

Understanding Inflation and the Implications for Monetary Policy




A Phillips Curve Retrospective

EDITED BY

Jeff Fuhrer,
Yolanda K. Kodrzycki,
Jane Sneddon Little,
and Giovanni P. Olivei

FOREWORD BY

Paul A. Samuelson



Understanding Inflation and the Implications for Monetary Policy

Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective

edited by

**Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little,
and Giovanni P. Olivei**

foreword by

Paul A. Samuelson

The MIT Press
Cambridge, MA
London, England

©2009 Massachusetts Institute of Technology. All rights reserved.

No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Sabon by Sztrecska Publishing and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Understanding inflation and the implications for monetary policy : a Phillips curve retrospective / edited by Jeff Fuhrer ... [et al.] ; foreword by Paul A. Samuelson.

p. cm.

Papers presented at a meeting held in June 2008 and sponsored by the Federal Reserve Bank of Boston.

Includes bibliographical references and indexes.

ISBN 978-0-262-01363-5 (hardcover : alk. paper)

1. Inflation (Finance)—Congresses. 2. Phillips curve—Congresses. 3. Monetary policy—Congresses. 4. Unemployment—Effect of inflation on—Mathematical models—Congresses. I. Fuhrer, Jeffrey C.

HG229.U43 2009

332.4'1—dc22

2009010490

Contents

Acknowledgments	<i>ix</i>
Foreword	<i>xi</i>
Paul A. Samuelson	
1 The Phillips Curve in Historical Context	1
Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei	
2 Fifty Years of the Phillips Curve: A Dialog on What We Have Learned	69
Robert M. Solow and John B. Taylor N. Gregory Mankiw, Moderator	
3 Forecasting Inflation	99
Phillips Curve Inflation Forecasts	101
James H. Stock and Mark W. Watson	
Comments	187
Adrian Rodney Pagan	
Comments	195
Lucrezia Reichlin	
4 The Labor Market and the Phillips Curve	205
A New Method for Estimating Time Variation in the NAIRU	207
William T. Dickens	
Comments	231
Olivier Blanchard	
Comments	237
Christopher A. Pissarides	

5 Inflation Expectations	245
Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy	247
Christopher A. Sims	
Comments	267
Michael T. Kiley	
Comments	281
Athanasios Orphanides	
6 Evidence on Price Determination	289
Implications of Microeconomic Price Data for Macroeconomic Models	291
Bartosz Maćkowiak and Frank Smets	
Comments	333
Robert G. King	
Comments	351
Virgiliu Midrigan	
7 Is the Phillips Curve Vertical in the Long Run?	359
Hysteresis in Unemployment	361
Laurence Ball	
Comments	383
V.V. Chari	
Comments	395
Jordi Galí	
8 Lessons for Central Bankers: A Panel Discussion	401
Israeli Monetary Policy and the Phillips Curve	403
Stanley Fischer	
Lessons for Central Bankers from a Phillips Curve Framework	415
Donald L. Kohn	
The Phillips Curve and the European Central Bank	427
Jürgen Stark	
Monetary Policy at the Riksbank and the Phillips Curve	433
Lars E.O. Svensson	

9 The Phillips Curve Going Forward:	445
What We Still Need to Learn	
Outstanding Issues in the Analysis of Inflation	447
Ben S. Bernanke	
Contributors	457
Author Index	473
Subject Index	481

Acknowledgments

The Research Department of the Federal Reserve Bank of Boston and the four editors are grateful to the many people who contributed to this book, which serves as a permanent record of our June 2008 conference. In particular, we thank Patricia Geagan for superb management of the conference arrangements. She was ably assisted by Selva Bahar Baziki, Tom DeCoff, Donna Dulski, Heidi Furse, Nancy Gillespie, Fabienne Madsen, Ana Patricia Muñoz, Ralph Ragsdale, Teresa Foy Romano, and Sam Shuman. As managing editor, Elizabeth Murry worked creatively with the authors to shape the preliminary conference presentations into final copy; her contributions were key to the book. In this endeavor, she was assisted by Denny Lie and Puja Singhal. Sally and Steve Sztrecska typeset the manuscript with their usual good cheer. Krista Magnuson provided expert proofreading, and Tobiah Waldron once again assembled an impressive index in record time.

At the MIT Press, we are especially appreciative of the efforts and kind indulgence of Deborah Cantor-Adams, Terry Lamoureux, and Jane Macdonald. They are wonderful collaborators.

Foreword

Paul A. Samuelson

When members of the research department at the Federal Reserve Bank of Boston first thought to organize a June 2008 meeting of top economists and money managers in the quiet of Cape Cod, they may not have realized how timely and hot such a gathering would prove to be. By August 2008, the U.S. and the global economy seemed to be in a perfect storm.

OPEC oil and energy prices then constituted a serious unfavorable supply-side shock. Reinforcing that was a near-bubble in the prices of grains and metals. Naturally, there soon followed chronic weakness and chaos in Wall Street's pricing of stocks and bonds.

Leading the parade of bad economic news was a global burst bubble in residential and commercial real estate. For the first time in economic history, such a drop in home prices took place when new fiendish instruments of Frankenstein financial engineering were fanning the flames of looney subprime mortgage lending and borrowing.

The timing could not have been better for a virtual Woodstock gathering of frontier economists. This edited volume based on the Boston Fed's 2008 conference provides a rare opportunity to get a clear peek at how the practice and theory of macroeconomics itself evolves. The invited speakers are on everybody's short list of macroeconomic specialists.

Independent central banks have converged in the last quarter century on a common synthesis. Broadly speaking, in most modern charters, central bankers need to be *two-eyed* monitors: their central banks are, at the same time, (1) to target against too high or too low growth rates in the overall indexes of inflation; and also (2) to seek to target against too high unemployment and excess capacity rates. (Some central banks like the European Central Bank and the Bank of England are exceptional in that

their charters mention only an inflation target—presumably in the hope that this *necessary* condition will also be a *sufficient* condition to achieve best real output growth.)

Successive Federal Reserve Bank chairmen—Paul Volcker, Alan Greenspan, and Ben Bernanke—along with central bankers on half a dozen continents, have favored using only *one* macroeconomic control variable to monitor *two* targets. If this practice raised a pure mathematician’s eyebrow at Chatham, Stanford’s John Taylor was there to explicate the Taylor rule, which seeks to minimize some weighted average of (a) unemployment gaps and (b) unsustainable growth rates of price levels. That rule seems to have given a fair description of actual central bank reactions at the turn of the new century.

Between 1980 and 2006–2007 this new synthesis seemed to work out pretty well. Bestselling textbooks began to hail the “Great Moderation” as a description for the 1980–2005 business cycles: recovery periods were longer than recessionary periods, while the amplitudes of boom and bust cycles seemed to have considerably dampened down.

Federal Reserve Chairman Bernanke inherited the task of publicly *quantifying* transparent target goals: maybe inflation in the cost of living should be kept within, say, a 1 percent to a 2 or 3 percent annual rate? And, depending upon a nation’s trade unions and other labor market institutions, unemployment rates of 3 percent to 6 percent might be hoped for achievements.

This central bank synthesis faces problems everywhere *whenever supply-side shocks generate severe degrees of stagflation*.

During the 1970s’ stagflation period, oil and harvest shortages were present. But what had peculiarly accelerated inflation in the 1970s was President Nixon and Fed Chairman Arthur Burns’s excessive creation of credit for the Machiavellian purpose of ensuring Nixon’s 1972 reelection. That on a small scale was a “banana republic” creation of too much fiat money. By contrast, much of 2008’s excessive inflation may have traced considerably to how much total factor productivity was being pushed down by supply-side shocks from abroad.

To quench such “micro-caused” inflation, today’s central bankers must assay the Herculean task of reducing nominal general wage rates—always and everywhere a most difficult political caper.

No wonder the 2008 get-together was planned to center around the Phillips curve trade-off between high unemployment and high inflation. Earlier during stagflationary times like the 1970's, raising the central bank's official interest rate might well moderate the "flation" in stagflation, but at the same time it would worsen the "stag."

This published record of the Chatham meetings will stand out as a historic gem. Expert economists expounded their wisdom there: Stanford's John Taylor, MIT's Robert Solow, Harvard's Greg Mankiw, along with many other economists from the Bank of Sweden and the European Central Bank, among other central banks. If Alan Greenspan was not present in person, his nuanced viewpoints were well represented. Past doughty debates on the Phillips curve by Milton Friedman and James Tobin were definitely there as living memories.

Some disagreements took place, of course. But the record of this meeting does show an admirable degree of civility among "rational expectationists" of the Robert Lucas and Thomas Sargent school, inflation hawks, and optimistic activists.

Live snapshots of science's history in the making are rare. To be able to read the present book is the next best thing to having been at this historic economics Woodstock.

1

The Phillips Curve in Historical Context

The Phillips Curve in Historical Context

Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little,
and Giovanni P. Olivei

In the spring of 2008, U.S. policymakers confronted a rather unappealing confluence of macroeconomic factors: falling employment and tepid final sales for the two quarters spanning the turn of the year suggested a weak real economy, while stupendous surges in oil and food prices pushed inflation above 5 percent, with so-called core inflation measures (excluding food and energy prices) rising above 3 percent. The faltering financial sector—and recall that the full extent of the financial meltdown was not anticipated at that point—added downside risk to the real economy. The federal personal income tax rebate provided a glimmer of hope, but the size and timing of the response to the tax rebate checks that were deposited beginning in May 2008 were quite uncertain. Meanwhile, oil prices remained stubbornly high, breaching \$130 per barrel in late May and heading further upward, thus raising the risk that inflation would not recede from its elevated level any time soon. Stagflation seemed a clear and present danger.

The economic environment changed dramatically in September 2008, as a number of systemically important financial institutions failed or came very close to it, equity prices declined dramatically, data on the real economy weakened sharply, and the price of oil dropped to less than one-half of its early July peak. Economic forecasters converged on recession, with many forecasts expecting unemployment to peak above 7 percent.¹ Concerns over elevated inflation rates retreated rapidly, and were soon replaced with concerns that inflation would fall below the Federal Reserve's (unofficial) "comfort zone," or even more disconcerting, below zero percent.

Both before and after the September 2008 watershed, economists would have liked to have had a clearer understanding of the determinants of inflation. If the economy remained weak and the degree of resource slack rose, how much disinflationary pressure would be exerted, if any? In other words, to what extent would a Phillips curve-type mechanism come into play? How would the rapid rise and subsequent decline in the relative prices of food and energy feed through to the general price level? Could one see signs of relative price pass-through in inflation expectations or wage-setting? Would the Federal Reserve erode its credibility if it wound up presiding over a period during which the annual inflation rate remained persistently above or below the presumed comfort zone of 1–2 percent? In turn, how might that breach affect inflation expectations in the medium-to-long run?

As economists and policymakers gathered on Cape Cod in June 2008, the first set of circumstances—the threat of stagflation—formed the immediate economic backdrop. One might have felt more confident about the answers to these questions if inflation modeling procedures were reasonably agreed upon and settled, and inflation was easy to forecast. But that would not be an accurate depiction of the current state of affairs in macro and monetary economics.

The ongoing need to provide better answers to these questions prompted the Federal Reserve Bank of Boston to organize the conference, “Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective.” Given the central role of the Phillips curve in many economic forecasters’ analytical arsenal, the fiftieth anniversary of the famous article that introduced this remarkable yet controversial relationship provided a strong motivation for examining some enduring macro and monetary policy questions. These issues and conundrums have taken on greater resonance in the ensuing year, as the United States and the rest of the world grapples with what is now the worst global financial crisis and economic downturn since the Great Depression. As background, this chapter’s first main section provides an intellectual history of the Phillips curve, while the second main section offers a summary of the revised conference papers and comments, placing this material within the history of thought regarding the Phillips curve paradigm.

1. An Intellectual History of the Phillips Curve: Theory and Empirics

A.W. Phillips's Basic Correlation (1958)

New Zealand-born economist Alban W. Phillips's seminal 1958 paper, "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957," posited and documented a negative correlation between the change in money wage rates and unemployment. While it is not widely recognized, Phillips's paper also discussed a number of other wage determinants that have since received considerable attention in the literature on wage and price determination. For example, he suggested the possibility of what is now called a "speed limit" effect, whereby not only the level but also the *change* in the rate of unemployment affect the change in nominal wages. Phillips also suggested that a cost-of-living effect, proxied by changes in retail prices, might affect the rate of change of money wages, although this effect was not generally present in his data, except when retail prices rose rapidly due to the effects of imported goods or domestic agricultural prices. This cost-of-living effect could mask the underlying negative correlation, and Phillips took some care to identify years in which the rate of increase of import prices was large enough to obscure the wage-unemployment correlation. Phillips also anticipated a reluctance on the part of workers to accept nominal wage cuts when unemployment is high, suggesting a relationship that is "highly non-linear."²

Phillips estimated a log-log relationship between nominal wage changes and unemployment from 1861 to 1913 as

$$(1) \log(\Delta w_t + 0.9) = 0.984 - 1.394 \log(U_t),$$

and examined, via scatter plots, the wage-unemployment correlation for subperiods, pointing to times when the change in import prices was sufficient to push nominal wages off the estimated curve, and emphasizing the clear presence of "speed limit" effects in several years. Note that all of Phillips's analysis involves money wages or nominal wages, not because Phillips believed that unemployment is related to nominal rather than real wages, but because he lived in a world in which it was reasonable to assume that prices would remain relatively stable, temporary disrup-

tions from import prices notwithstanding. This omission, innocuous for the first part of the twentieth century but grossly counterfactual for the century's second half, was taken up by Friedman and Phelps, and is discussed below.

Phillips then superimposed the scatter plot of the U.K. data from 1913–1947 and 1948–1957 on the estimated wage-unemployment curve. Again, Phillips provided a detailed accounting of the effects of imported goods prices on the change in wages, and in the latter period, documents a lagged relationship between the two, perhaps establishing the first instance of a dynamic Phillips relationship.

The results from this analysis are at once familiar and alien to modern practitioners. One is not surprised that the data show frequent deviations of wage changes from the estimated Phillips curve, due in large part to outsized surges in the prices of imported goods. On the less familiar side, one might be hard-pressed to reject the stability of the estimated curve—using data from 1861 to 1913, and depicted in figure 1.1—based on the “out of sample” scatter plots for the ensuing 45 years. Since that era we have not seen such an extended period of stability in the underlying correlation Phillips found between inflation and unemployment.

A Key Theoretical Insight: Friedman (1968) and Phelps (1968)

The presence of a reasonably reliable correlation between unemployment and nominal wage (or price) inflation might imply a trade-off between the two that policymakers could exploit by choosing pairs of inflation/unemployment outcomes that they deem socially desirable. For example, a desire to maintain very low unemployment might be achieved by accepting a moderately high but stable rate of inflation. The extent to which monetary policy can exploit the trade-off between inflation and unemployment has dominated the aggregate supply literature at least since the 1960s.

The policy implications of an exploitable Phillips correlation were widely discussed in U.S. policy circles in the 1960s (for examples, see the accounts in Akerlof, Dickens, and Perry 2000; Primiceri 2006; and Sargent, Williams, and Zha 2005). It was Samuelson and Solow (1960) who first noted an empirical trade-off between wage inflation and unemployment for the United States (though the relationship was not as tight as in




Figure 1.1
Rate of Change of Wage Rates and Percentage Unemployment in the United Kingdom, 1861–1913

Source: Redrawn from Phillips (1958), figure 1.

Note: The light grey dots give an approximation to the rate of change of wages associated with the indicated level of unemployment if unemployment were held constant at that level (see Phillips, 1958, 290).

