survey data to analyze whether inflation reports of central banks reach different audiences equally well. They find that the same inflation report is perceived differently by different respondents, and that interest rate surprises tend to increase with the divergence in perceptions. Another example comes from De Haan, Fabian Amtenbrink, and Sandra Waller (2004), who report substantial differences between newspaper reports published the day after ECB policy decisions in the Financial Times (FT) and the Frankfurter Allgemeine Zeitung (FAZ). The British-based FT, which is critical of the money growth pillar in the ECB's monetary policy strategy, tends to pay relatively little attention to it. In sharp contrast, the home-town FAZ clearly supports the idea that money should have a prominent role in the ECB's strategy, and accordingly gives that pillar substantial attention.

A similar phenomenon can arise when the identical message is interpreted differently at different points in time. As one example, Berger, Ehrmann, and Fratzscher (2006) study the quantity and tone of reporting on the ECB's press conference in 57 international and national newspapers in the euro area. They find that the press seems to take its cues from both prior market expectations and the inflationary environment. The tone of the reporting is generally more critical when a policy decision takes markets by surprise. The decisions are also discussed less favorably if the most recent euro area inflation figures exceed 2%, thus violating the ECB's definition of price stability.

5. The impact of central bank communication on inflation performance

We turn now to the harder but perhaps more important question of whether and how central bank communication affects the ultimate goal variables of monetary policy. As we have noted, the combination of long lags in monetary policy and the myriad influences on macro outcomes make it virtually impossible to isolate a link between a specific communication event (such as a speech) and, say, inflation. So this section takes a longer-term perspective by summarizing the literature on how (if at all) differences in communication *strategies*, either over time or across central banks, influence economic outcomes. Since there is almost no evidence on the effects of central bank communication strategies on the real economy, we focus entirely on actual and expected inflation.

5.1 Anchoring inflation expectations

Since a large number of central banks have adopted formal inflation targeting, and others (such as the Swiss National Bank and the ECB) provide quantitative definitions of price stability, a substantial body of research has emerged to assess the effects of an explicit numerical inflation target on inflation outcomes. The central, but not the only, questions are whether clear and consistent communication on the long-term inflation objective (a) anchors the public's long-run inflation expectations and (b) reduces inflation forecast errors.

Before discussing the research findings, a number of caveats are in order. Michael Dueker and Fischer (2006) identify various problems encountered by studies comparing inflation targeters with non-targeters. First, even if a central bank is recognized as an inflation targeter, the dating of the start of the new policy regime is often contentious. For example, Bernanke *et al*. (1999) date the start of inflation targeting (IT) regimes by the *announcement* of the target, while Laurence Ball and Niamh Sheridan (2005) date it by its

implementation. Second, the classification of inflation targeters is not always clear. Some central banks might declare themselves to be inflation targeters, but in practice lack some of the institutional features that are generally associated with inflation targeting. Others might insist that they are not inflation targeters, but share in the characteristic practices (Kenneth Kuttner 2004). Third, results seem to depend crucially on the choice of control group, which varies widely across studies.

Last, and certainly not least, there is a potentially serious endogeneity problem. Just ask yourself: Does the adoption of IT cause inflation to decline, or does high inflation cause a country to adopt IT? Another version of this question is a bit more complex. The introduction of IT can be thought of as a commitment device that changes the central bank's policy and, therefore, economic outcomes. But an alternative interpretation is that the formal adoption of IT is a manifestation of a country's intent to focus more vigorously on reducing inflation. For example, Stephen Cecchetti and Ehrmann (2002) argue that central bankers' aversion to inflation variability increased during the 1990s in many countries, including both inflation targeters and non-targeters.

Subject to these caveats, what does the evidence show? A number of studies find that IT has succeeded in anchoring inflation expectations. Various approaches appear in the literature. David R. Johnson (2003) estimates the determinants of expected inflation (π^e) in the period *prior to* the announcement of an inflation target, and then uses the estimated model to predict π^e under the IT regime. The difference between actual and predicted π^e is then taken to measure the effect of inflation targeting. He finds large reductions in the level of expected inflation in the period after the announcement of inflation targets in Australia, Canada, New Zealand, and Sweden, but not in the United Kingdom.

Another approach, adopted in Johnson (2002) and in Andrew Levin, Fabio Natalucci, and Jeremy Piger (2004), compares targeting to non-targeting countries. Controlling for country, year, trend-inflation, and business-cycle effects, Johnson detects a reduction in inflationary expectations in the IT countries but not in the control group. Levin *et al.* provide evidence that long-term inflation forecasts exhibit a highly significant correlation with a three-year moving average of lagged inflation in the control group, but not in the IT group, suggesting that the inflation-targeting central banks have successfully de-linked expectations from realized inflation.

Gürkaynak, Levin, and Swanson (2006) take a different approach to the same question. They show that long-term inflation expectations derived from index-linked bonds are less responsive to macroeconomic data releases and monetary policy announcements in Sweden and the United Kingdom, two inflation targeting countries, than they are in the United States, which has no explicit inflation objective. Table 3 reports their results for the United States. The first column shows that one-year nominal interest rates respond to the surprise component of macroeconomic releases. The second column indicates that market expectations of inflation over 10 years are also affected significantly. Gürkaynak *et al.* find similar results for the United Kingdom prior to the Bank of England's independence, as shown by the first set of results in Table 4. However, since the BoE became an independent, inflation-targeting central bank in 1998, short-term nominal rates in the UK are still responsive to macroeconomic news but long-term inflation expectations are not—indicating that expectations have been successfully anchored. Interestingly, this evidence suggests that it is not just the existence of an explicit inflation objective that matters, but also the institutional setup of the central bank.