Phillips’s data from the United Kingdom). They then discussed the policy implications of this trade-off between inflation and unemployment, and speculated that such a trade-off could be exploited, if at all, only in the short run. They pointed out that in the long run several factors could lead to shifts in the Phillips curve that would greatly complicate any policy effort aimed at choosing a specific point along the short-run Phillips curve.

Milton Friedman, in his December 1967 presidential address to the American Economic Association, was especially influential in stating the most serious flaw in arguments for an exploitable inflation-unemployment trade-off: surely labor markets would operate so that nominal wages *relative to price inflation* were relatively high when excess demand for labor was large, and vice versa.³ Friedman traced out the mechanisms

by which a monetary policy that aims to lower unemployment via a monetary expansion can only achieve that goal temporarily, as lower interest rates stimulate spending, raise the marginal product of labor, and increase employment and output. In Friedman's view, prices will rise before wages, lowering the real wage received, thereby prompting increased nominal wage demands by labor. Ultimately, wage increases will match accumulated price increases, and the rising real wage rate will bring unemployment back to its "natural" rate. As Friedman puts it, "there is always a temporary trade-off between inflation and unemployment; there is no permanent trade-off."⁴ From this point forward, monetary policy's ability or inability to influence inflation re-emerged as a central theme in macroeconomics, as it was in earlier debates. But any policy outcome is intimately tied to the precise form taken by the Phillips curve.

Edmund Phelps took a related tack, drawing on his earlier work (Phelps 1967), and posited that "the Phillips curve ... shifts uniformly upward by one point with every one point increase of the expected percentage price increase" (Phelps 1968, p. 682). A consequence is that the long-run or equilibrium unemployment rate is independent of the rate of inflation. In his early papers, Phelps employed an adaptive expectations framework, which implies that the unemployment rate U is linked to the *change* in the rate of inflation π :

$$(2) \quad \pi_t = \pi_t^e - aU_t = \pi_{t-1} - aU_t$$

$$\Delta\pi_t \equiv \pi_t - \pi_{t-1} = -aU_t.$$

This so-called accelerationist Phillips curve—in which the acceleration or second time-derivative of prices is related to unemployment—embodied two critical innovations in the literature. First, it eliminated the long-run trade-off between inflation and unemployment that was inherent in the original Phillips curve model. Second, it began to emphasize the importance of expectations in the price-setting process, a change that was to have dramatic implications for the evolution of inflation models for the next four decades.

While Friedman and Phelps consider the long-run or "natural" rate of unemployment, meaning the rate to which unemployment returns in equilibrium independent of the level of inflation, research on the Phillips

curve has focused on the concept of the non-accelerating inflation rate of unemployment (NAIRU). One can slightly alter equation (2) above to highlight the role of the NAIRU in the Phillips curve:

$$\Delta\pi_t \equiv \pi_t - \pi_{t-1} = -a(U_t - U^N).$$

This formulation makes it clear that when the unemployment rate equals the NAIRU (which is implicitly zero in equation 2), here denoted by U^N , the change in the inflation rate is zero. More generally, when the unemployment rate equals the NAIRU, inflation equals expected inflation, which in Phelps's paper is proxied by lagged inflation.

Introducing Rational Expectations: Lucas (1973) and Sargent-Wallace (1975)

With explicit expectations beginning to play a more central role in models of price determination, the earlier introduction of Muth's (1961) rational expectations principle into the macroeconomics literature was taken up following the Friedman-Phelps critique of Phillips's original framework. The policy implications of the Phillips curve trade-off took on greater policy urgency as U.S. inflation accelerated in the late 1960s and early 1970s. Muth made the simple but profound observation that in economic models, expectations were often proxied by *ad hoc* mechanisms that were inconsistent with the equations that researchers wrote down to determine the evolution of the key variables. A more internally consistent method is to assume that expectations are formed in a way that is derived from the model that the researcher posits. To take a simple example, imagine that prices p_t depend on the previous period's expectation of prices in period t , p_t^e , plus an adjustment for current excess demand conditions D_t :

$$p_t = ap_t^e + bD_t.$$

One could assume that expected prices, p_t^e , are formed adaptively, which loosely speaking makes them a function of past prices,

$$p_t^e = cp_{t-1},$$

so that the resulting price equation becomes

$$p_t = acp_{t-1} + bD_t.$$

Alternatively, one can assume that the original equation for prices determines how expectations will be formed. In this case, the expected prices in period t are a function of expected prices in period t using information up to period $t - 1$ and expected excess demand in period t . Expected prices are given by

$$p_t^e = ap_t^e + bD_t^e = \frac{b}{1-a} D_t^e,$$

and thus the evolution of prices is determined by the equation,

$$p_t = \frac{ab}{1-a} D_t^e + \frac{b(1-a)}{1-a} D_t = \frac{ab}{1-a} (D_t^e - D_t) + \frac{b}{1-a} D_t.$$

The rational expectations assumption in this case bears important implications for the evolution of prices. Under adaptive expectations, prices depend explicitly on past prices, imparting some inertia to the subsequent evolution of prices. Under rational expectations, prices move proportionately and immediately in response to excess demand and excess demand surprises.⁵

Lucas (1973) employed the rational expectations assumption in an imperfect information model of aggregate supply, in which price misperceptions cause output to deviate from full-employment output.⁶ Producers are unable to perfectly disentangle the extent to which a movement in the price they observe for their product is a relative price change, which should elicit a production response, versus an aggregate price change, induced by an increase in the money supply, which should not elicit a production response. The slope of the output/price relationship depends on the ratio of variances in firm-specific price shocks versus aggregate price shocks: in the limit, as all relative price shocks become aggregate price shocks, the slope of the supply relation becomes vertical. The implication is that even in the short run, monetary policy can influence output only by causing unanticipated movements in the price level. Thus even the short-run trade-off between inflation and unemployment outlined by Friedman and Phelps is ephemeral in this class of models incorporating rational expectations.

Sargent and Wallace (1975) derive a very similar result. With a simple *ad hoc* macroeconomic model comprising a Lucas-style Phillips curve, an IS curve, an LM curve, and an equation describing productive capac-

ity, they find that output is related to unexpected movements in prices. A corollary is that if errors in anticipating prices are not serially correlated, both output and the price level will not exhibit any serial correlation. That is, the paper implied a very flexible price level. Foreshadowing a vigorous discussion in decades to come, Sargent and Wallace deride their own model as quite *ad hoc*, in their own words describing it as “not derived from a consistent set of assumptions about individuals’ and firms’ objective functions and the information available to them,” a feature that they consider “deplorable” (p. 241).

Essentially, the short-run trade-off that Phelps and Friedman posited arose as long as the monetary authority could create unanticipated growth in the money supply. Under adaptive expectations, anticipation errors could persist for some time. Under rational expectations, as long as wage-setters know the money growth rule, such forecasting errors are unlikely to persist, and thus the influence of monetary policy on employment and output is limited, and the price level is flexible.

Yet in responding to Sargent and Wallace (1975) as well as to the oil shocks and the positive correlation of inflation and unemployment in the 1970s, Robert Gordon (1977) points out that the argument contending that monetary policy cannot even briefly influence unemployment unless such policy is unpredictable requires that the price level respond instantaneously to any change in the market-clearing price. But this argument flies in the face of strong empirical evidence that U.S. prices adjust only sluggishly.⁷ Building on this critique, in the late 1970s and early 1980s Gordon led the development of a “resolutely Keynesian” coherent dynamic aggregate supply and demand framework that came to be known as the “triangle” model⁸ (see Gordon 1982, Gordon and King 1982, Gordon 2008). This framework incorporates as basic tenets the long-run neutrality of monetary policy and an explicit role for supply shocks. Gordon’s triangle model interprets past inflation as reflecting not just the formation of inflation expectations, but also a generalized inertia stemming from implicit and explicit wage-price contracts and lengthy supply chains. This mainstream backward-looking specification, which Gordon points out can be consistent with rational expectations, enjoyed some empirical success and became a workhorse model widely used for forecasting purposes, particularly at central banks.

Rational Expectations with Price Inertia: Fischer (1977) and Gray (1977)

Work by Stanley Fischer and Jo Anna Gray laid the groundwork for a long tradition in macroeconomics that, by positing a variety of wage and price rigidities, finds a role for monetary policy in rational expectations models. These earliest papers in the genre assume *ad hoc* one- or two-period nominal wage contracts. Making wages or prices predetermined for some time allows anticipated monetary policy to have an effect on employment and output, even under rational expectations. It also imbues wages and prices with some persistence (in general it implies $n - 1$ period serial correlation, where n is the number of periods for which wages or prices are held fixed).

The intuition behind such price rigidity is straightforward. Say the nominal wage rate is held fixed for two periods, and the monetary authority is free to change the money supply in response to information received after the wage is set. Then monetary policy can affect the price level and thus the real wage before the nominal wage is able to adjust to these actions. Because output will generally be a (negative) function of real wages, monetary policy is now able to affect real output during the period that the nominal wage is fixed.⁹ However, as suggested above, the duration of the monetary policy effect on output is limited to the length of the longest wage contract. The observed duration of the employment and output effects in business cycles suggested that this represented an empirically significant limitation of the models.¹⁰ Nevertheless, at this point in the late 1970s, the literature focused on developing theoretical frameworks with rational expectations in which anticipated policy had or did not have lasting effects on output. The empirical validation of these models was scant.

Staggered Contracts with Multi-Period Rigidity: Taylor (1980)

Taylor's seminal paper, "Aggregate Dynamics and Staggered Contracts," broke the strict correspondence between the length of the longest wage contract and the duration of monetary policy effects on output and employment through two innovations. First, in Taylor's model contracts were "staggered," meaning these were not all renegotiated at the same time. Second, and just as importantly, the contracts were made with ref-

erence to the contracts that had been set previously, to the extent that those contracts would remain in effect for part of the life of the contract currently being negotiated. The weights that past and future contracts receive in influencing the current wage contract depend on how the previous and future contracts overlap with the current contract.

Figure 1.2 displays the distribution of these contract weights for three different contracting models. The weights sum to one in all cases, and in general reflect the diminishing importance to today's contract of contracts set further in the past and those contracts expected to be set further in the future. The top panel shows the Fischer model, in which contracts last for two periods and can overlap. Yet the contracts are set without reference to the wage rates embedded in other contracts still in effect or expected to be in effect. In contrast, the Taylor model, shown in the middle panel of figure 1.2, shows that neighboring contracts from the past and in the (expected) future influence the setting of today's contract wage. Because contract lengths are all the same in Taylor's framework, the pattern of contract weights takes a symmetric triangular shape.

Because Taylor's contracts are set relative to these overlapping contracts, the effects of a change in the money supply today affects not only today's contract, but contracts for the next several periods, which will be set partly in reference to today's contract. Those future contracts in turn will serve as reference points for contracts set even further in the future. In this manner, monetary policy will have very long-lasting (in principle, infinite) effects on future real wages and thus on output. In addition, note that in the few periods immediately following the shock, the effects of monetary policy *rise*, as the shock affects (in the case of four-period contracts) first the current, then the current and next period's, then three and maximally four sets of overlapping contracts. This hump-shaped response to monetary policy conforms with the perceived effects of monetary policy on the economy, as illustrated in the middle panel of figure 1.2.

New Formulations, with Partial Micro Foundations: Calvo (1983) and Rotemberg (1982, 1983)

By the early 1980s, the appeal of including wage and price rigidities in macroeconomic models was evident, but as Sargent and Wallace (1975)




Figure 1.2
The Influence of Neighboring Wage Contracts on the Current Contract Wage
Source: Authors' calculations.

noted, the microeconomic foundations for such behavior were less clear. Why, in the face of changing economic conditions, would firms hold wages or prices fixed in nominal terms? The search for plausible microeconomic foundations for such models began in earnest.

Three papers that arrived on the scene around the same time provided partial answers to this question regarding the existence of rigid (or “sticky”) wages and prices. The two authors, Calvo (1983) and Rotemberg (1982, 1983), offered different explanatory rationales, but when stripped to their respective cores, their models are nearly identical and bear essentially the same implications for macroeconomic dynamics.¹¹

Calvo’s paper, as suggested by its title “Staggered Prices in a Utility-Maximizing Framework,” provides partial microeconomic foundations for aggregate movements. The model assumes that firms may change prices only upon receipt of a price-change signal, an event that future authors whimsically described as being “tapped by the Calvo fairy.” The exogenous probability of receiving such a signal was modeled as drawn from a geometric distribution, chosen by Calvo for its analytic tractability and expressed as:

Probability (receiving a signal, h periods hence) = $\delta e^{-\delta h}$,

which implies that the mean duration of a price contract is $\frac{1}{\delta}$. Equivalently, this probability implies a geometric distribution of price contract lengths, with shorter durations most likely and longer durations increasingly unlikely. As in Taylor (1980), firms set their contract price (when tapped by Calvo’s price-change signal) in reference to overlapping contracts, and the result is a price level that depends on a geometric weighted average of the infinite past and future contract prices.¹² The effect of overlapping contracts in the Calvo model is displayed in the bottom panel of figure 1.2. Because it implies qualitatively similar features for price (or wage) contracts, Calvo’s model bears the same implications for the effectiveness of anticipated monetary policy actions as in Taylor.

In Calvo’s paper, utility maximization arises in specifying consumer demand for various firms’s products, which carry different prices. Rotemberg (1982, 1983) arguably makes greater advances in providing optimizing foundations for price-setting *per se*. He assumes that when adjusting prices, individual firms face quadratic costs, both relative to

previous prices and relative to the price that would obtain in the absence of adjustment costs, p^* . One can express the firm's optimization problem as, per Roberts (1995, p. 976),

$$(3) \quad \min_p E_t \sum_{s=t}^{\infty} \beta^{s-t} [(p_s - p_s^*)^2 + c(p_s - p_{s-1})^2].$$

The first-order conditions for this optimization problem can be simplified to obtain the now-canonical form of the New Keynesian Phillips curve, a form that can also be derived from the Calvo model above:

$$(4) \quad \pi_t = \beta E_t \pi_{t+1} + c x_t + \varepsilon_t,$$

where x_t in the Rotemberg model represents the deviation of the firm's actual price from the firm's optimal (absent adjustment costs) price, whereas in the Calvo model this deviation stands for excess demand.

The virtue of these models is that they incorporate rational expectations, provide some underlying microeconomic foundations for pricing decisions, and allow for a nontrivial role for anticipated monetary policy. As discussed below, in most incarnations the Calvo/Rotemberg models impose strongly counterfactual implications, but these implications are best revealed in a richer macroeconomic environment that articulates the behavior of the central bank and the private spending decisions of agents. To anticipate these later developments, note that one can “iterate forward” the canonical New Keynesian Phillips curve—that is, use the definition of inflation at period $t + 1$ to substitute for the value of inflation that appears on the right-hand-side of the equation, and so on—to obtain a solution for the inflation rate in terms of expected future output or excess demand; thus

$$(5) \quad \pi_t = c E_t \sum_{i=0}^{\infty} \beta^i x_{t+i} + \varepsilon_t.$$

This rendering of the Calvo/Rotemberg models implies that inflation is a purely forward-looking variable. As a consequence, it can move frictionlessly in response to shocks to the driving variable x . In addition, it will be serially correlated only to the extent that x is serially correlated. These features, which bear important and testable implications, will be addressed in more detail below.

Deeper Micro-Foundations and Empirical Testing

More recently, several authors have formulated an aggregate supply equation for the economy that is derived from the firm's optimization problem.¹³ In this specification, each firm faces a Calvo-style restriction on its ability to reset prices in a monopolistically competitive setting, where each firm supplies a differentiated good. The resulting aggregate supply equation yields an intuitively appealing version of the New Keynesian Phillips curve; the firms that can reset their prices (those that have been tapped by the "Calvo fairy") set their price so as to maximize profits over the price's expected duration, and thus set their price to the expected average marginal cost of production over that period.¹⁴ This implies that the rate of inflation will be a function of expected real marginal cost, x ,

$$(6) \quad \pi_t = \beta E_t \pi_{t+1} + \kappa x_t \Rightarrow \kappa E_t \sum_{i=0}^{\infty} \beta^i x_{t+i}.$$

Galí and Gertler (1999) and Sbordone (2002) examine this version of the inflation specification, taking the real average unit labor cost as a proxy for the real marginal cost. This assumption is equivalent to using labor's share of income,

$$(7) \quad x_t \approx \left(\frac{w_t}{y_t / L_t} \right) / p_t = \frac{w_t L_t}{p_t y_t},$$

where w is the nominal average wage rate, y is nominal output, L is the labor input, and p is the price level. This proxy has become the most common determinant in empirical inflation specifications. These authors find considerable empirical support in favor of this version of the model, estimating a significant and positive value for κ .¹⁵

Galí and Gertler also consider augmenting this New Keynesian Phillips curve with a backward-looking element that is motivated by the presence of some firms who follow a simple rule of thumb in setting prices. Christiano, Eichenbaum, and Evans (2005) derive a similar specification under the assumption that price-setters who are unable to reset prices instead index their prices to last period's inflation rate. The Fuhrer-Moore model (1995) employs a relative price-contracting specification to derive a simi-

lar two-sided hybrid Phillips curve. All of these variants imply a so-called hybrid New Keynesian Phillips curve of the form

$$(8) \quad \pi_t = \gamma^f E_t \pi_{t+1} + \gamma^b \pi_{t-1} + \kappa x_t,$$

which gives rise to an interesting question regarding the relative magnitude of γ^f and γ^b . Galí and Gertler's estimates span a range from near zero to a bit over 0.5, but the modal estimate γ^b is about 0.25. This set of results implies a statistically significant but economically limited role for the rule-of-thumb price setters, and perhaps more importantly, a limited need for including lags of inflation in the inflation specification.

That finding stands at odds with the empirical results found by other researchers. The jury is still out on the empirical success of the purely forward-looking New Keynesian Phillips curve, and the ongoing debate remains lively. As Rudd and Whelan (2006) put it, "the observation that lagged inflation plays an important role in empirical inflation regressions poses a major challenge to the rational expectations sticky-price models that underpin the new-Keynesian Phillips curve" (p. 318). For most of the past 45 years, the inflation rate in the United States has been a very persistent series, characterized by a sum of autoregressive coefficients of 0.7 to 0.9.¹⁶ Any model that wishes to explain the behavior of U.S. inflation in the last half-century must grapple with this first-order empirical fact about inflation. So a key question is then established: where does the persistence in inflation come from?

The crux of the issue can be seen by inspecting equations (4) and (5), which define the canonical New Keynesian Phillips curve. Inflation will generally inherit the autocorrelation properties of output (or marginal costs). Both of these series exhibit high degrees of autocorrelation. Thus, the question becomes whether inflation adds its own intrinsic persistence to that of the output process.¹⁷ If the persistence in the output process is sufficient to explain the persistence in inflation, then the coefficients on the lagged inflation term in the hybrid New Keynesian Phillips curve of equation (8) should be zero.¹⁸

Why is the size of this lag coefficient—and thus the degree of intrinsic persistence—so important? The more intrinsic persistence inflation embodies, the more difficult it will be for monetary policy to move inflation around. If inflation itself is inertial, then a given monetary policy

action that changes output will have a smaller effect on inflation. Thus it is important for the central bank to know how much of the observed persistence of inflation is an artifact of the persistence of output, and how much of this persistent inflation is *sui generis*.

In a series of papers, Fuhrer and Moore (1995), Fuhrer (1997, 2006), and Rudd and Whelan (2006, 2007) provide evidence bearing on this question. Their combined analyses suggest that the purely forward-looking New Keynesian Phillips curve as described by equation (6) performs quite poorly. They employ a variety of tests, all of which come to the same conclusion: inflation appears to embody a sizable amount of intrinsic persistence; that is, persistence beyond what is inherited from the output gap or the real marginal cost.

A key insight into disentangling intrinsic and inherited inflation persistence lies in the shock term, ε_t , which appears on the right-side of equations (4) and (5). Without this shock, inflation would be identically equal to the discounted sum of future output gaps or marginal cost, and thus its behavior would be entirely determined by the behavior of the driving variable. But in the presence of shocks to the New Keynesian Phillips curve, inflation can either respond inertially—implying a lagged inflation term in the New Keynesian Phillips curve—or it can respond immediately. Thus a key to identifying the absence or presence of a lagged term is the importance of the shock term in the New Keynesian Phillips curve. Empirically, the New Keynesian Phillips curve appears to be buffeted by shocks of significant magnitude. This result could reflect a serious mismeasurement of the gap or the marginal cost measure, or it could reflect the importance of supply shocks in the determination of prices.¹⁹ The proper characterization of the inflation process in this third-generation descendant of the original Phillips curve remains an open research question.

Integrating the Phillips Curve into Newer Macroeconomic Models

So far we have discussed empirical work that estimates the Phillips curve as a stand-alone equation. However, now there is a large and growing literature encompassing the Phillips curve estimation within a general equilibrium representation of the macroeconomy. Rotemberg and Woodford (1997) provide one of the earliest examples of a truly micro-founded

optimizing model that jointly determines prices, output, and interest rates. There is a very stylized general equilibrium model, in that just three equations describe the economy. The output side of the model is straightforward; it derives from the first-order condition for a utility-maximizing consumer, which equates the intertemporal ratio of marginal utilities of consumption to the product of the discount rate and the real rate of return. This condition can be log-linearized to yield the “optimizing IS” equation,²⁰

$$(9) \quad \tilde{y}_t = \beta E_t \tilde{y}_{t+1} - \sigma(R_t - E_t \pi_{t+1}) + \varepsilon_{y,t},$$

where \tilde{y} is the output gap and R the short-term interest rate. A purely forward-looking Phillips curve of the same form as in (4), with \tilde{y} replacing x , describes the dynamics of prices, and thus the aggregate supply side of the economy. The model is then closed by a feedback rule à la Taylor (1993) for interest rates,

$$(10) \quad R_t = R^* + \sum_{i=0}^j \alpha_{\pi,i} (\pi_{t-i} - \pi^*) + \sum_{i=0}^k \alpha_{y,i} \tilde{y}_{t-i} + \sum_{i=1}^l \alpha_{R,i} (R_{t-i} - R^*) + \varepsilon_{R,t},$$

where R^* and π^* are long-run target values for the short-term interest rate and for inflation, respectively.

Rotemberg and Woodford’s model is characterized by purely forward-looking aggregate demand and supply relationships. Yet the model achieves some empirical success, in that the model’s impulse responses match the empirical impulse responses from a benchmark three-variable vector autoregression (VAR) in $[\tilde{y}, \pi, R]$ to an identified monetary policy shock reasonably well. The model’s ability to fit other features of the data, however, is achieved in part by allowing the time-series properties of the shock processes in the aggregate demand and aggregate supply equations to take on an arbitrarily complex structure. In other words, the shock processes appear to play an important role in characterizing the dynamics of output and inflation, while the empirical content of the driving processes in the aggregate demand and supply equations appears to be very limited.²¹

Rotemberg and Woodford’s work has spurred the development of more ambitious dynamic stochastic general equilibrium (DSGE) models of the macroeconomy. These models provide a more disaggregated rep-

resentation of the demand side of the economy by explicitly treating consumption and investment as separate variables. The price Phillips curve relationship features real marginal costs as its driving process. As a result, these DSGE models also include an equation characterizing the dynamics of the real wage rate. The nominal wage-setting process follows a Calvo-style setup where workers face a constant probability of re-optimizing their nominal wage every period. Given this setup, when workers have the ability to re-optimize their nominal wage, they will take into account expected changes in future inflation and the evolution of current and future marginal rates of substitution between consumption and leisure. The implied wage Phillips curve is more complicated than its price Phillips curve counterpart, though conceptually very similar. One benefit of explicitly modeling wage dynamics is that it becomes possible to investigate the relative importance of price and wage rigidities. Moreover, unlike Rotemberg and Woodford's model, the newer DSGE models are not purely forward-looking. Instead, these more recent models include a backward-looking component in price and wage inflation through price and wage indexation, and these DSGE models allow for real rigidities in consumption and investment via habit formation in consumption and adjustment costs in investment, respectively.

The most notable examples in this new generation of DSGE models are Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). The underlying DSGE model in the two papers is very similar, but the estimation strategy is not. Using the same limited information estimation strategy as in Rotemberg and Woodford, Christiano, Eichenbaum, and Evans estimate some of their model's structural parameters by matching the model's impulse responses to the empirical impulse responses to an identified monetary policy shock in a VAR. Smets and Wouters use a Bayesian likelihood approach. In their setup, shock processes to the model's equations can have an ARMA structure.

The two papers yield somewhat different implications for the Phillips curve. In Christiano, Eichenbaum, and Evans, wage and price indexation is assumed to be complete. For the price Phillips curve, this implies that the specification takes the form of equation (1.8), with $\gamma^f = \gamma^b = 0.5$. In Smets and Wouters, the degrees of wage and price indexation are free parameters that are estimated. In this latter case, the price indexation

parameter is estimated to be very low, so that the estimated price Phillips curve features a limited role for a lagged inflation term. However, this feature of the Smets and Wouters model implies that in order to match the dynamics of inflation, the shock process for inflation plays a very important role and is estimated to be quite persistent.

Both Christiano, Eichenbaum, and Evans and Smets and Wouters stress the importance of generating a driving process for price inflation that is persistent. This persistence is achieved through real rigidities and nominal wage stickiness. These two features contribute to producing a persistent process for the labor share. In particular, Christiano, Eichenbaum, and Evans note that wage stickiness—in addition to wage indexation—plays a crucial role in fitting the model to the data. Yet once price indexation is accounted for, the degree of price stickiness is much less important. Smets and Wouters's different estimation technique yields more nuanced conclusions in this regard. Yet they still estimate that wage indexation is higher than price indexation, a conclusion that again points to the need for articulating a persistent process in which real marginal costs match the inflation dynamics.

Overall, while these more sophisticated DSGE models achieve some empirical success, the empirical relevance of the Phillips curve remains an open question in these models. Smets and Wouters show that the shock process for inflation is a very important determinant of inflation dynamics in the short term, and as a result that the driving process for inflation plays a limited role. Another way of putting this is that the estimated slope of the Phillips curve, in both the work of Christiano, Eichenbaum, and Evans (especially in follow-up work by Altig et al. 2005), and in Smets and Wouters, is very small. Movements in real marginal costs have to be large and persistent in order to play some significant role in explaining inflation dynamics. This limited connection between real economic activity (as measured by real marginal costs) and inflation is problematic in a Phillips curve framework, and is inconsistent with some features of the data. Altig et al. show that while Christiano, Eichenbaum, and Evans's DSGE model matches the empirical response to a monetary policy shock relatively well, it does less well for a productivity shock. In particular, the estimated empirical response of inflation to a productivity shock is not as inertial as it is in the case of a monetary policy shock. The implication

is that DSGE models, estimated with a small slope for the New Keynesian Phillips curve, cannot match the empirically large and relatively fast response of inflation to the productivity shock. Instead, DSGE models produce a response that is too small and too inertial. The observation that the speed with which prices adjust appears to differ according to the type of shock hitting firms features prominently in more recent work on inflation dynamics, as discussed in the next section.

The Phillips Curve and Emerging Micro-Founded Alternative Explanations

So far, we have described the microeconomic foundations of the Phillips curve's new formulations mainly in the context of the Calvo framework. Analytical tractability has made this setup highly popular, and indeed most of the established micro-founded work on inflation dynamics relies on Calvo's framework. But this framework is now coming under increasing scrutiny. As already mentioned, the Calvo model generates a purely forward-looking Phillips curve, and adjustments to the setup to allow for lagged inflation (for example via indexation) are perceived as unsatisfactorily *ad hoc*. In addition, this setup implies that the frequency of price adjustment is independent of the type of shock affecting a firm, an observation that seems at odds with recent empirical evidence.

By addressing some of its shortcomings, current research is now providing alternatives to the Calvo setup. One example is the sticky information model of Mankiw and Reis (2002). The model assumes that acquiring information is costly, and as a result information about macroeconomic conditions diffuses slowly through the population. Specifically, Mankiw and Reis assume that in each period a fraction of firms acquires complete (perfect) information about the current state of the economy, and sets prices optimally based on this information. The remaining firms continue to set prices based on outdated information. Mankiw and Reis's model shares the Calvo feature that the probability of acquiring information about the state of the economy at a certain point in time is exogenous. Their model's implications, however, are different: Mankiw and Reis posit that what matters now for current inflation is not current expectations about future economic conditions, but past expectations

about current economic conditions. The Phillips curve specification in this sticky-information context takes the form

$$(11) \quad \pi_t = \frac{\alpha\tau}{1-\tau} \tilde{y}_t + \tau \sum_{j=0}^{\infty} (1-\tau)^j E_{t-1-j}(\pi_t + \alpha\Delta\tilde{y}_t),$$

where $\Delta\tilde{y}_t = \tilde{y}_t - \tilde{y}_{t-1}$. Inflation depends on the current output gap and on a geometric sum of past expectations of current inflation and output growth relative to potential.

The presence of past expectations of current inflation makes the Phillip curve representation somewhat similar to the Fischer (1977) contracting model. The sticky information Phillips curve specification, unlike the pure forward-looking Calvo-style specification, can generate a delayed inflationary response—that is, inflation inertia—to a monetary policy shock. The qualitative features of inflation’s response to a monetary policy shock under the sticky-information specification (equation 11) match those of a hybrid New Keynesian Phillips curve as given in equation (8), calibrated with a sizable weight on the backward-looking component of inflation. In contrast to the pure forward-looking Calvo-style specification, which allows for the possibility of disinflationary booms, Mankiw and Reis also show that in a setting characterized by sticky information, disinflation is accompanied by a recession.

While the sticky information Phillips curve better matches certain features of the data than the purely forward-looking Calvo specification, the Mankiw and Reis model still has drawbacks. Short-lived supply shocks generate little inflation inertia in the sticky information setup, but this is in marked contrast to the empirical evidence. This defect is likely one of the reasons why, at least so far, the sticky information model has found limited empirical success. Kiley (2007), in particular, shows that lagged inflation still enters significantly when included as an additional regressor in the estimation of (11). It is possible, though, that the sticky information model complements other forms of price rigidity, as suggested in recent empirical work by Dupor, Kitamura, and Tsuruga (2008).

Mankiw and Reis’s sticky information setup is still grounded in an environment in which agents can acquire and process *all* relevant information, albeit intermittently. A different class of models is based instead on the assumption that agents have limited information processing

capacities, and therefore cannot attend perfectly to all available information. Differences then arise between publicly available information and the private information agents use in their decisionmaking. Since information processing capacity is limited, agents employ this capacity optimally. This “rational inattention” framework, proposed in a series of papers by Sims (1998, 2003, 2006), underpins recent work that tries to explain both the macroeconomic and microeconomic features of price dynamics.