Finally, Jansen and De Haan (2007) find evidence that communication by the ECB regarding risks to price stability (measured by the frequency and intensity of the keyword "vigilance") reduced high-frequency inflationary expectations slightly in the second half of 2005, when a change in the ECB's policy stance became increasingly likely.

Table 3
Responsiveness of US interest rates and inflation expectations to macroeconomic news, 1998-2005

	1-year nominal rate	Inflation compensation
Capacity Utilization	1.57 ***	0.52
Consumer Confidence	1.46 ***	0.43
core CPI	0.98	1.42 **
real GDP (advance)	2.17 **	1.77 **
Initial Jobless Claims	-1.13 ***	-0.51 **
NAPM/ISM Manufacturing	2.28 ***	1.48 **
New Home Sales	0.53	1.42 ***
Nonfarm Payrolls	4.44 ***	0.52
Retail Sales	1.69 ***	0.88
Unemployment Rate	-0.97	-0.11
Monetary Policy	0.23 **	-0.12

Source: Gürkaynak *et al.* (2006).
The coefficients denote the response of, respectively, 1-year spot nominal rates and inflation compensation (based on the different reactions between the nominal and real one-year forward rate from nine to ten years ahead) to standardized macroeconomic surprises.
***/** indicate significance at the 99%/95% level.

Table 4

Responsiveness of UK interest rates and inflation expectations to macroeconomic news

	1993-1997		**1998-2005**	
	1-year nominal rate	Inflation compensation	1-year nominal rate	Inflation compensation
Average Earnings	3.23 ***	0.15	1.81 ***	-0.26
Real GDP (Preliminary)	1.75	1.80 **	2.04 ***	-0.49
Manufacturing Production	0.76	0.33	1.26 ***	-0.04
Producer Price Index	2.13 ***	2.22 **	0.21	-0.22
Core Retail Price Index	2.39 ***	2.60 ***	2.60 ***	-0.76
Retail Sales	2.17 ***	-0.19	1.58 ***	-1.18 ***
Monetary Policy	0.67 ***	-0.60 ***	0.72 ***	-0.13

Source: Gürkaynak *et al.* (2006).
The coefficients denote the response of, respectively, 1-year spot nominal rates and inflation compensation (based on the different reactions between the nominal and real one-year forward rate from nine to ten years ahead) to standardized

macroeconomic surprises.
***/** indicate significance at the 99%/95% level.

These studies all suggest notable effects of inflation targeting on inflation expectations. However, other authors find that expectations are equally well anchored in their control groups, casting doubt on whether the effect identified in the previous studies is causal. Efrem Castelnuovo, Sergio Nicoletti-Altimari, and Diego Rodriguez Palenzuela (2003) find that long-term inflation expectations are well-anchored in all countries in their sample except Japan—regardless of whether the central bank has an inflation target, a quantitative definition of price stability, or no quantified target at all.

A related strand of the literature focuses on the *uncertainty of inflation forecasts* and *inflation forecast errors* rather than on the forecast itself. Ippei Fujiwara (2005) demonstrates that the publication of the Bank of Japan (BoJ)'s inflation and growth forecasts influences those of professional forecasters; in particular, the heterogeneity across forecasters decreases after the BoJ's forecast is published. Pierre Siklos (2002) reports that the adoption of an inflation target leads to lower forecast errors in a panel study of 12 countries.

But subsequent studies do not support his result. Johnson (2002) compares the standard deviation of expected inflation and the average absolute size of inflation forecast errors in a panel of 11 countries, five with inflation targets and six without, over 1984 to 2000. He finds that, when actual inflation falls, uncertainty about future inflation falls in *both* targeting and non-targeting countries. There is also little evidence that average absolute forecast errors are lower in targeting countries. These findings are in line with those of Vittorio Corbo, Oscar Landerretche, and Klaus Schmidt-Hebbel (2002), who analyze how one-step-ahead inflation forecast errors (constructed from estimated country VARs) have

evolved over time. They find that countries that adopted inflation targeting have converged to levels of accuracy similar to those observed in the control group of non-targeters.

Overall, inflationary expectations appear to be generally well anchored, and inflation forecast errors small, in IT countries. And studies of countries undergoing regime changes suggest a causal link between adopting IT and anchoring inflation expectations. However, cross-sectional comparisons yield more ambiguous results; the choice of the control group is apparently crucial. So communication of an explicit inflation target is surely not the only way to anchor expectations.

5.2 Inflation and its dynamics

Building greater trust by credibly communicating a long-term inflation objective may also affect inflation itself. Most obviously, the introduction of IT might reduce the average *level* of inflation—which was certainly the intent of the inventors of IT in New Zealand and most (if not all) of its early adopters. IT might also affect the time-series properties of inflation. If agents believe that inflation will return to its target level more quickly when it deviates from that target, then inflation should be *less persistent*. King (2002) argues that UK inflation has been lower, more stable, and less persistent since inflation targeting was introduced. In support of this argument, Kuttner and Posen (1999) report that the introduction of inflation targeting in Canada and the UK led to lower inflation; but their findings for New Zealand are more mixed.