Work by Maćkowiak and Wiederholt (2008) considers a setting in which, because of limited information capabilities, price-setting firms must decide whether to pay attention to idiosyncratic or to aggregate conditions. If idiosyncratic shocks are much larger than aggregate shocks, firms will rationally devote more attention to idiosyncratic shocks than to aggregate shocks. As a result, prices respond quickly to idiosyncratic shocks and slowly to aggregate shocks. The model can thus generate inflation inertia in response to a monetary policy shock, even if firms are able to change prices in every period. The model is also consistent with empirical evidence provided by Boivin, Giannoni, and Mihov (2009), who show that sectoral prices respond quickly to sector-specific shocks, and slowly to monetary policy shocks.

Paciello (2008) complements Maćkowiak and Wiederholt’s work by examining a general equilibrium framework in which the only friction present is given by the firm’s limited information-processing ability. In contrast to Maćkowiak and Wiederholt, Paciello considers two aggregate sources of shocks, those stemming from either a technology shock or a monetary policy shock. Firms opt to be better informed about technology shocks because these disturbances are more volatile than monetary policy shocks, and thus affect profit-maximizing prices relatively more. As a result, inflation is more responsive to productivity shocks induced by technical change than to monetary policy shocks, a finding consistent with the empirical evidence in Altig et al. Whether Paciello’s model is consistent with other empirical evidence remains to be seen. This model delivers strong predictions regarding changes in the response of inflation to productivity and to monetary policy shocks—as a function of changes in the volatility of monetary policy shocks relative to the volatility of productivity shocks over time. Other things equal, if the volatility of mon-

etary policy shocks relative to technology shocks has declined over time, Paciello's model would predict a more inertial inflationary response to a monetary policy shock.

More generally, this class of theoretical models based on rational inattention still needs to undergo more comprehensive empirical testing. But the development of rational inattention models represents a promising avenue of research, which has already shown the potential to explain some empirical findings that are hard to reconcile within the more standard Calvo-style New Keynesian Phillips curve setup.

Microeconomic Evidence on Price-Setting Behavior: Do the Theoretical Models Square with the Empirical Evidence?

Theoretical developments on inflation dynamics since the early 1980s have stressed the importance of providing micro-foundations to describe the Phillips curve relationship. Much of the empirical work that has tried to fit micro-founded versions of the Phillips curve discusses the implied degree of price stickiness in the estimated Phillips curve; in other words, these models try to estimate the frequency with which firms, on average, change their prices. This frequency of price adjustment is a crucial deep structural parameter in micro-founded Phillips curve relationships, as it governs the size of the inflation-activity trade-off.²² Yet, until recently, there were a limited number of studies on micro price dynamics. Some of these studies looked at newspapers and retail catalogs (see, for example, Cecchetti 1986 and Kashyap 1995), while others looked at the prices of intermediate products in manufacturing (Carlton 1986). These papers documented that certain wholesale and retail prices could go unchanged for several months. More recently, a broader set of micro price data has become available, which has made it possible to obtain broader evidence on the extent of price rigidity and its implications for inflation dynamics.

Bils and Klenow (2004) use unpublished data from the U.S. Bureau of Labor Statistics (BLS) for 1995 to 1997 on the monthly frequency of price changes for 350 categories of consumer goods and services comprising around 70 percent of consumer expenditures. In contrast to the previous literature, Bils and Klenow find that prices change fairly frequently, with half of prices lasting 4.3 months or less. Structural esti-

mates of the New Keynesian Phillips curve—whether or not embedded into a DSGE model—usually produce a much lower frequency of price adjustment. Bills and Klenow also show that the standard Calvo pricing model produces price changes that are much more persistent and less volatile than in micro price data. This is especially true for those goods with less frequent price changes.

Subsequent work has built on Bills and Klenow and, in addition to further exploring the frequency and size of price adjustment, has documented other features of the BLS dataset. Nakamura and Steinsson (2008a) and Klenow and Kryvtsov (2008) note that the issue of how frequently prices change is complicated by the presence of sales and forced item substitutions in micro price data. For example, Klenow and Kryvtsov show that when sale-related price changes are removed, the estimated median price duration increases from 3.7 to 7.2 months. Nakamura and Steinsson note that sale price changes are more transient than regular price changes, and in most cases a price returns to its original level after a sale price offer ends. The estimated median price duration further increases when forced item substitutions (usually a product upgrade or a model changeover) are excluded from the data. The relevance of sale prices and forced item substitutions for explaining the dynamics of inflation at an aggregate level is still open to question. Some sales and forced item substitutions are likely a function of the business cycle, so that calibrating aggregate Phillips curve relationships with a median price duration that excludes sales and forced item substitutions may not be entirely justified.

Another feature that emerges from the micro price data is that, when prices change, they tend to change by a large amount, on average. Klenow and Kryvtsov document that the median absolute size of a price change is 11.5 percent, versus an average monthly inflation rate of 0.2 percent over the sample period they consider. Golosov and Lucas (2007) develop a menu-cost model with idiosyncratic and aggregate shocks to match the micro price data features in Klenow and Kryvtsov. Yet the Golosov and Lucas model is not consistent with large real effects of monetary policy shocks. Intuitively, in Golosov and Lucas's state-dependent pricing model, even if firms do not react to the monetary shock because the monetary shock alone is not large enough to justify paying the menu cost incurred in changing prices, many firms still engage in re-pricing because

of the presence of large idiosyncratic shocks. Once a firm decides to re-price for any reason, it will take the monetary policy shock into account when choosing the new price. As a result, monetary policy shocks have large and rapid effects on aggregate prices and little impact on economic activity.

The evidence from micro price data on the presence of relatively flexible prices contrasts with the well-documented persistence in aggregate inflation. Consequently, Phillips curve specifications that try to match a persistent aggregate inflation process have difficulties in matching the more flexible disaggregated price data. In contrast, models such as Golosov and Lucas that calibrate relatively flexible individual prices generate predictions for aggregate shocks that do not square well with most of the extant empirical evidence. This conundrum has found a potential resolution in Boivin, Giannoni, and Mihov's aforementioned research. Their work addresses a limitation common to the recent literature on micro price data; namely that these studies do not distinguish between idiosyncratic and aggregate sources of price changes. As such, micro-based models of price-setting behavior do not answer the question of whether disaggregated prices respond differently to idiosyncratic and aggregate shocks. The Golosov and Lucas model, for example, implies that there is essentially no difference in the responsiveness of prices to idiosyncratic and aggregate shocks. Boivin, Giannoni, and Mihov address this issue from an empirical standpoint, and reach a different conclusion. They show that sectoral prices respond sizably and rapidly to sector-specific shocks, but respond only sluggishly to aggregate shocks such as a monetary policy shock.

The differential responses of sectoral prices to sector-specific versus aggregate shocks limits the ability of evidence from micro price data to inform the development of an aggregate micro-founded Phillips curve relationship. Even so, the micro price data should still provide some discipline at the macro level. For example, the Calvo price setup implies that older prices, when altered, should change by a larger amount than prices subject to more frequent alteration. This feature of the Calvo setup finds no support in the empirical micro price data. More generally, micro price data reveal little correlation between the size of price changes and inflation. Instead, the frequency of price adjustments is strongly correlated

with the level of inflation. These findings run counter to the Calvo pricing model, which predicts a perfect correlation between the average size of price changes and inflation.

Time-Varying Inflation Targets and the Phillips Curve

The central bank controls the rate of inflation in the long run, and no explanation of the behavior of inflation can abstract from the role played by the monetary policy authority. Indeed, the general equilibrium models we have surveyed explicitly account for a monetary policy rule, often in the form of a reaction function à la Taylor. Yet these models usually constrain the central bank to maintaining a constant inflation target. The micro-founded Phillips curves we have discussed are also obtained as a log-linearization around a zero steady-state level of inflation. Both of these assumptions regarding inflation are counterfactual, and the question is whether relaxing these assumptions leads to different implications for the dynamics of the Phillips curve.

Kozicki and Tinsley (2002) consider a Phillips curve specification in which the nominal inflation anchor is not zero. The nominal anchor is estimated from a four-variable VAR (the variables included are inflation, the output gap, the ten-year Treasury yield, and the federal funds rate) with shifting endpoints using Kalman-filtering techniques to deal with the time-varying inflation target, assumed to evolve as a random walk. After retrieving an estimate of the time-varying inflation target, Kozicki and Tinsley estimate alternative Phillips curve specifications, in which inflation is expressed as a deviation from the estimated target. In principle, the time-varying inflation target, if varying enough, could lead to estimates in Phillips curve specifications that differ from the corresponding specifications in which inflation is not expressed as a deviation from the target. There is ample evidence that from 1960 to the present, the implicit inflation target set by the Federal Reserve has changed. When not explicitly accounted for, this change in the target could result in overstating the degree of persistence in inflation.²³ Still, even after accounting for low-frequency changes in the inflation target, the general conclusion in Kozicki and Tinsley is that while shifts in the long-run inflation anchor have contributed to the observed persistence of U.S. inflation, such shifts

do not appear to explain all of the historical persistence in inflation. Hybrid Phillips curve specifications of the deviation of inflation from its nominal anchor explain the historical behavior of inflation better than purely forward-looking specifications.

Cogley and Sbordone (2008) reach different conclusions. They explicitly derive a Calvo-style price Phillips curve that allows for a time-varying (and thus non-zero) steady-state for inflation. The specification takes the form

$$(12) \quad \hat{\pi}_t = y_t^f E_t \hat{\pi}_{t+1} + y_t^b \hat{\pi}_{t-1} + \kappa_t x_t + \text{other terms},$$

where $\hat{\pi}$ denotes the deviation of inflation from the time-varying steady state and x represents real marginal costs. There are two differences with respect to the standard Calvo-style specification. Additional terms appear on the right side of equation (12). These include innovations to steady-state inflation, higher order leads of inflation expectations, and terms involving the discount factor and the growth rate of output. Empirically, Cogley and Sbordone show that these additional terms are not important, though in principle their omission could lead to biased inflation estimates. The second and more important modification is that, because of the time-varying steady state for the rate of inflation, the coefficients in the Phillips curve are now time-varying. When the steady-state rate of inflation changes, the parameters drift too.

Cogley and Sbordone estimate the steady-state rate of inflation as the Beveridge-Nelson trend component of inflation from a reduced-form VAR (which includes as variables inflation, output growth, real marginal costs, and the federal funds rate). They interpret movements in trend inflation as changes in the Federal Reserve's inflation target. After having estimated the time-varying inflation target, Cogley and Sbordone proceed to estimate the Phillips curve specification (12). Their estimates indicate no role for lagged inflation: the coefficient y_t^b is always estimated to be close to zero. In sum, once removing trend inflation, Cogley and Sbordone conclude that the inflation process is well captured by a purely forward-looking Phillips curve specification.

It is not clear at this point what accounts for the different results in Cogley and Sbordone's findings relative to those in Kozicki and Tinsley. The estimation method is different, and Cogley and Sbordone's specifica-

tion embeds time-varying coefficients, which are not a feature of Kozicki and Tinsley's specifications. In all, Cogley and Sbordone allow for a more flexible specification, and time-varying coefficients could be one reason for the different findings. The more general point that both Kozicki and Tinsley and Cogley and Sbordone make remains well taken—that when thinking about inflation dynamics, it is important to account explicitly for low-frequency movements in inflation that result from a changing inflation target. It is likely that future empirical tests of micro-founded Phillips curve specifications will encompass such a feature. Moreover, time-varying coefficients in the Phillips curve specification allow for changing inflation dynamics, which, as we discuss in the next section, appear to have played an important role for at least part of the last 45 years.

Empirical Challenges and Pragmatic Implementation Issues

This brief intellectual history of the Phillips curve illustrates that much work has been done, both theoretically and empirically, since Phillips's seminal 1958 paper. It is also clear that economists have yet to converge to a widely agreed specification that is satisfactory both from a theoretical and an empirical standpoint. The Calvo-style New Keynesian Phillips curve setup has been employed extensively in theoretical frameworks and has been the subject of numerous empirical studies. Still, recent theoretical and empirical advances suggest that the Calvo-style New Keynesian Phillips curve could soon be displaced by alternative specifications. The current lack of consensus is disappointing, but it has generated a large and varied body of work that is leading to a much better understanding of the empirical features that a micro-founded model of inflation should ideally match. These features have been discussed in the previous sections. The persistence of inflation, the dynamic response of inflation to different macroeconomic shocks, and the response of sectoral inflation to sector-specific idiosyncratic shocks are all features against which to assess the empirical relevance of micro-founded Phillips-curve models. In addition, potential changes in the nominal inflation anchor and in the relative importance of different sources of shocks point to the need for having empirical specifications that can adequately capture changes in inflation dynamics. Here we underscore that most of the recent work has been

focused on price Phillips curves, but theoretical and empirical advances in explaining wage dynamics are also needed. Price and wage inflation are related in that the key driver for price inflation, real marginal costs, is also a function of wage dynamics. Thus, a better understanding of wage behavior could also lead to better models of price inflation.

Since the 1960s, the Phillips curve has been playing a central role in policymakers' understanding of the macroeconomy and in the formulation of monetary policy. It is not surprising then that empirical challenges in estimating a Phillips curve relationship have been closely intertwined with challenges in conducting monetary policy. Time variation in the NAIRU and/or in the potential rate of economic growth makes measuring the activity gap difficult in real time, thus posing important consequences for Phillips curve-based inflation forecasts. Several studies have attributed part of the increase in U.S. inflation experienced in the early 1970s to policymakers taking time to learn about an upward shift in the NAIRU and about a productivity slowdown (see, among other studies, Romer and Romer 2002, Orphanides 2003, and Orphanides and Williams 2005b). In the context of the Phillips curve, the failure to detect an upward shift in the NAIRU (or a decline in the potential rate of growth of the economy) results in inflation forecasts that are too optimistic and, thus, to an overly accommodative monetary policy stance, other things remaining equal.

Other studies have emphasized not only monetary policymakers learning about the NAIRU, but also their learning about the value of the other coefficients in a Phillips curve relationship. Primiceri (2006) interprets the run-up in U.S. inflation in the 1960s and 1970s and the subsequent disinflation of the early 1980s to policymakers learning about the persistence of inflation, the inflation-unemployment trade-off, and the NAIRU. Primiceri assumes that the true specification for the inflation process that the monetary authority is learning over time is a standard backward-looking Phillips curve,

$$(13) \quad \pi_t = \beta(L)\pi_{t-1} - \theta(L)(U_{t-1} - U_{t-1}^N) + \varepsilon_t,$$

where $\beta(1) = 1$. Primiceri argues that the early run-up in inflation was caused not just by the policymakers' misperception of the NAIRU but also, and more importantly, by the policymakers' underestimation of

the persistence of inflation. In terms of the Phillips curve specification in equation (13), policymakers took time to learn about an upward shift in U^N and about $\beta(1)$ being equal to unity. The monetary policymaking authority was operating with an estimate $\hat{\beta}(1) < 1$, which meant that policymakers believed that the inflation process was less persistent than in reality. Then during the mid-1970s, according to Primiceri, policymakers also became unduly pessimistic about the size of the inflation-unemployment trade-off, as measured by $\theta(L)$. In other words, policymakers thought that the sacrifice ratio—the increase in unemployment necessary to bring down inflation by one percentage point—was extremely high. As a result, the monetary authority did not lean strongly against the high levels of inflation during this period because it believed that an inflation-fighting policy would be too costly in terms of unemployment. The disinflation of the early 1980s reflected the Federal Reserve's better understanding of the true parameters in equation (13), most notably the Phillips curve's long-run verticality and the upward shift in the NAIRU. Further, by the early 1980s the Federal Reserve had come to believe that the sacrifice ratio needed to achieve disinflation was smaller than it had estimated previously.

Primiceri's interpretation of the rise and fall of U.S. inflation hinges on the policymaking authority operating with the correct Phillips specification, as in equation (13), but still having to learn about the relationship's true coefficients. Other studies have instead posited that the monetary policymaker operates with a misspecified Phillips curve (see, for example, Sargent, Williams, and Zha 2006). Overall, while economists disagree about what caused the rise and fall in postwar U.S. inflation, all of these studies highlight the crucial role played by some type of Phillips curve relationship in the conduct of monetary policy, and the importance of practical difficulties in obtaining an accurate real-time estimate of the Phillips curve.

How have the theoretical advances on micro-founded versions of the Phillips curve affected the contemporary conduct of monetary policy? At this point, micro-founded versions of the Phillips curve can be viewed as complementary to standard backward-looking specifications. So far, there is little evidence suggesting that forward-looking Phillips curve specifications provide more accurate inflation forecasts than a stan-

dard backward-looking specification, as in equation (13). As a result, the traditional backward-looking specification, possibly augmented to account for supply shocks, continues to play a role in shaping the inflation outlook and the conduct of monetary policy. Still, the importance of expected future inflation for determining current inflation is finding its way into the policy discourse.

We conclude this history of the Phillips curve by mentioning some empirical issues pertaining to the traditional backward-looking specification that bear on its usefulness as a tool for monetary policy.²⁴ This is not to diminish the importance of the recent theoretical and empirical advances. In contrast, we do so in order to highlight some empirical challenges that more micro-founded models will likely have to confront when modeling inflation dynamics.

A number of empirical studies have documented shifts in the backward-looking Phillips curve parameters.²⁵ The changes appear most pronounced for the effect that the relative price of oil has on core inflation. There is also evidence of a shift in the parameters on lagged inflation that seems widely accepted by economists; by contrast, evidence that the effect of the real activity variable may have diminished in recent decades is more contentious. The parameter shifts appear to be concentrated in the early 1980s, while the Phillips curve seems to have been *relatively* stable in the past 20 years. Accounting for a change in parameters in the 1980s is important in that, as shown in Fuhrer, Olivei, and Tootell (2009), it can dramatically improve the out-of-sample forecasting performance of the Phillips curve. Overall, these findings point to potentially important long-run changes in the dynamics of U.S. inflation. Yet an open question that deserves more study is finding the best way to accommodate changing inflation dynamics in the traditional backward-looking Phillips curve specification.

Interpreting shifts in the backward-looking Phillips curve can be difficult, as the framework is not explicit about many structural features of price-setting behavior. This is particularly true for the lagged inflation terms present in the traditional Phillips curve. In addition, the framework is mute regarding the behavior of other aspects of the economy that bear on inflation—notably the systematic behavior of monetary policy and the transmission channel from monetary instruments to output to inflation.

More micro-founded structural models of inflation and of the macro-economy will have to shed light on the nature of these shifts.

Still, even if economists achieve a better understanding of the reasons behind the shifts in the Phillips curve parameters, some of the movements in inflation, especially in the post-1984 period, are likely to remain difficult to explain in the context of a traditional Phillips curve framework. The forecasting performance of the traditional backward-looking Phillips curve over some of the post-1984 period was not different from the forecasting performance of a time-varying univariate autoregressive process for predicting inflation, even when accounting for changing inflation dynamics in the Phillips curve (see Fuhrer, Olivei, and Tootell 2009). This is true for the late 1990s, and for the most recent period (in this last instance, more so for core PCE than for core CPI inflation). These episodes highlight that all is not well even with the traditional backward-looking Phillips curve. For example, from mid-2003 until mid-2005, the Phillips curve would have predicted a fall in inflation, as the unemployment rate was relatively high. Yet contrary to the Phillips curve prediction, inflation picked up over this period. This discrepancy suggests that a more empirically satisfactory model of the inflation process will necessarily involve a deeper understanding of the determinants of the structural shock to the Phillips curve relationship, as this shock is playing an important role in inflation developments.

2. Overview of the Book

As the foregoing history illustrates, 50 years after its debut in 1958, the Phillips curve framework remains a key expository and forecasting tool in academic and policymaking circles. Yet despite important theoretical developments and the availability of rich new data on micro pricing behavior, economists have yet to agree on a satisfactory form for a micro-founded model of inflation. Foremost among the remaining challenges are developing micro-based macroeconomic models that can 1) match the empirical features of disparate aggregate and micro price behavior, 2) incorporate the (yet to be established) determinants of inflation expectations, and 3) reflect ongoing changes in inflation dynamics and structural shocks. Far from being purely academic issues, these are practical

matters that official forecasters grapple with on a daily basis. In 2007 and early 2008, for example, policymakers were struggling to determine the impact of surging oil prices on actual and expected core inflation, and to establish whether an explicit inflation target helps to anchor inflation expectations.

Accordingly, the Boston Fed invited leading academics and policymakers to Cape Cod in June 2008 to review the 50-year evolution of the Phillips curve and to assess what we know about inflation; our hope was to help stretch the boundaries of our knowledge and, thus, to strengthen the conduct of monetary policy. The ensuing conference sessions covered a range of challenging issues, including Stock and Watson's discussion of the predictability of inflation and the relative performance of alternative Phillips curve- and non-Phillips curve-based forecasting models; Dickens's paper on improving estimates of the NAIRU and, thus, the unemployment gap via the Beveridge curve; Sims's presentation on finding attractive alternatives to rational expectations models and their implications for inertial inflation behavior and monetary policy; Maćkowiak and Smets's review of promising ways to model inflation while matching both the macro and micro evidence on price behavior; Ball's rethinking of the hypothesis that actual unemployment can shift the NAIRU; and the panel discussion by Fischer, Kohn, Stark, and Svensson, all central bank policymakers, about the driving need to better understand and address exogenous supply and structural shocks and inflation expectations. The rest of the book starts with a dialog between Solow and Taylor, moderated by Mankiw, on the first 50 years of the Phillips curve, and ends with remarks by Chairman Bernanke on unresolved questions pertaining to inflation. The rest of this section provides brief summaries of this volume's contents, highlighting how the issues discussed contribute to our evolving understanding of inflation and the practice of monetary policy.

Fifty Years of the Phillips Curve: A Dialog on What We Have Learned

The conversation between Bob Solow and John Taylor, moderated by Greg Mankiw, began by recalling the earliest days of the Phillips curve and these economists' first reactions to Phillips's 1958 article. As Solow explained, he and Paul Samuelson invented its name, "the Phillips

curve,” and Solow confessed finding the article’s “amazingly” stable empirical relationship between unemployment and inflation “remarkable.” In retrospect, John Taylor was struck that he was never tempted to try to exploit the long-run trade-off between employment (output) and inflation that occurs in the original Phillips curve (absent inflation expectations). Partly this was because Phillips himself did not view the curve as something for monetary policymakers to exploit. But Taylor also remarked that following Milton Friedman’s 1967 presidential address, the idea that shifts in the Phillips curve over time would eliminate any long-run trade-off spread rapidly, and policymakers generally adopted models with an expectations-augmented Phillips curve and slowly adaptive expectations.

In reference to Friedman’s 1967 address, Solow noted his own puzzlement as he gradually realized that Friedman had reversed—without sounding any bells or whistles—the direction of causality in the Phillips framework. Setting aside our sophisticated general-equilibrium quibbles, Solow suggested that we all know Phillips viewed causality as running from disequilibrium in the labor market to inflation. In Friedman’s version, the only way to push the unemployment rate away from its “natural” rate is to create an inflation rate that differs from expectations—a “contradictory” kind of causality that Solow does not find “plausible, not remotely.”

Turning to current efforts to understand the relationship between inflation and unemployment based on the Taylor-Calvo versions of the New Keynesian Phillips curve, Mankiw asked whether the profession is heading in the right direction. Taylor pointed out that following Friedman’s adaptive expectations came Lucas’s rational expectations, and then the need to explain the puzzling persistence in observed inflation and the observed effect of monetary policy. In searching for an explanation, economists began to observe that prices and wages are not reset every period but instead last a while—in Taylor’s version because of staggered contracts.²⁶ Taylor thinks that measured by the standard of what we get out of it, this approach has been very useful since, among other things, it leads to simple equations with an important role for inflation expectations, and a prediction that the more aggressive the monetary policy, the less the inflation inertia.

Solow's assessment was less positive because he does not believe the New Keynesian Phillips curve premise that if inflation is constant and is expected to remain constant, output settles at the "natural" rate, now or later. What he does like about the New Keynesian Phillips curve is that it allows economists to embed what looks like a Phillips curve (but isn't²⁷) in a model with an IS curve and a Taylor rule. As a result, researchers can talk rigorously about causality, which they couldn't otherwise. Solow also doubts that the current New Keynesian Phillips curve can produce adequate inflation persistence.

Reminded that Solow had also expressed skepticism that the long-run unemployment rate is unaffected by the rate of inflation pegged by the Fed, Taylor affirmed his strong belief that the natural rate of unemployment is invariant to monetary policy—as are trend productivity growth and the real interest rate—as a principle and an approximation. These values come from the real economy, and he finds the classical Phillips curve dichotomy to be useful in making this distinction. Indeed, the more we can convince people that the natural rate is invariant to monetary policy, the better. In rejoinder, Solow pointed out that we are talking about a theory in which the two central concepts—the natural rate of unemployment/output and the expected rate of inflation—escape observation, elude clear definition, and jump around a lot. This kind of instability causes difficulty for economists, who are left to explain that inflation is accelerating because the unemployment rate is below the "natural" rate. How do we know? Because the rate of inflation is rising . . .

In the mid-2008 context of soaring oil prices, Mankiw asked whether we have a good way to think about how relative price changes fit into overall inflation. Solow replied that while some economists would argue that inflation is everywhere and always a monetary phenomenon and there is no reason why a relative price increase should affect the aggregate price level as long as other prices fall "enough," we all know that it may be hard to achieve the relative price change required by the market without a rise in the general price level. But, Taylor replied, it is also important to remember that commodity price shocks tend to pass through to other prices, and, empirically, the amount of pass-through to general inflation tends to be lower in countries where monetary policy is focused on delivering low inflation.

In drawing their conversation to a close, Mankiw asked Solow and Taylor to identify the big unanswered questions waiting for the next generation of macroeconomists. Solow suggested that large structural shifts in the economy, like the U.S. economy's relative shift from goods to services and the increased importance of global competition, deserve careful study. He is convinced that such changes affect price behavior, particularly the aspects the Phillips curve tries to capture. Taylor hopes for more efforts to test price- and wage-setting models against the micro data. In addition, he wonders whether the current search for *the* microeconomic theory of price adjustment is ill-advised and if we are searching for something we are never going to find. Perhaps economists should be looking for better macroeconomic equations, particularly macro price equations that incorporate many different types of price adjustment at the micro level.

Phillips Curve Inflation Forecasts

Phillips's 1958 paper examines data across a great sweep of history (1861 to 1957) and in Solow's term, documents an "amazing" empirical relationship between the change in nominal wages and unemployment over the first century of the modern industrial era. As Sims points out in his paper included in this volume, Phillips's insight gave Keynesian economists a much-needed way to measure how far the economy was from capacity and to make quantitative forecasts of how aggregate demand would affect inflation. Since then, the Phillips curve has remained a staple framework in most policymakers' tool boxes. Continuing this tradition, in their paper for this volume, Stock and Watson also examine data for an extended period, 1953:Q1 to 2008:Q1, to evaluate the relative success of inflation forecasts that use a Phillips curve-type activity measure and those that do not. They conclude that while forecasting inflation is hard, the evidence suggests that Phillips curve forecasts do not generally improve on *good* univariate models. Nevertheless, "the backward-looking Phillips curve remains a workhorse of large macroeconomic forecasting models and continues to be the best way to understand policy discussions about the rates of unemployment and inflation."

In addition to a period of simultaneously high inflation and high unemployment, the 1970s also ushered in the powerful idea of rational expecta-

tions, Milton Friedman's reverse-direction Phillips curve, and the premise that the best monetary policy is a predictable policy. As a result, academic economists and policymakers tended to go their separate ways. Although much of the academic literature of the 1970s focused on introducing wage and price rigidities into models with rational expectations, monetary policymakers tended to eschew this new framework, and continued to rely on models in which current inflation was a function of lagged inflation and the unemployment rate. From the policymaking standpoint, the major challenge in modeling inflation was to build in a role for exogenous changes in food and energy prices, such as those that occurred in the mid-to-late 1970s. Once supply shocks were built into macroeconomic models, the older models appeared to provide reasonably reliable explanations of observed inflation, at least through the early 1980s. In the U.S. context, for example, the decline in inflation that accompanied the dramatic rise in unemployment under Federal Reserve Chairman Paul Volcker was consistent with the predictions of a backward-looking Phillips curve model.

By the mid-1990s, scholars began to detect a noticeable deterioration in Phillips curve-based inflation forecasts. In particular, an influential paper by Atkeson and Ohanian (2001) concluded that since 1985 forecasts of U.S. inflation based on a Phillips curve specification did not improve upon forecasts based on a simple univariate model. To investigate the predictability of inflation more systematically, James Stock and Mark Watson undertake a comparison of the out-of-sample performance of alternative models of inflation using a single consistent data set for the United States in the period spanning 1953 to 2008. Their study encompasses different measures of inflation (for example, core versus total, consumption-based versus economy-wide) as well as a variety of univariate and multivariate model specifications.

For the sample period as a whole, the Stock-Watson (2007) unobserved components-stochastic volatility (UC-SV) model has better overall performance than the other univariate models and all of the multivariate models. In the UC-SV model, inflation has a permanent component z_t and a temporary component ε_t :

$$\pi_t = z_t + \varepsilon_t, \text{ where } \varepsilon_t = \sigma_{\varepsilon,t} \zeta_{\varepsilon,t},$$

$$z_t = z_{t-1} + u_t, \text{ where } u_t = \sigma_{u,t} \zeta_{u,t},$$

$$\ln \sigma_{\varepsilon,t}^2 = \ln \sigma_{\varepsilon,t-1}^2 + v_{\varepsilon,t},$$

$$\ln \sigma_{u,t}^2 = \ln \sigma_{u,t-1}^2 + v_{u,t},$$

where $\zeta_t = (\zeta_{\varepsilon,t}, \zeta_{u,t})$ is independent and identically distributed (i.i.d.) $N(0, I_2)$, $v_t = (v_{\varepsilon,t}, v_{u,t})$ is i.i.d., $N(0, \gamma I_2)$, ζ_t and v_t are independently distributed, and γ is a scalar parameter.

Stock and Watson find that Phillips curve-based models provide reliably superior forecasts only during the 1970s and early 1980s. The authors also conclude that the choice of the activity variable (unemployment, output, or the principal component of many economic activity indicators) in such a model is secondary to the choice of whether to use an activity-based model for making inflation projections.

Reexamining the findings according to the size of the gap between the actual rate of unemployment and the NAIRU, Stock and Watson find that univariate models tend to provide better forecasts when the unemployment gap is small, as compared to models that incorporate a Phillips curve. On the other hand, Phillips curve models perform better than univariate models when the gap is large—that is, around economic turning points. Thus, Stock and Watson's findings suggest that central banks may be justified in lowering their expectations of inflation during recessions. During less extreme phases of the business cycle, unemployment and other economic activity variables tend to be unreliable in gauging the likely direction of inflation.

Adrian Pagan, the first discussant, maintains that although purely statistical models may win forecasting competitions, central banks need to incorporate economic variables into their projections of inflation in order to explain their policy decisions to the public. Thus, models based on the Phillips curve serve a communications function, even if policymakers use univariate models in the background to refine their projections for inflation.