By contrast, taking a cross-sectional perspective, Ball and Sheridan (2005) find no empirical evidence that inflation targeting improves performance once you control for regression to the mean. (High inflation tends to come down.) IT leads neither to lower levels of inflation, nor to less inflation variability, nor to less persistence. It is true that IT-adopting

countries saw their inflation levels drop after targeting began. But, controlling for the initial level of inflation, the decline is similar for targeters and non-targeters. Ball and Sheridan suggest that this finding reflects the endogeneity issue mentioned earlier: Countries that adopted IT had above-average inflation prior to adoption.

Luke Willard (2006), after dealing with the endogeneity problem in a variety of ways, supports Ball and Sheridan's conclusions. Similarly, Hu Lin and Haichun Ye (2007) study 22 industrial countries over the period 1985–1999 and address the selection problem by using a variety of recently-developed propensity score matching methods. Their main result is that inflation targeting has no significant beneficial effects on targeting countries' inflation or inflation variability. However, Marco Vega and Diego Winkelried (2005), using a similar methodology as Lin and Ye, but for a larger set of countries, do find that IT has helped reduce the level and volatility of inflation. Their results are robust to alternative definitions of treatment and control groups. Analysing ten OECD countries, Alvaro Angeriz and Philip Arestis (forthcoming) also find that IT has gone hand in hand with low inflation. However, IT was often introduced after inflation had already begun its downward trend. They also find that non-targeting countries have succeeded in achieving and maintaining low inflation, but their control group consists of only two cases: the US and the European Economic and Monetary Union.

What is to be made of these disparate results? Mishkin and Schmidt-Hebbel (2007) emphasize that a judgment on the effectiveness of inflation targeting in bringing down inflation depends crucially on the choice of the control group. In particular, IT does not appear to lead to performance superior to that of a group of successful non-targeters. The main benefit they see in inflation targeting is as a disciplinary device that helps potentially

wayward countries move closer to the performance of this successful control group. Bernanke *et al.* (1999) even argue that inflation targets have often been adopted *after* central banks have succeeded in lowering inflation. The primary purpose is to lock in earlier disinflationary gains.

Taking a broader perspective, Antonio Fatás, Ilian Mihov, and Andrew Rose (2007) find for a sample of 42 countries over 40 years that inflation is lower in the presence of *some* quantitative target when other determinants of inflation are taken into account. According to their results, it is *quantification* that matters, not so much the exact target; inflation, exchange rate, or monetary targets are all linked to lower inflation.

Turning to *inflation persistence*, Levin *et al.* (2004) and Luca Benati (forthcoming) find it to be considerably lower—sometimes even negative! —in the presence of inflation targets. Table 5 reports estimates of inflation persistence for a number of countries, separated by monetary regimes according to Benati's classification. It is evident that inflation persistence is particularly low under inflation targeting. But it is also low under the Bretton Woods regime.

Table 5

Measures of inflation persistence

	Canada		**Sweden**		**New Zealand**		**UK**	
	ρ	*90% confidence interval*	ρ	*90% confidence interval*	ρ	*90% confidence interval*	ρ	*90% confidence interval*
Interwar period	0.59	*[0.38;0.80]*	0.70	*[0.42;1.02]*	1.00	*[0.68;1.05]*	--	--
Bretton Woods	0.71	*[0.54;0.88]*	0.29	*[-0.01;0.58]*	0.29	*[0.03;0.55]*	--	--
1971 to inflation targeting	0.90	*[0.72;1.04]*	0.52	*[0.12;1.02]*	0.82	*[0.68;1.01]*	0.78	*[0.62;0.98]*
Inflation Targeting	-0.33	*[-0.76;0.14]*	0.36	*[-0.11;0.92]*	0.41	*[0.00;0.86]*	-0.13	*[-0.45;0.18]*

Source: Benati (forthcoming).
The parameter ρ is the sum of the autoregressive coefficients in a regression of inflation on its own lags. If the confidence interval includes a value of 1, the null hypothesis of a unit root cannot be rejected.

In conclusion, the evidence suggests that adopting an inflation target may have beneficial effects by lowering inflation, by de-linking long-run inflation expectations from

short-run data, and by reducing inflation persistence. However, these estimated benefits may reflect a kind of selectivity bias: They seem to accrue primarily to countries that succeed in stabilizing inflation. There appears to be no systematic difference in the economic performance of low-inflation countries with and without explicit inflation targets.

Accordingly, we conclude that inflation targeting is one way, but certainly not the only way, to control inflation and inflationary expectations. One clear alternative is establishing an anti-inflation track record that allows economic agents to make reasonably accurate inferences about the central bank's objectives and strategy. For example, Linda Goldberg and Michael Klein (2005) analyze the response of euro-area asset prices to US inflation news, and find that the ECB's credibility in financial markets improved after it went into its first tightening cycle.

One issue that remains open is the role of communication during a disinflation. One interesting comparison was made by Michael Bordo, Christopher Erceg, Levin, and Ryan Michaels (2007), who show in three U.S. disinflation episodes that it is particularly important for a central bank to communicate an aggressive policy stance, if it starts with relatively low credibility. It would be nice to know whether the introduction of an inflation target can facilitate the disinflation process under these circumstances, and what alternatives a central bank should (and should not) consider.