Pagan notes that an essential feature of the data for the United States and other nations is that the inflation-generating process varies over time. Stock and Watson's UC-SV model captures this time variation through its stochastic volatility feature. However, stochastic volatility may not be an essential feature of a superior forecasting model. Pagan proposes alternative models with time-varying parameters or time-varying estimation

windows that are likely to be more palatable for central banks—especially those with an inflation target. Moreover, research applying such methods to data from Australia and the United Kingdom suggests that time-variation patterns are affected by the state of the real economy.

Lucrezia Reichlin, the second discussant for the Stock and Watson paper, points out that the mid-1980s marked not only the relative deterioration in the usefulness of the Phillips curve for forecasting inflation, but also the start of the so-called Great Moderation, in which output volatility declined. A less-known fact is that the ability of accepted economic models to predict output growth also declined during this period. Reichlin argues that the performance of the Phillips curve should be evaluated in this broader macroeconomic context.

Reichlin uses a macroeconomic VAR model to investigate the causes of the changes in volatility and predictability of both inflation and GDP. She concludes that the patterns since 1984 can be explained by changes in how shocks are propagated through the economy, rather than by changes in the variability of shocks. Thus, it seems plausible that improvements in macroeconomic policy have brought about smoother but less predictable movements in both inflation and output over time.

Obtaining More Precise Measures of the NAIRU

A theme running throughout our discussion of the Phillips curve, particularly the New Keynesian version, is that its central concepts—the natural rate of unemployment (alternately, the NAIRU) or output (or marginal cost) and the expected rate of inflation—are hard to define, hard to measure, and, as far as anyone can tell, not very stable; these are “three suspicious characteristics,” as Bob Solow describes them in the opening session. These difficulties help to explain why Phillips curve-based models can be hard to interpret and why Phillips curve-based forecasts are not always successful. In response to this challenge, William Dickens explores a way to improve our measures of one of these key unobservable concepts: the NAIRU.

For the concept of the natural rate of unemployment or the NAIRU to provide a meaningful guide for policymaking purposes, it must be measured reasonably accurately, as just suggested. Unfortunately, the accepted

practice of backing out an estimate of the natural rate or NAIRU through an econometric estimation of the Phillips curve relationship has serious limitations, as noted by Ball and Mankiw (2002) and other authors.

To ground these ideas, consider the following basic form of the Phillips curve,

$$\pi_t = \pi_t^e - \alpha(U_t - U_t^*) + \varepsilon_t,$$

in which inflation is equal to its expected value, minus the difference between the actual and (potentially time-varying) natural rates of unemployment multiplied by a parameter plus a supply shock denoted by ε_t . This relationship can be used to solve for the natural rate of unemployment:

$$U_t^* = U_t + (1/\alpha) \{[\pi_t - \pi_t^e] - \varepsilon_t\}.$$

To obtain a numerical estimate of U^* , one must first specify expected inflation in terms of observables. As a long literature in macroeconomics has noted, this is by no means a straightforward proposition. Next, there is the issue of how to specify supply shocks over time. Without further restrictions, the distinction between ε_t and U_t^* is arbitrary: a change in either one shifts the Phillips curve. However, these terms represent distinctly different concepts. Some authors have distinguished between the two by assuming that supply shocks are relatively high-frequency movements attributable to factors such as oil price shocks or exchange rate movements, while the natural rate moves at lower frequencies in response to changes in labor market practices and institutions. Finally, regardless of how one chooses to identify supply shocks, there remains the difficulty of obtaining an estimate of the parameter, α . Jodi Galí discusses alternative approaches to estimating this parameter and the various shortcomings of these methods in his discussion of Laurence Ball's paper.

Given the various sources of uncertainty about how to derive the NAIRU or natural rate of unemployment from the Phillips curve relationship, it should not be surprising that the resulting estimates are quite imprecise. For example, Staiger, Stock, and Watson (1997) estimated the 95 percent confidence band for the NAIRU in the United States in 1990 to be between 5.1 percent and 7.7 percent. Yet shortly after this result was published, new estimates of the NAIRU using similar methodologies dropped the estimate below 5 percent for the 1990s.

In his paper for this book, William Dickens proposes a new methodology for deriving estimates of the NAIRU from the Beveridge curve, which is named for the British economist who first noticed a negative relationship between the unemployment rate and the job vacancy rate (defined as unfilled jobs relative to the size of the labor force). Movements along the Beveridge curve are indicative of cyclical conditions in the labor market: strong labor demand implies low unemployment and a high rate of unfilled jobs, and vice versa. Shifts of the curve stem largely from factors associated with the efficiency of matching workers with jobs (Blanchard and Diamond 1989), which should correspond to changes in the natural rate of unemployment. Thus the location of the Beveridge curve can provide additional information about the locus of the NAIRU.

To implement the new methodology, Dickens derives a specification for the Beveridge curve from a gross flows model in which jobs are created as new firms are formed and jobs destroyed as existing firms cease production, together with a hypothesized functional form of the process by which unemployed workers are matched to jobs. Econometric estimation using U.S. data for time periods during which the unemployment-vacancy relationship appears to be stable yields plausible, precise parameter values. The estimates are virtually unchanged when Dickens uses different methods to account for shifts in the unemployment-vacancy relationship.

If changes in match efficiency coincide with changes in the natural rate of unemployment, then U_t^* should enter the expression for the Beveridge curve. Thus estimates of the Beveridge curve should augment Phillips curve-based information about the location of the NAIRU, potentially improving the accuracy of estimates derived from the Phillips curve alone. Dickens jointly estimates the Beveridge curve and Phillips curve relationships. For the 1961–2007 estimation period as a whole, Dickens is able to reduce the uncertainty of the NAIRU estimates by about 30 percent, compared to those derived using the Phillips curve alone.

Dickens's study is hampered by the lack of consistent measures of job vacancies over time. For the earlier years of his study, Dickens's must infer vacancy rates from data on help-wanted advertising. For more recent

years, he is able to obtain vacancy data directly from the Job Openings and Labor Turnover Survey (JOLTS), initiated by the U.S. Bureau of Labor Statistics in 2000. Dickens expects the precision of the estimates to improve as the available observations on vacancies and unemployment increase over time.

Despite their limitations, the new estimates provide information about the timing of shifts in the natural rate of unemployment. As Dickens writes, “While NAIRU values much above 6 percent can be ruled out during the 1960s and mid-to-late 1990s, values less than that can be ruled out for the decade starting in 1978. This provides more guidance to policymakers than past estimates” (see p. 225).

In his discussion, Olivier Blanchard points to how the methodology adopted by Dickens casts light on the relative importance of various structural shifts in labor markets. Many observers have attributed the low inflation and unemployment throughout much of the 1990s and 2000s to the effect of globalization in reducing the bargaining power of workers. Indeed, that is a theme included in Paul Samuelson’s foreword to this volume. By attributing the recent decline in the natural rate to a shift in the Beveridge curve rather than in the Phillips curve, Dickens implies that globalization may not have been the main driver of the inflation and unemployment patterns seen over the past decade or so. Instead, the likely drivers were an increase in the efficiency of matching of unemployed workers to jobs (for example, as a result of new Internet-based technologies) or reduced flows of workers into unemployment (as a result of either decreased worker separations from jobs or increased hiring of workers from outside the labor force or from among the already-employed labor force). The U.S. data strongly suggest that the answer lies mostly in reductions in worker separations, which in turn were caused mostly by diminished rates of job destruction rather than of worker quits. Blanchard concludes that a challenge for future research is to investigate the causes of this decline in job destruction.

Christopher Pissarides remains skeptical that Dickens has identified changes in the natural rate. In his view, the method fails to uncover changes in the natural rate that occur while the economy remains on a fixed Beveridge curve—that is, when match efficiency remains unchanged.

Pissarides advocates developing and estimating a model of the labor market that accounts for the endogeneity of job separations.

Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy

As noted earlier, expected inflation is another (increasingly) central concept in Phillips curve analysis that is unobservable, ill-defined, and hard to pin down. Christopher Sims's paper takes up this second concept and traces how the treatment of inflation expectations in the Phillips curve framework has evolved over time.

Sims posits that inflation expectations first entered Phillips curve equations in a sustained manner after Lucas, in part with Rapping, developed a reversed-direction, rational expectations model in the late 1960s and early 1970s—but contends that this early treatment of expectations was either too abstract and unrealistic or too simple and innocent of theory or micro foundations when included in policymaking models. While the New Keynesian Phillips curve—with its continuum of monopolistically competitive firms, rational expectations, and Taylor- or Calvo-type price-setting frictions—attempted to fill the gap, Sims points out that the New Keynesian Phillips curve approach merely moves non-neutrality from agent behavior to the pricing frictions—in other words, to the contract lengths, which are “not constants of nature” and “will surely change systematically with the level, variability, and forecastability of inflation.” A further problem, Sims argues, is that once inflation expectations enter the Phillips curve framework, it becomes possible in principle for a disturbance to impact inflation directly through the expectations term rather than indirectly through its effect on real tightness. Looking for empirical evidence regarding the relative importance of a Phillips curve-type mechanism (i.e., some measure of tightness) in determining inflation, Sims then presents a set of monetary structural VARs and concludes that while monetary policy is definitely not neutral in its effects on output, thinking about the determinants of inflation in terms of the New Keynesian Phillips curve does not seem particularly helpful.

Where do we go from here? Looking ahead, Sims's answer to the question “where does the persistence in inflation come from?” suggests

looking at the implications of models with learning, behavioral economics, intermittent observation and—particularly promising, he suggests—models with rational inattention and models in which rational agents share the same information and the same range of outcomes but disagree about the probability distribution for those states. Sims points out that in a world characterized by rational inattention, agents will behave as if they observe market signals with error and, because these agents have different incentives to invest in processing a given bit of new information, they will have different probability distributions for a given set of possible outcomes.²⁸ While it is clearly hard to model rational inattention or heterogeneous assumptions regarding probability distributions, it is even harder to imagine that economic agents do not behave in these ways. Thus, it is important to incorporate these assumptions in future models so as to provide appropriate guidance for monetary policy.

Michael Kiley agrees that information constraints play a crucial role—along with sticky prices and other nominal rigidities—in explaining U.S. inflation dynamics. And to Sims's emphasis on the costs of acquiring and processing information, Kiley would add 1) the nontrivial cost of calculating optimal actions under uncertainty and highly nonlinear objective functions as well as 2) imperfect knowledge about the central bank's goals. In the latter case, where the central bank's objectives are not explicit or well understood, households and firms will need to infer the inflation goal from the central bank's actions, albeit with delays and mistakes. Thus, he suggests, the nature of the monetary policy regime is an important determinant of inflation expectations. Orphanides, who has done groundbreaking work with Williams (2005) on the role of learning in the formation of inflation expectations,²⁹ makes a similar point about the role of central bank communications in clarifying its inflation objective.

While Sims concludes that something like the Phillips curve will continue to have a role in general equilibrium models as a way of drawing the links between costs, prices, wages and output, he argues that the rational inattention perspective suggests that locating inertia only in that one equation may be a mistake, since the same limits on information processing may also be at work in the slow reaction of consumption to income or investment to interest rates. Sluggish responses of various kinds may

be related through their dependence on a common resource constraint. Recognizing that commonality may lead to new ways to assess the welfare implications of monetary policies designed to achieve price stability.

Implications of Microeconomic Price Data for Macroeconomic Models

As pointed out in the above history of the Phillips curve, economists have sought to improve the micro foundations of Phillips curve analysis since the early 1980s, but, until recently, data limitations have constrained their efforts. Over the past decade, however, the situation has changed markedly—with big improvements in the breadth, detail, and frequency of micro-level price data sparking a surge of new work in this area.

Reviewing what economists have learned about micro pricing behavior, Maćkowiak and Smets in this volume examine a number of papers that explore the wonderfully rich U.S. and European data from the Bureau of Labor Statistics (BLS) and Inflation Persistence Network (IPN), as well as survey results and newly available scanner data.³⁰ While contending that the question “what do the micro data say?” has no simple answer, Maćkowiak and Smets also see a number of regularities across the U.S. and European data that confirm several of Taylor’s 1999 findings.³¹ First and of key importance, in both the United States and the euro area the data reveal much heterogeneity across sectors in the frequency and size of price changes and in the frequency and form of sales and forced item substitutions.³² Still, in many sectors, prices remain constant for extended periods—primarily, according to the survey data, because firms want to avoid disrupting long-term relationships with their customers. By contrast, menu and information costs are generally reported to be relatively unimportant. Finally, as in Taylor (1999), prices change a lot relative to inflation, on average, and, in cross-country regressions, the frequency of price change depends positively on the average rate of inflation. There is little evidence of synchronization.

Providing further detail, the authors cite a related study by Maćkowiak, Moench, and Wiederholt (2008), which finds that most of the considerable variation in sector price indexes is triggered by sector-specific shocks and occurs within a month—meaning that sector price indexes are not sticky at all. By contrast, sector price indexes respond only slowly to

aggregate shocks; just 15 percent of the long-run response occurs within a month. Thus, the degree of price stickiness appears to depend on the source of the shock. Further, Maćkowiak, Moench, and Wiederholt (2008) also observe that the frequency of sector-specific price changes helps to explain cross-sector differences in the speed of impulse responses of prices to macro shocks—as could be consistent with the menu cost model, the imperfect information model (Reis 2006) and the rational inattention model (Maćkowiak and Wiederholt 2009).³³

How well do standard macroeconomic models of price setting (for example, the Calvo model and the menu cost model of Golosov and Lucas [2007]) actually relate to the new micro data? Maćkowiak and Smets point out that while the micro data support the basic premise underlying both the New Keynesian and the Neoclassical Synthesis that many prices stay fixed for extended periods, the micro data are so detailed and the models are so simple that some aspect of each is bound to be rejected. As a result, just how models of price rigidity that fit the micro data can imply the relatively slow impulse responses to macro shocks seen in the aggregate data remains a matter of much controversy—although, as King points out in his comments included in this volume, the micro data also provide useful discipline and should help to distinguish between the macro models.

Ideally, of course, macroeconomists would like DSGE models that allow much heterogeneity and match both the detailed micro data and the macro data as well. Realistically, however, Maćkowiak and Smets believe the best we can hope for right now is a model that matches the macro data well and tells a “reasonable story...broadly in line” with the micro data. In pursuit of such a model, Maćkowiak and Smets examine the outcome of calibrating several menu-cost and other state-dependent models to match some features of the micro data and find the results to be problematic. For example, as Midrigan concurs, menu-cost models like Caplin and Spulber (1987) and Golosov and Lucas (2007) produce monetary neutrality (with the aggregate price level responding one-for-one with the growth of money) because money has a strong selection effect in these models; the firms that choose to raise prices at a given point are those that need the largest price change. As a result, while these models can match the 10-percent average price change found in the BLS

data (Klenow and Kryvtsov 2008), the aggregate price level becomes more flexible than individual prices, contrary to the empirical evidence.

Further, as Maćkowiak and Smets, Midrigan, and King—indeed a growing consensus—agree, menu costs alone are unlikely to be large enough to produce a sizable monetary transmission mechanism. In support, Midrigan, one of the discussants, notes that sale prices usually return promptly to their exact presale level, that firms with sticky prices and firms with more flexible prices both *choose* whether to change prices with nominal exchange rates, and that if menu costs were the only friction, matching the observed slow response of the aggregate price level to nominal shocks would require that individual firms adjust prices every ten quarters, instead of every two to three quarters, as found in the data.³⁴

As for the Calvo model, many economists view its inherent lack of inflation persistence as a flaw, which some recent DSGE models have tried to address by adding “dynamic indexation.” Under such a scheme, a fraction of firms adjust their price each time period, with a small subset adjusting optimally and the rest adjusting by inflation at $t - 1$. But as King points out in his response to Maćkowiak and Smets’s paper, dynamic indexation is highly inconsistent with the micro data that show intervals of constant nominal prices, price declines as well as gains, and no tendency for price changes to cluster at last month’s inflation rate. Drawing on unpublished research from Nakamura and Steinsson (2008b), King shows that there is no strong relationship between the average size of price increases and inflation, as the Calvo model would predict. By contrast, inflation is strongly associated with the fraction of firms choosing to raise prices. Thus, King encourages that more effort be made to understand the timing rather than the size of micro price adjustments.

Since prices turn out to be less sticky than assumed in many DSGE models, Maćkowiak and Smets turn to “promising” approaches that reflect their observation that firms find it optimal to change prices by large amounts in response to firm- and sector-specific shocks but by small amounts in response to aggregate shocks. Nakamura and Steinsson (2008b) achieve this effect by introducing intermediate inputs (as in Basu 1995) while Kryvtsov and Midrigan (2009) introduce real rigidity at the macro level via sluggish wages. But in the end it is the rational inat-

tention and sticky information models that Maćkowiak and Smets find particularly appealing.³⁵ These models build on Lucas's much criticized idea that real effects of nominal shocks reflect imperfect information, buttressed by Sims's (2003) point that if agents have a limited capacity to process information, publicly available information may not be fully reflected in agents' decisions. As an example, Maćkowiak and Wiederholt (2008) develop a model in which information about the current state of monetary policy is widely available, but agents find it optimal to devote almost all of their limited information-processing capacity to monitoring idiosyncratic conditions and pay very little attention to macro policy shocks. In such a world, prices respond strongly and quickly to idiosyncratic shocks and weakly and slowly to aggregate shocks; the real effects of nominal shocks are strong and persistent, and the welfare costs of increased macro volatility are likely large.³⁶

What Determines the Natural Rate of Unemployment?

With the major economies now entering what could be an unusually long recession, Laurence Ball's paper on the determinants of the natural rate of unemployment addresses a topic of renewed policy concern. More generally, it also addresses the type of structural shock that Solow and others have urged deserves more research attention.

Most of the macroeconomics literature of the past four decades has accepted the Friedman-Phelps premise that monetary policy can move unemployment away from its natural rate only temporarily. The term "natural rate" is understood to be the level of unemployment consistent with aggregate production being at its long-run equilibrium level, given the structure of labor and product markets. Macroeconomists initially treated it as time-invariant, but this assumption became increasingly untenable in light of empirical evidence. European joblessness rose dramatically in the decade from the mid-1970s to the mid-1980s. More recently, the United States managed to reduce unemployment to an exceptionally low rate in the late 1990s without triggering an acceleration in inflation.

Prompted by such sustained movements in unemployment, economists turned to studying why the natural rate appears to change over time, as

well as why it appears to vary across countries. Most hypotheses focused on specific supply-side or exogenous influences, such as the demographic composition of the labor force, skill-biased technological progress, institutional factors such as legal and administrative restrictions on layoffs, and the structure of unemployment insurance benefits. By contrast, Blanchard and Summers (1986) introduced a more general explanation called “hysteresis”—the notion that the natural rate of unemployment can be influenced by the path of actual unemployment. If hysteresis were confirmed in the data, this could suggest that monetary policy has longer-lasting effects on unemployment than many economists had come to believe. Evidence of hysteresis could also be used to argue against having central banks focus exclusively on inflation, since doing so could have the unintended consequence of exacerbating unemployment over an extended period of time.

Building on his previous research, Laurence Ball’s paper in this book studies the relationship between unemployment and the NAIRU—which should move up and down with the natural rate—for a panel of 20 OECD nations for the period 1980 to 2007. Assuming that inflation expectations are determined on the basis of lagged inflation, the Phillips curve relationship posits that falling inflation is a sign that unemployment exceeds the NAIRU. Conversely, rising inflation indicates that unemployment is below the NAIRU. Ball derives NAIRU estimates from this framework, using a modified version of the method in Ball and Mankiw (2002). He then compares the estimated NAIRU series to actual unemployment to determine if increases (decreases) in the latter are followed by increases (decreases) in the former. If so, that might imply that high (low) unemployment caused a higher (lower) NAIRU.

The evidence supports hysteresis to some extent, but is not conclusive. Ball finds that all eight episodes with a substantial increase in the NAIRU were associated with a major disinflation, which is consistent with hysteresis. On the other hand, at most only five of the nine episodes with a substantial decrease in the NAIRU were preceded by sizable increases in inflation.

Ball calls for a renewal of research interest in the mechanisms underlying hysteresis. One possible explanation, originally suggested by Blanchard and Summers, concerns the behavior of the long-term unemployed. The

argument is that if the economy undergoes sustained weakness in aggregate demand, long-term unemployment is likely to increase. Workers who have been unemployed for an extended period of time become somewhat detached from the labor market, and therefore exert less downward pressure on wage rates than newly unemployed workers who are actively searching for a job. Thus, measured unemployment increases while wage inflation stabilizes. Ball finds this explanation quite plausible, and in his discussion of Ball's paper, Jordi Galí suggests testing the hypothesized mechanism directly by not including the long-term unemployed in the computation of joblessness. More generally, Galí anticipates that the current period of sharply rising unemployment will prompt new research that advances the understanding of hysteresis.

By contrast with Ball and Galí, in his remarks V.V. Chari maintains that the evidence to date—drawn from many countries and time periods—strongly rejects the plausibility that monetary policy has real lasting effects on the economy. Chari presents data indicating that real output growth is remarkably stable across a wide variety of policy regimes. Moreover, countries that have adopted inflation targeting in the last two decades have been able to achieve reductions in inflation without introducing any material changes in real-side variables. The disagreement expressed here indicates some of the rifts existing among contemporary macroeconomists.

Lessons for Central Bankers: A Panel Discussion among Monetary Policymakers

Since the original article appeared in 1958, the usefulness of the Phillips curve as a policy tool has been a topic of intense debate. How—if at all—are policymakers using Phillips curve analysis today, and what do central bankers view as the primary challenges to their use of this framework?

As a practical matter, most policymakers—including all who spoke at the Boston Fed conference in June 2008—appear to use Phillips curve-like or Phillips curve-type models to generate forecasts for their policy deliberations. As might be expected, they use more recent versions of the Phillips curve approach, and employ it as just one among several forecasting tools. For example, while Donald Kohn reports that models in the Phillips

curve tradition remain at the core of how he and other policymakers think about inflation,³⁷ he notes that the original Phillips curve has evolved over time to recognize the importance of expectations, the possibility of structural change, and the uncertainty surrounding wage and price dynamics. Moreover, while the Phillips curve framework incorporates expectations, supply shocks, and resource utilization—which Kohn views as the key drivers of inflation—he points out that the utilization-inflation link at the heart of the Phillips curve approach seems to account for a rather modest part of observed inflation fluctuations. In that light, how these inflation drivers interact becomes a pressing question. In particular, with analysts assigning an increasingly central role in the inflation process to inflation expectations and how these are formed, measuring these expectations, identifying their determinants, and keeping them “well anchored” appear to be high-priority issues for most central bankers.

One reason why central bankers have built eclectic arsenals, as Kohn suggests, may be that relatively successful forecasting exercises based on the Phillips curve framework frequently use reduced-form regressions with proxies for key, hard-to-measure variables (such as lagged inflation for inflation expectations and the unemployment or output gap for resource utilization). While Kohn considers such regressions to be among the best forecasting tools available, he also points out that lagged inflation is a very imperfect measure of inflation expectations. In particular, and despite the fact that reduced-form regressions imply that sharp jumps in oil prices have only modest effects on future inflation (expectations *given*), Kohn is concerned that repeated increases in energy prices may actually lead to a rise in long-term inflation expectations. In addition, Stanley Fisher, who also puts a good deal of weight on inflation expectations in setting policy, describes the difficulties of choosing between inconsistent measures of expected inflation and of trying to make policy in the wake of a significant and abrupt change in the monetary transmission mechanism.³⁸

Of course, for over 30 years policymakers have recognized the desirability of looking beyond reduced-form exercises—ever since Lucas stressed the need for structural models in analyzing the impact of any shock, like a change in the policy regime, that affects the decisions/behaviors of economic agents.³⁹ But today that route is strewn with challenges because economists have developed many structural models, each emphasizing a

different imperfection, bolstered by different amounts of empirical support, and conveying different policy implications (cf., Maćkowiak and Smets in this volume). Given these circumstances, Kohn suggests that the best approach for policymakers may be to look for the common lessons to be drawn from these models. Fortunately, he notes, many structural models of nominal wage-price adjustment imply the same general conclusions regarding the appropriate response of monetary policy to sharp increases in commodity prices. That is, in the face of an oil price shock, these models concur that policy should allow a *temporary* increase in both unemployment and in inflation—to balance the harmful effects of higher oil prices on both employment and prices—*provided that* long-run inflation expectations remain well anchored. Similarly, Governor Fischer notes that the Bank of Israel’s DSGE model and their Keynesian-type model give fairly consistent results when the unemployment rate is far from the natural rate—although at other times the messages tend to differ.⁴⁰ In Sweden, moreover, where the Riksbank uses a whole set of models ranging from a state-of-the-art DSGE model to a few indicator and single-equation models, Lars Svensson reports that the board and staff practice a “kind of informal averaging” of the resulting forecasts (to the mean or median, not the mode), applying a good deal of their own judgment.

From the perspective of the European Central Bank (ECB), Jürgen Stark also advises being wary of reduced form models that short-circuit the workings of a complex economy, have no role for the money supply, and assume away shocks that originate in the money market or the financial sector. At the ECB, Stark points out that policy analysis is supported by two pillars—an economic pillar and a monetary pillar. Under the economic pillar, the ECB’s staff prepares projections of growth and inflation using a range of models, including those based on the Phillips curve. But they also look at monetary dynamics and monetary aggregates and rely on a large DSGE model with a developed credit market to reveal inflationary trends, potential financial imbalances, and the risks of financial turmoil that would not show up in models where inflation and output move only because of innovations in real activity or cost shocks.⁴¹ Stark reports that since the start of the financial tensions in August 2007, the ECB has found monetary analysis to be crucially important.

Since monetary policy actions work with a lag, these central bankers uniformly stress the need for forward-looking analysis and policy decisions—even, as Stark notes, when short-term forces threaten to distract them. The goal of monetary policy must remain to minimize the costs of fluctuations in *future* activity and *future* inflation. Or as Lars Svensson puts it, what matters for private sector decisions is less the current policy rate, and more the *expected path* of the policy rate and, thus, expectations about inflation and the real economy. As a result, the Riksbank practices “forecast targeting,” choosing and publishing a policy-rate path that produces a forecast that “looks good”—in the sense that resource use achieves the “normal” level and the inflation rate hits its target within two to three years.⁴² By comparison, Fischer reports that the Israelis give themselves just one year because Israeli inflation has been very volatile, and they have limited faith in their forecast more than a year ahead.

Agreeing on the importance of grounding inflation expectations, these policymakers tend to view their models and forecasts as communications tools. While Swedish policymakers use several models, they put particular reliance on a DSGE model with New Keynesian Phillips curve elements in the supply bloc. Because they now publish the forecasted path for the policy rate, in Svensson’s view discussion among Riksbank Board members stays oriented toward the future while their key model’s general equilibrium perspective encourages a systematic treatment of alternative assumptions. Similarly, the Riksbank publishes uncertainty intervals around its forecasts to remind the public that forecasting uncertainty abounds, and that the forecast is a forecast, not a promise.

Like a growing number of institutions, three of the four central banks represented on the conference panel practice inflation targeting and view an explicit inflation target as effective in helping to anchor inflation expectations. Elsewhere in this volume, Michael Kiley and Athanasios Orphanides provide supporting evidence regarding this proposition and suggest that an explicit inflation target may be especially useful as a communications tool in the presence of learning or rational inattention.⁴³ As a result, Orphanides concludes that “clarity regarding the central bank’s

price stability objective may improve macroeconomic performance,” even in the presence of a series of adverse supply shocks and financial disturbances. Or as Lucas noted in his famous 1973 critique, “it appears that policy makers, if they wish to forecast the response of citizens, must take the latter into their confidence.” Increasingly, central bankers are trying to do so.

The Phillips Curve Going Forward: What We Still Need to Learn about Inflation

Although our understanding of the inflation process has changed and expanded considerably over the past 50 years, many gaps, puzzles, and unanswered questions remain. In Federal Reserve Board Chairman Ben Bernanke’s view, the most pressing issues for policymakers relate to the interaction of commodity prices and inflation, the role of labor costs in setting prices, problems stemming from the need to make policy on the basis of highly uncertain real time data, and, once again, the determinants and impact of inflation expectations.

In elaborating on this list, Bernanke noted that the extraordinary and largely unexpected volatility of oil and other commodity prices in recent months underscores our need for better forecasts for this sector. Recognizing that commodity futures provide very little information about future spot prices, he encouraged additional efforts to identify the fundamental determinants of commodity prices and their structural relationships. While the traditional Phillips curve and much empirical work treat oil prices as exogenous, Bernanke pointed out that the breadth of the recent commodity price gains suggests that aggregate and sector-specific developments both play a role in determining these prices. Indeed, he wondered whether the link between global growth and commodity prices suggests a place for global—in addition to domestic—slack in the Phillips curve framework, and asked what the behavior of commodity prices can tell us about the state of the world’s economy.

Turning to the second item on his list, the role of labor costs in the inflation process, Bernanke pointed out that analysts naturally expect marginal cost (of which labor comprises a large share) to play a key

role in firms' pricing decisions; however, the empirical evidence for this link is not strong, in part, most likely, because neither labor compensation nor labor productivity is well measured. In addition, time-varying markups could be hiding the links between prices and unit labor costs. Further empirical work to clarify these relationships would be welcome.

Next Bernanke took up one of the themes running through the historical overview and the conference discussions—the difficulties of making policy decisions in real time in the face of considerable uncertainty and on the basis of indicators, like the output gap, that are hard to measure. Because economists have accumulated much evidence suggesting that economic slack does in fact affect inflation, the Chairman urged researchers to continue the search for better ways to measure the relevant gaps as well as to disentangle transitory from persistent changes in inflation. He also asked policy analysts to consider better procedures for making policy decisions when information about the state of the economy is limited and knowledge of how the economy works is incomplete.

Regarding inflation expectations, Bernanke noted that traditional models with rational expectations have no role for learning, whereas in fact the public lacks full knowledge of the state and workings of the economy and of policymakers' objectives, all of which change over time. Thus, he expressed a particular need for gaining a better understanding of how learning shapes the public's inflation expectations and how policymakers' words and actions can influence this process. Another important issue relates to how inflation expectations affect actual inflation. Is it through the wage channel or, given Blinder et al.'s (1998) puzzling finding that expected aggregate inflation plays a limited role in firms' pricing decisions, is it through a route that is less direct? Finally, while policymakers have several measures of inflation expectations, they have little information regarding the expectations of the price-setters, the firms, and little guidance on how to weight differing measures of expected inflation.

In all, Chairman Bernanke presented the economics profession with a challenging set of compelling questions. We hope that this volume proves helpful to the economists who seek to respond.