6. Assessment and issues for future research

To what extent, then, is central bank communication a useful instrument for policy makers? In 2001, Blinder *et al*. (p. 9) wrote: "To date, there is no research to report on." Since then, there has been a veritable explosion of research on central bank communication. While the evidence reviewed here is not entirely one-sided, it seems safe to conclude that

communication can be an important and powerful part of the central bank's toolkit since it has the ability to move financial markets, to improve the predictability of monetary policy, and the potential to help the monetary authorities achieve macroeconomic objectives such as low and stable inflation.

More specifically, our review of the rapidly-growing empirical literature, virtually all of which was written in this decade, suggests that:

- No consensus has yet emerged on what communication policies constitute "best practice" for central banks. Practices, in fact, differ substantially and are evolving.
- The predictability of monetary policy decisions has improved notably in many countries. With only a few exceptions, the empirical studies to date suggest that more and better central bank communication contributed to this improvement by "reducing noise." We are tempted to call this issue "closed" but, of course, future research can always overturn a previous conclusion.
- That said, the predictability of monetary policy appears to be degraded somewhat when central banks speak with too many conflicting voices—as sometimes happens at the Federal Reserve, for example.
- What might be called "short-run" central bank communication—that is, disclosing central bank views on, e.g., the outlook for the economy and monetary policy—has a wide variety of impacts on financial markets, including on the levels and volatilities of various interest rates. Official statements, reports, and minutes appear to have the clearest and most consistent empirical effects on financial markets. The evidence on the impacts of speeches is more mixed. But

it, too, is mainly supportive of the idea that central bank communication "creates news." The distinctions among different forms of communication merit further empirical evaluation (e.g., what works best?)..

- The limited number of studies that try to assess the directional intent of the central bank's messages generally find that markets move in the "right" direction—that is, what used to be called "announcement effects" help the central bank rather than hinder it. But there has been relatively little such research to date. We need more evidence in order to be convinced on this point, plus further analysis of how, if at all, creating this type of directional news raises welfare.
- Regarding what might be called "long-run" central bank communication, the empirical evidence so far is largely limited to one set of questions: the effects of inflation targeting (IT) on inflation outcomes. Those are good questions, but not the only ones. So research on the links between communication and other macro variables is essential.
- For a variety of reasons, isolating clear effects of IT turns out to be harder than might be expected. But there is clear evidence that IT helps anchor inflationary expectations. (However, it is not the only way to do so.) The evidence that adopting IT leads to lower or less variable inflation is far less compelling. On such questions, the choice of the control groups seems to be critical.

This list of research findings constitutes a quantum leap over what we knew at the start of the decade, which was almost nothing. But there is a lot more to learn. For example, the publication of projected paths for the central bank's policy rate appears to be the "new

frontier" in central bank communication. But it has been practiced in so few countries for so few years that we have little empirical knowledge of its effects as yet. As more data accumulates, this should be a high-priority area for future research.

Another important, but barely explored issue is what constitutes "optimal" communication policy, and how that depends on the institutional environment in which a central bank operates, the nature of its decisionmaking process, and the structure of its monetary policy committee. For example, the Bank of England has a committee of nine members who are individually accountable, while both the FOMC and the ECB Governing Council are larger committees with group accountability, and the Reserve Bank of New Zealand still has a single decisionmaker. Central banks need to tailor their communication strategies to these and other institutional features. Research on that important topic has barely begun.

Finally, virtually all the research to date has focused on central bank communication with the financial markets. It may be time to pay some attention to communication with the general public. Admittedly, studying communication with the general public will pose new challenges to researchers—not least because financial market prices will be less relevant. But the issues are at least as important. In the end, it is the general public that gives central banks their democratic legitimacy, and hence their independence.

References

Amato, Jeffery D., Stephen Morris, and Hyun Song Shin. 2002. "Communication and Monetary Policy." *Oxford Review of Economic Policy*, 18(4): 495-503.

Andersson, Malin, Hans Dillén, and Peter Sellin. 2006. "Monetary Policy Signaling and Movements in the Term Structure of Interest Rates." *Journal of Monetary Economics*, 53(8): 1815-55.

Angeriz, Alvaro, and Philip Arestis. Forthcoming. "Assessing inflation targeting through intervention analysis." *Oxford Economic Papers.*

Archer, David. 2004. "Communication with the Public." In *Practical Experience with Inflation Targeting*, Prague: Czech National Bank, 145-55.

Ball, Laurence, and Niamh Sheridan. 2005. "Does Inflation Targeting Matter?" In *The Inflation-Targeting Debate,* ed. Ben S. Bernanke and Michael Woodford. Chicago: University of Chicago Press, 249-76.

Beine, Michel, Gust Janssen, and Christelle Lecourt. 2006. "Should Central Bankers Talk to the FX Market?" Unpublished, available at http://homepages.ulb.ac.be/~mbeine/

Benati, Luca. Forthcoming. "Investigating Inflation Persistence across Monetary Regimes." *Quarterly Journal of Economics.*

Berardi, Michele, and John Duffy. 2007. "The Value of Central Bank Transparency When Agents Are Learning." *European Journal of Political Economy*, 23(1): 9-29.

Berger, Helge, Michael Ehrmann, and Marcel Fratzscher. 2006. "Monetary Policy in the Media." ECB Working Paper 679.

Berger, Helge, Jakob De Haan, and Jan-Egbert Sturm. 2006. "Does Money Matter in the ECB Strategy? New Evidence Based on ECB Communication.", CESifo Working Paper 1652.

Bernanke, Ben. 2004. "Fedspeak." Remarks at the Meetings of the American Economic Association, San Diego, available at: http://www.federalreserve.gov/boarddocs/speeches/2004/200401032/default.htm.

Bernanke, Ben. 2007. "Federal Reserve Communications." Remarks at the Cato Institute 25th Annual Monetary Conference, Washington, D.C., 14 November 2007, available at http://www.federalreserve.gov/newsevents/speech/bernanke20071114a.htm.

Bernanke, Ben, Thomas Laubach, Frederic S Mishkin, and Adam S. Posen. 1999. *Inflation Targeting: Lessons from the International Experience.* Princeton, New Jersey: Princeton University Press.

Bernanke, Ben, Vincent R. Reinhart, and Brian P. Sack. 2004. "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment." *Brookings Papers on Economic Activity,* 2: 1-100.

Bernoth, Kirsten, and Jürgen von Hagen. 2004. "The Euribor Futures Market: Efficiency and the Impact of ECB Policy Announcements." *International Finance,* 7 (1): 1-24.

Blinder, Alan S. 1998. *Central Banking in Theory and Practice*. Cambridge MA: MIT Press.

Blinder, Alan S. 2004. *The Quiet Revolution: Central Banking Goes Modern*. New Haven, CN: Yale University Press.

Blinder, Alan S. 2007. "Monetary Policy by Committee: Why and How?" *European Journal of Political Economy*, 23(1), 106-23.

Blinder, Alan S., Charles Goodhart, Philipp Hildebrand, David Lipton, and Charles Wyplosz. 2001. *How Do Central Banks Talk?* Geneva Reports on the World Economy 3, Geneva/London: ICMB/CEPR.

Blinder, Alan S., and Ricardo Reis. 2005. "Understanding the Greenspan Standard." In *The Greenspan Era: Lessons for the Future*, Kansas City: Federal Reserve Bank of Kansas City, 11-96.

Bordo, Michael D., Christopher Erceg, Andrew T. Levin and Ryan Michaels. 2007. "Three Great American Disinflations." NBER Working Paper 12982.

Boschen, John F., and Leonard O. Mills. 1995. "The Relation between Narrative and Money Market Indicators of Monetary Policy." *Economic Inquiry*, 33(1): 24-44.

Brand, Claus, Daniel Buncic, and Jarkko Turunen. 2006. "The Impact of ECB Monetary Policy Decisions and Communication on the Yield Curve." ECB Working Paper 657.

Brunner, Karl. 1981. "The Art of Central Banking." Center for Research in Government Policy and Business, University of Rochester, Working Paper GPB 81-6.

Buiter, Willem. 1999. "Alice in Euroland." *Journal of Common Market Studies*, 37(2): 181–209.

Burkhard, Lukas, and Andreas M. Fischer. 2007. "Communicating Policy Options at the Zero Bound." CEPR Discussion Paper 6563.

Castelnuovo, Efrem, Sergio Nicoletti-Altimari, and Diego Rodriguez Palenzuela. 2003. "Definition of Price Stability, Range and Point Inflation Targets: The Anchoring of

Long-Term Inflation Expectations." In *Background Studies for the ECB's Evaluation of Its Monetary Policy Strategy,* ed. O. Issing. Frankfurt-am-Main, Germany: European Central Bank, 43-90.

Cecchetti, Stephen, and Michael Ehrmann. 2002. "Does Inflation Targeting Increase Output Volatility? An International Comparison of Policymakers Preferences and Outcomes." In *Monetary Policy: Rules and Transmission Mechanisms,* ed. Norman Loayza and Klaus Schmidt-Hebbel. Santiago: Central Bank of Chile, 247-74.

Chadha, Jaghit S., and Charles Nolan. 2001. "Inflation Targeting, Transparency and Interest Rate Volatility: Ditching 'Monetary Mystique' in the UK." *Journal of Macroeconomics*, 23(3): 349-66.

Chappell, Henry W., Rob Roy McGregor, and Todd A. Vermilyea. 2004. *Committee Decisions on Monetary Policy. Evidence from Historical Records of the Federal Open Market Committee.* Cambridge (MA): MIT Press.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37(4): 1661-1707.

Connolly, Ellis, and Marion Kohler. 2004. "News and Interest Rate Expectations: A Study of Six Central Banks." In *The Future of Inflation Targeting*, ed. Christopher Kent and Simon Guttman. Sydney: Reserve Bank of Australia, 108-34.

Corbo, Vittorio, Oscar Landerretche, and Klaus Schmidt-Hebbel. 2002. "Does Inflation Targeting Make a Difference?" In *Inflation Targeting: Design, Performance, Challenges*, ed. Norman Loayza and Raimundo Soto. Santiago: Central Bank of Chile, 221-69.

Cruijsen, Carin A.B. van der and Sylvester C.W. Eijffinger. 2007. "The Economic Impact of Central Bank Transparency: A Survey," CEPR Discussion Paper, No. 6070.

Cukierman, Alex. 2008. "The Limits of Transparency." ECB Working Paper.

Cukierman, Alex, and Allan Meltzer. 1986. "A Theory of Ambiguity, Credibility and Inflation under Discretion and Asymmetric Information." *Econometrica*, 54(4): 1099-1128.

De Haan, Jakob, and Sylvester C.W. Eijffinger. 2000. "The Democratic Accountability of the European Central Bank: A Comment on Two Fairy-tales." *Journal of Common Market Studies,* 38(3): 393-407.

De Haan, Jakob, Fabian Amtenbrink, and Sandra Waller. 2004. "The Transparency and Credibility of the European Central Bank." *Journal of Common Market Studies*, 42(4): 775-94.

Demiralp, Selma, and Oscar Jorda, 2002. "The Announcement Effect: Evidence from Open Market Desk Data." *Federal Reserve Bank of New York Economic Policy Review*, 8: 29-48.

Dominguez, Kathryn M., and Jeffrey A. Frankel. 1993. "Does Foreign Exchange Intervention Matter? The Portfolio Effect." *American Economic Review,* 83(5): 1356–69.

Dominguez, Kathryn M., and Freyan Panthaki. 2007. "The Influence of Actual and Unrequited Interventions." *International Journal of Finance and Economics*, 12(2): 171-200.

Dueker, Michael J., and Andreas M. Fischer. 2006. "Do Inflation Targeters Outperform Non-targeters?", *Federal Reserve Bank of St. Louis Review*, 88(5): 431-50.

Eggertsson, Gauti, and Michael Woodford 2003. "The Zero Bound on Short-Term Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 1:139–211.

Ehrmann, Michael, and Marcel Fratzscher. 2007a. "Explaining Monetary Policy Decisions in Press Conferences." ECB Working Paper 767.

Ehrmann, Michael, and Marcel Fratzscher. 2007b. "Communication by Central Bank Committee Members: Different Strategies, Same Effectiveness?" *Journal of Money, Credit, and Banking*, 39(2-3): 509-41.

Ehrmann, Michael, and Marcel Fratzscher. 2007c. "The Timing of Central Bank Communication." *European Journal of Political Economy*, 23(1): 124-45.

Ehrmann, Michael, and Marcel Fratzscher. 2007d. "Transparency, Disclosure and the Federal Reserve." *International Journal of Central Banking*, 3(1): 179-225.

Ehrmann, Michael, and Marcel Fratzscher. 2007e. "Social value of public information: testing the limits to transparency." ECB Working Paper 821.

Ehrmann, Michael, and Marcel Fratzscher. 2008. "Purdah – On the Rationale for Central Bank Silence Around Policy Meetings." Forthcoming, *Journal of Money, Credit, and Banking*.

Eusepi, Stefano, and Bruce Preston. 2007. "Central Bank Communication and Expectations Stabilization." NBER Working Paper 13259.

Fatás, Antonio, Ilian Mihov, and Andrew Rose. 2007. "Quantitative Goals for Monetary Policy." *Journal of Money, Credit, and, Banking*, 39(5): 1163-76.

Faust, Jon, and Lars E.O. Svensson. 2001. "Transparency and Credibility: Monetary Policy with Unobservable Goals." *International Economic Review,* 42(2): 369-97.

Fracasso, Andrea, Hans Genberg, and Charles Wyplosz. 2003. "How do Central Banks Write?" *Geneva Reports on the World Economy 2.* Geneva/London: ICMB and CEPR.

Fratzscher, Marcel. 2004. "Communication and Exchange Rate Policy." ECB Working Paper 363.

Fratzscher, Marcel. Forthcoming. "Oral Interventions versus Actual Interventions in FX Markets–An Event-Study Approach," *The Economic Journal*.

Fratzscher, Marcel. 2006. "On the Long-Term Effectiveness of Exchange Rate Communication and Interventions." *Journal of International Money and Finance*, 25(1): 146-67.

Fujiwara, Ippei. 2005. "Is the Central Bank's Publication of Economic Forecasts Influential?" *Economics Letters*, 89(3): 255–61.

Garfinkel, Michelle and Seonghwan Oh. 2005. "When and How Much to Talk: Credibility and Flexibility in Monetary Policy with Private Information," *Journal of Monetary Economics*, 35(2): 341-357.

Geraats, Petra. 2002. "Central Bank Transparency." *The Economic Journal*, 112 (483): 532–65.

Gerlach, Stefan. 2007. "Interest Rate Setting by the ECB, 1999-2006: Words and Deeds." *International Journal of Central Banking*, 3(3): 1-46.

Gerlach-Kristen, Petra. 2004. "Is the MPC's Voting Record Informative about Future UK Monetary Policy?" *Scandinavian Journal of Economics,* 106(2): 299-313.

Goldberg, Linda, and Michael Klein. 2005. "Establishing Credibility: Evolving Perceptions of the European Central Bank." NBER Working Paper 11792.

Goodfriend, Marvin. 1986. “Monetary Mystique: Secrecy and Central Banking.” *Journal of Monetary Economics*, 17(1): 63-92.

Goodfriend, Marvin. (1991). “Interest Rates and the Conduct of Monetary Policy.” *Carnegie-Rochester Conference Series on Public Policy*, 34: 7-30.

Goodhart, Charles A.E.. 2001. “Monetary policy transmission lags and the formulation of the policy decision on interest rates”. *Federal Reserve Bank of St. Louis Review*, 83(4), pp. 165-81.

Goodhart, Charles A.E.. 2005. “Dear Jean-Claude.” *Central Banking*, 16(1): 32-36.

Gosselin, Pierre, Aileen Lotz, and Charles Wyplosz. 2007. “Interest Rate Signal;s and Central Bank Transparency,” CEPR Discussion Paper 6454.

Gürkaynak, Refet S. 2005. “Using Federal Funds Futures Contracts for Monetary Policy Analysis.” Federal Reserve Board, Finance and Economic Discussion Series 2005-29.

Gürkaynak, Refet S., Andrew Levin, and Eric Swanson. 2006. “Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-Term Bond Yields in the US, UK and Sweden.” Federal Reserve Bank of San Francisco Working Paper 2006-09.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements.” *International Journal of Central Banking*, 1(1): 55-93.

Guthrie, Graeme, and Julian Wright. 2000. “Open Mouth Operations.” *Journal of Monetary Economics*, 46(2): 489-516.

Heinemann, Friedrich, and Karin Ullrich. 2007. "Does it Pay to Watch Central Bankers Lips? The Information Content of ECB Wording." *Swiss Journal of Economics and Statistics* 3: 155-85.

Holsti, Ole R. 1969. *Content Analysis for Social Sciences and Humanities*. Reading: Addison-Wesley.

Issing, Otmar 1999. "The Eurosystem: Transparent and Accountable, or 'Willem in Euroland'", *Journal of Common Market Studies*, 37(3): 503–519.

Issing, Otmar. 2005. "Communication, Transparency, Accountability – Monetary Policy in the Twenty-First Century." *Federal Reserve Bank of St. Louis Review*, 87(2): 65-83.

Jansen, David-Jan. 2008. "Does the Clarity of Central Bank Communication Affect Volatility in Financial Markets?" unpublished paper, De Nederlandsche Bank, March.

Jansen, David-Jan, and Jakob De Haan. 2005. "Talking Heads: The Effects of ECB Statements on the Euro-Dollar Exchange Rate." *Journal of International Money and Finance*, 24(2): 343-361.

Jansen, David-Jan, and Jakob De Haan. 2006a. "Does ECB communication help in predicting its interest rate decisions?" CESifo Working Paper 1804.

Jansen, David-Jan, and Jakob De Haan. 2006b. "Look Who's Talking: ECB Communication During the First Years of EMU." *International Journal of Finance and Economics*, 11(3): 219-228.

Jansen, David-Jan, and Jakob De Haan. 2007. "Were Verbal Efforts to Support The Euro Effective? A High-frequency Analysis of ECB Statements." *European Journal of Political Economy*, 23(1): 245-259.

Jansen, David-Jan, and Jakob De Haan. 2007. "The Importance of Being Vigilant: Has ECB Communication Influenced Euro Area Inflation Expectations?" CESifo Working Paper 2134.

Jensen, Henrik. 2002. "Optimal Degrees of Transparency in Monetary Policy Making." *Scandinavian Journal of Economics*, 104(3): 399-422.

Johnson, David R. 2002. "The Effect of Inflation Targeting on the Behavior of Expected Inflation: Evidence from an 11 Country Panel." *Journal of Monetary Economics*, 49(8): 1493–1519.

Johnson, David R. 2003. "The Effect of Inflation Targets on the Level of Expected Inflation in Five Countries." *Review of Economics and Statistics*, 55(4): 1076–81.

Kahnemann, Daniel. 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review,* 93(5): 1449-75.

King, Mervyn. 2000. "Monetary Theory: Policy in Practice." Address to the joint luncheon of the American Economic Association and the American Finance Association, Boston, available at

http://www.bankofengland.co.uk/publications/speeches/2000/speech67.htm.

King, Mervyn. 2002. "The Inflation Target Ten Years On." Speech delivered to the London School of Economics, available at

http://www.bankofengland.co.uk/publications/speeches/2002/speech181.pdf

King, Mervyn. 2005. "Monetary Policy: Practice Ahead of Theory." 2005 Mais Lecture, delivered at Cass Business School, London; available at:

http://www.bankofengland.co.uk/publications/speeches/2005/speech245.pdf

Kohn, Donald L., and Brian Sack (2004). "Central Bank Talk: Does it Matter and Why?" In: *Macroeconomics, Monetary Policy, and Financial Stability*, Ottawa: Bank of Canada, 175-206.

Kuttner, Kenneth. 2004. "A Snapshot of Inflation Targeting in its Adolescence." In *The Future of Inflation Targeting,* ed. C. Kent and S. Guttmann. Sydney: Reserve Bank of Australia, 6-47.

Kuttner, Kenneth N., and Adam S. Posen 1999. "Does Talk Matter After All? Inflation Targeting and Central Bank Behavior." Federal Reserve Bank of New York Staff Report 88.

Lange, Joe, Brian Sack, and William Whitesell. 2003. "Anticipations of Monetary Policy in Financial Markets." *Journal of Money, Credit, and Banking,* 35(6): 889-909.

Lapp, John S., and Douglas K. Pearce. 2000. "Does a Bias in FOMC Policy Directives Help Predict Inter-Meeting Policy Changes?" *Journal of Money, Credit, and Banking*, 32(3): 435–41.

Levin, Andrew T., Fabio M. Natalucci, and Jeremy M. Piger. 2004. "The Macroeconomic Effects of Inflation Targeting." *Federal Reserve Bank of St. Louis Review*, 86(4): 51–80.

Lin, Shu, and Haichun Ye. 2007. "Does Inflation Targeting Really Make a Difference? Evaluating the Treatment Effect of Inflation Targeting in Seven Industrial Countries." *Journal of Monetary Economics*, 54(8): 2521-33.

Meade, Ellen E., and D. Nathan Sheets. 2005. "Regional Influences on FOMC Voting Patterns." *Journal of Money, Credit, and Banking*, 27(4): 661-78.

Meyer, Larry. 1998. "Come With Me to the FOMC. The Gillis Lecture." Willamette University, available at http://www.federalreserve.gov/boarddocs/speeches/1998/199804022.htm.

Mishkin, Frederic S. 2004. "Can Central Bank Transparency Go Too Far?" In *The Future of Inflation Targeting*, ed. Christopher Kent and Simon Guttmann. Sydney: Reserve Bank of Australia, 48-65.

Mishkin, Frederic S. and Klaus Schmidt-Hebbel. 2007. "Does Inflation Targeting Make a Difference?" NBER Working Papers 12876.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review*, 92(5): 1521-1534.

Musard-Gies, Marie. 2006. "Do ECB's Statements Steer Short-Term and Long-Term Interest Rates in the Euro-Zone?" *The Manchester School*, 74 (Supplement): 116–39.

Orphanides, Athanasios 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review*, 91: 964-85.

Orphanides, Athanasios, and John C. Williams (2004). "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation Targeting Debate*, ed. Ben Bernanke and Michael Woodford. Chicago: University of Chicago Press, 201-34.

Pakko, Michael R. 2005. "On the Information Content of FOMC Policy Statements: Evidence from a Taylor-rule Perspective." *Economic Inquiry*, 43(3): 558–69.

Poole, William. 2001. "Expectations." *Federal Reserve Bank of St. Louis Review*, 83(1): 1–10.

Poole, William, and Robert H. Rasche 2003. "The Impact of Changes in FOMC Disclosure Practices on the Transparency of Monetary Policy: Are Markets and the FOMC Better "Synched"?", *Federal Reserve Bank of St. Louis Review*, 85(1): 1-10.

Reeves, Rachel, and Michael Sawicki. 2007. "Do Financial Markets React to Bank of England Communication?" *European Journal of Political Economy*, 23(1): 207-27.

Reinhart, Vincent, and Brian Sack 2006. "Grading the Federal Open Market Committee's Communications." Unpublished.

Romer, Christina, and David Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." In *NBER Macroeconomics Annual 1989*, ed. Olivier Jean Blanchard and Stanley Fischer. Cambridge (MA): MIT Press, 121-70.

Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90(3): 429–457.

Rosa, Carlo. 2008. "Talking Less and Moving the Markets More: Is this the Recipe for Monetary Policy Effectiveness? Evidence from the ECB and the Fed." Unpublished, available at http://carlorosa1.googlepages.com/research.

Rosa, Carlo, and Giovanni Verga. 2007. "On the Consistency and Effectiveness of Central Bank Communication: Evidence from the ECB." *European Journal of Political Economy*, 23(1): 146-75.

Rozkrut, Marek, Krzysztof Rybiński, Lucyna Sztaba, and Radosław Szwaja. 2007. "Quest for Central Bank Communication. Does It Pay To Be "Talkative"? *European Journal of Political Economy*, 23(1): 67-87.

Rudebusch, Glenn, and John Williams. 2008. “Revealing the Secrets of The Temple: The Value of Publishing Central Bank Interest Rate Projections.” In *Asset Prices and Monetary Policy*, ed. John Y. Campbell. Chicago: University of Chicago Press.

Sarno, Lucio, and Mark P. Taylor 2001. “Official Intervention in Foreign Exchange Market: Is It Effective and, If So, How Does It Work?, *Journal of Economic Literature*, 39(3): 839–68.

Sibert, Anne C. 2006. “Is Central Bank Transparency Desirable?” CEPR Discussion Paper 5641.

Siklos, Pierre L. 2002. *The Changing Face of Central Banking: Evolutionary Trends Since World War II*, Cambridge University Press.

Siklos, Pierre L. and Martin T. Bohl. Forthcoming. "Policy Words and Policy Deeds: The ECB and the Euro," forthcoming in *International Journal of Finance and Economics.*

Swanson, Eric T. 2006. “Federal Reserve Transparency and Financial Market Forecasts of Short-Term Interest Rates.” *Journal of Money, Credit, and Banking*, 38 (3): 791-819.

Svensson, Lars E.O. (2003). “What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules.” *Journal of Economic Literature*, 41(2): 427-477.

Svensson, Lars E.O. 2006a. “Social Value of Public Information: Morris and Shin (2002) Is Actually Pro Transparency, Not Con.” *American Economic Review*, 96(1): 448-451.

Svensson, Lars E.O. 2006b. "The Instrument-Rate Projection under Inflation Targeting: The Norwegian Example." In *Stability and Economic Growth: The Role of Central Banks.* Mexico: Banco de Mexico, 175-98.

Thornton, Daniel. L., and David. C. Wheelock. 2000. "A History of the Asymmetric Policy Directive" *Federal Reserve Bank of St. Louis Review*, 82(5): 1–16.

Vega, Marco and Diego Winkelried. 2005. "Inflation Targeting and Inflation Behavior: A Successful Story?" *International Journal of Central Banking*, 1(3): 153-175.

Willard, Luke. 2006. "The Effect of Inflation Targeting on Inflation: A Reassessment," Ph.D. dissertation essay, Princeton University.

Woodford, Michael. 2001. "Monetary Policy in the Information Economy." In *Economic Policy for the Information Economy.* Kansas City: Federal Reserve Bank of Kansas City, 297-370.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

Woodford, Michael. 2005. "Central-Bank Communication and Policy Effectiveness." In *The Greenspan Era: Lessons for the Future.* Kansas City: Federal Reserve Bank of Kansas City, 399-474.