Notes

1. For example, the median of the October 2008 Consensus Forecast for the civilian unemployment rate peaked at 7.4 percent in the fourth quarter of 2009. At the end of May 2009, U.S. unemployment stood at 9.4 percent, up from 7.6 percent in January, and the May 2009 Consensus Forecast for unemployment peaked at 9.8 in early 2010.
2. Philips (1958), p. 283.
3. The canonical first-order condition for labor in a perfectly competitive environment yields a similar conclusion: the nominal wage will be set equal to the nominal marginal product of labor, or equivalently, the real wage equals the real marginal product.
4. Friedman (1968), p. 11.
5. Immediate adjustment is *not* a property of rational expectations *per se*, but it is a property in this simple model.
6. While the original work of Phillips, Samuelson and Solow, and Friedman focused on the wage-unemployment correlation, much of the subsequent literature centers on the *inflation*-unemployment link. Implicitly, this switch achieves two goals. By focusing on price inflation, it devotes attention to the variable that is of more direct relevance to monetary policy. And by switching its focus to inflation, the literature sidesteps the difficult link from wages to prices, which depends on the behavior of productivity and the markup of prices over labor costs.
7. Gordon argues that if actual prices do not drop instantly when the market-clearing price falls, firms will accumulate (presumably unwanted) inventory endlessly—as long as the assumed Lucas supply function, which defines changes in output and employment as *voluntary* responses to the gap between actual and expected inflation, is retained.
8. Gordon applied the label “triangle” model because it contains three sets of explanatory variables: a measure of excess demand, which usually takes the form of the deviation of the unemployment rate from the NAIRU or output from potential; supply shock variables, such as changes in relative oil or import prices and changes in trend productivity growth; and lags of inflation (with the restriction that the sum of the coefficients of lagged inflation equals 1).
9. Fischer points out that a particular type of indexation can cause long-term contracts to replicate the behavior of one-period contracts, reinstating the policy neutrality result. But the form of indexation that he described does not correspond to any indexing schemes observed in the economy.
10. Models could overcome this restriction mechanically by assuming serially correlated shocks to the real economy. But constructing a model that implied *endogenous* persistence of the type observed in the macroeconomic data remained an aspiration.

11. Roberts (1995) derives the isomorphism between the two specifications.

12. Calvo's formulation makes the current contract price a geometrically weighted average of future price levels, adjusted for excess demand in the future. Denoting the contract price by V_t , the price level by P_t , and excess demand by E_t ,

$$V_t = \delta \int_t^{\infty} [P_s + \beta E_s] e^{-\delta(s-t)} ds.$$

The price index is a geometrically weighted average of past contract prices,

$$P_t = \delta \int_{-\infty}^t V_s e^{-\delta(t-s)} ds.$$

The combination of these two makes prices implicitly a mixed forward- and backward-looking function of contract prices, similar to Taylor (1980):

$$V_t = \int_0^{\infty} \beta_s [V_{t-s} + V_{t+s} + E_{t+s}] ds,$$

$$\beta_s = \int_0^{\infty} f_t f_{t+s} dr,$$

$$f_s = \delta e^{-\delta s}.$$

13. See, for example, Woodford (1996, 2003).

14. Galí and Gertler point out that under certain restrictions, marginal cost and the output gap are proportional. However, these conditions are not likely to be satisfied in U.S. data, and the evidence in their 1999 paper in part demonstrates differences between the two series that are critical in modeling inflation.

15. As Galí and Gertler note in their paper, a number of previous authors were unable to develop a positive and significant coefficient in the New Keynesian Phillips curve when using a measure of the output gap as a proxy for marginal cost.

16. More recent studies have documented the possibility that inflation persistence may have declined in recent years. See, for example Benati (2008).

17. See Fuhrer (2006) for a detailed discussion of this issue.

18. Equation 9 implies another complication to this issue: if some price-setters are backward-looking, then the forward-looking price-setters will take this inertia into account in forecasting future inflation, which will act to multiply this inertial effect on inflation.

19. Note that the optimizing framework employed to derive most New Keynesian Phillips curves implies that many candidates for supply shocks should be captured in a proper measure of marginal cost, and thus should not appear as additive shocks to the New Keynesian Phillips curve.

20. Rotemberg and Woodford's formulation shifts the timing of these key equations somewhat, due to their assumptions about predetermined components of spending, which they adopt to better match the empirical properties of their benchmark vector autoregression.

21. The accompanying comment by Fuhrer (1997) provides an analysis of the extent to which the model's success is achieved through these error processes.
22. In terms of the Phillips curve specification in (8), the slope κ is a positive function of the frequency with which firms change prices.
23. Hendry and Neale (1991) show that stationary series with step changes are often mistaken for I(1) processes, a finding that exaggerates the degree of persistence in the series.
24. This specification is Gordon's (1982) triangle model of inflation. The term "triangle" refers to a Phillips curve that depends on three elements: lags of inflation (with the restriction that the sum of the coefficients of lagged inflation equals unity), a measure of excess demand which usually takes the form of a deviation of the unemployment rate from the NAIRU, and supply-shock variables.
25. See Fuhrer, Olivei, and Tootell (2009) and the literature referenced therein.
26. But, Taylor noted, there are other versions, and the work continues. Recently, for instance, economists have been looking at state-dependent pricing.
27. In Solow's interpretation, in the New Keynesian Phillips curve the output gap represents aggregate marginal cost; it is not a measure of disequilibrium, as Phillips intended.
28. As Kiley reminds us in this volume, this idea is what Lucas had in mind when he argued that expectations are rational *subject to information constraints* that leave agents with imperfect knowledge of aggregate conditions.
29. Orphanides and Williams (2005a) show that if agents learn from recent economic outcomes in forming inflation expectations, an adverse supply shock can lead to more protracted inflation and recessions than in a perfect information, rational expectations economy. Learning behavior tends to impart additional persistence to inflation and complicates modeling efforts.
30. They focus particularly on Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008a and 2008b), and Dhyne et al. (2005) for Europe along with Alvarez (2008) and Alvarez and Hernando (2007), Blinder et al. (1998), and Zbaracki et al. (2004) for survey data.
31. Taylor's chapter for the *Handbook of Macroeconomics* (1999) reports that micro-level prices do not change more often than wages, that price and wage setting behavior show much heterogeneity, that neither price nor wage setting is synchronized, and that the frequency of price and wage changes is positively related to the pace of inflation.
32. In the United States the median consumer price lasts four to nine months (depending on whether sales prices and forced item substitutions are excluded) versus 11 months in the euro area.
33. Boivin, Giannoni and Mihov (2009) draw similar conclusions. In new work for their Boston Fed conference paper included in this volume, Maćkowiak and Smets confirm that the frequency of price change within a sector helps to explain the speed of impulse responses of prices to macroeconomic shocks across sectors.

This work, based on McCallum and Smets (2008), uses factor-augmented VAR methods from Bernanke, Boivin, and Eliasziw (2005).

34. King makes a similar point about the Calvo model. Although the Calvo model has many advantages (e.g., it can deliver nominal prices that are constant for uneven periods of time) as parameterized in the mid-1990s with 10 percent of firms adjusting prices each quarter, the average price was assumed to be sticky for 10 quarters.

35. Using a DSGE model from Smets and Wouters (2003), Maćkowiak and Smets draw hints from the importance of backward-looking elements that some form of imperfect information about macro shocks “matters” for macro dynamics; they emphasize that “the fact that prices change does not imply that prices reflect perfectly “all available information.”

36. To Maćkowiak and Smets’s list of appealing ways to span the gap between micro price flexibility and aggregate inertia (i.e., real rigidities and information frictions), Midrigan suggests adding a third: inventory-based models of money demand as in Alvarez, Atkeson, and Edmond (2008). Since real rigidities include those that reflect the slow response of aggregate marginal cost to fluctuations in output and since measuring the behavior of real marginal cost over the cycle is hard, Bils and Khan (2000) and Kryvtsov and Midrigan (2009) suggest that economists can learn a lot about the behavior of marginal cost—and the size of the related rigidities—from the cyclical behavior of inventories.

37. Indeed, he says that “alternative frameworks seem to lack solid economic foundations and empirical support” (p. 415 in this volume).

38. Governor Fischer explains that Israel’s inflationary history has resulted in many contracts, including those for rental housing, being denominated in dollars. Until recently, thus, the close link between the exchange rate and the price level (with an immediate pass-through of about one-third in a quarter) meant that monetary policy tended to work very fast because it affected the exchange rate. However, the recent strength of the shekel has led to a rapid decline in the share of contracts denominated in dollars and a disorienting change in the Israeli monetary transmission mechanism.

39. See Lucas (1976). In the context of rapidly rising oil prices, Kohn notes that Woodford (1994) uses the Lucas critique to argue that the tendency of commodity prices to forecast inflation may not be structural and could disappear under different regimes.

40. In June 2008, the Bank of Israel’s DSGE model (which uses a Hodrick-Prescott filter to measure the gap) showed the Israeli economy fluctuating around full employment while the Keynesian-style model was suggesting that the economy had been above full-employment for some time. Since the most recent price data had just revealed a surprising surge in inflation in almost every price group, Governor Fischer was ready to conclude that strong demand had paved the way for commodity price pressures to spread and that the Phillips curve was alive and well.

41. Sims, in writing about inflation expectations, rational inattention, and monetary policy for this book, also argues that it is important, though hard, to model the interaction of asset markets and monetary policy.
42. The Riksbank has published a forecasted policy-rate path since February 2007 and, as Svensson points out, is the first central bank with an “individualistic” policy board to do so. From Svensson’s perspective, reaching agreement on this forecasted path has proved easier than expected—in large part because extensive interactions with board members allow staff to identify the path and forecast that a majority of the Board members are likely to prefer.
43. Kiley (this volume) cites evidence that inflation expectations and inflation compensation appear more stable in countries with an explicit inflation target.

References

- Akerlof, George A., William T. Dickens, and George L. Perry. 2000. “Near-Rational Wage and Price Setting and the Long-Run Phillips Curve.” *Brookings Papers on Economic Activity* 1: 1–44.
- Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde. 2005. “Firm-Specific Capital, Nominal Rigidities and the Business Cycle.” Working Paper 2005-01. Chicago: Federal Reserve Bank of Chicago. Available at http://www.chicagofed.org/publications/workingpapers/wp2005_01.pdf.
- Alvarez, Fernando, Andrew Atkeson, and Chris Edmond. Forthcoming. “Sluggish Responses of Prices and Inflation to Monetary Shocks in an Inventory Model of Money Demand.” *Quarterly Journal of Economics*.
- Álvarez, Luis J., and Ignacio Hernando. 2007. “Competition and Price Adjustment in the Euro Area.” In *Pricing Decisions in the Euro Area: How Firms Set Prices and Why*, ed. Silvia Fabiani, Claire Loupias, Fernando Martins, and Roberto Sabbatini, 200–217. Oxford: Oxford University Press.
- Álvarez, Luis J. 2008. “What Do Micro Price Data Tell Us About the Validity of the New Keynesian Phillips Curve?” *Economics: The Open-Access, Open-Assessment E-Journal* 2(19). Available at <http://www.economics-ejournal.org/economics/journalarticles/2008-19>.
- Atkeson, Andrew, and Lee E. Ohanian. 2001. “Are Phillips Curves Useful for Forecasting Inflation?” *Federal Reserve Bank of Minneapolis Quarterly Review* 25(1): 2–11. Available at <http://www.minneapolisfed.org/research/QR/QR2511.pdf>.
- Ball, Laurence, and N. Gregory Mankiw. 2002. “The NAIRU in Theory and Practice.” *Journal of Economic Perspectives* 16(4): 115–136.
- Basu, Susanto. 1995. “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare.” *American Economic Review* 85(3): 512–531.
- Benati, Luca. 2008. “Investigating Inflation Persistence Across Monetary Regimes.” *Quarterly Journal of Economics*, 123(3): 1005–1060.

Bernanke, Ben S., Jean Boivin, and Piotr Elias. 2005. "Measuring Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120(1): 387–422.

Bils, Mark, and James A. Khan. 2000. "What Inventory Behavior Tells Us About Business Cycles." *American Economic Review* 90(3): 458–481.

Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.

Blanchard, Olivier Jean, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 2: 1–60.

Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual*, ed. Stanley Fischer, 15–90. Cambridge, MA: The MIT Press.

Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.

Boivin, Jean, Marc P. Giannoni, and Ilian Mihov. 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated U.S. Data." *American Economic Review* 99(1): 350–384.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–398.

Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102(4): 703–726.

Carlton, Dennis W. 1986. "The Rigidity of Prices." *American Economic Review* 74(4): 637–658.

Cecchetti, Stephen G. 1986. "The Frequency of Price Adjustment: A Study of Newsstand Prices of Magazines." *Journal of Econometrics* 31(3): 255–274.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.

Cogley, Timothy, and Argia M. Sbordone. 2008. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review* 98(2): 2101–2126.

Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünemann, Fabio Rünstler, and Jouko Vilmunen. 2005. "Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data." Working Paper 524. Frankfurt: European Central Bank. Available at <http://www.ecb.int/pub/pdf/scpwps/ecbwp524.pdf>.

Dupor, Bill, Tomiyuki Kitamura, and Takayuki Tsuruga. 2008. "Integrating Sticky Prices and Sticky Information." Manuscript. Available at http://web.econ.ohio-state.edu/dupor/sipc12_Sep17.pdf.

- Fischer, Stanley. 1977. "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule." *Journal of Political Economy* 85(1): 191–205.
- Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review* 58(1): 1–17.
- Fuhrer, Jeffrey. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy: Comment." In *NBER Macroeconomics Annual 1997*, ed. Ben S. Bernanke and Julio Rotemberg, 346–355. Cambridge, MA: The MIT Press.
- Fuhrer, Jeffrey. 2006. "Intrinsic and Inherited Inflation Persistence." *International Journal of Central Banking* 2(3): 49–86.
- Fuhrer, Jeffrey and George Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics*, 110(February): 127–159.
- Fuhrer, Jeff, Giovanni Olivei, and Geoffrey M.B. Tootell. 2009. "Empirical Estimates of Changing Inflation Dynamics." Manuscript. Boston: Federal Reserve Bank of Boston. Available upon request.
- Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44: 195–222.
- Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115(2): 171–199.
- Gordon, Robert J. 1977. "The Theory of Domestic Inflation." *American Economic Review* 67(1): 128–134.
- Gordon, Robert J. 1982. "Price Inertia and Policy Ineffectiveness in the United States, 1890–1980." *Journal of Political Economy* 90(6): 1087–1117.
- Gordon, Robert J. 2008. "The History of the Phillips Curve: An American Perspective." Paper presented at the Northwestern Macro Workshop, September 29. Available at http://faculty-web.at.northwestern.edu/economics/gordon/NUWkshop_Rev_080926.pdf.
- Gordon, Robert J., and Stephen R. King. 1982. "The Output Cost of Disinflation in Traditional and Vector Autoregressive Models." *Brookings Papers on Economic Activity* 1: 205–242.
- Gray, Jo Anna. 1977. "Wage Indexation: A Macroeconomic Approach." *Journal of Monetary Economics* 2(2): 221–235.
- Hendry, David F., and Adrian J. Neale. 1991. "A Monte Carlo Study of the Effects of Structural Breaks on Tests for Unit Roots." In *Economic Structural Change: Analysis and Forecasting*, ed. Peter Hackl and Anders H. Westlund, 95–119. New York: Springer-Verlag.
- Kashyap, Anil K. 1995. "Sticky Prices: New Evidence from Retail Catalogues." *Quarterly Journal of Economics* 110: 245–274.
- Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit and Banking* 39(1, Supplement): 101–125.

- Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter For Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(4): 863–904.
- Kozicki, Sharon, and Peter A. Tinsley. 2002. "Alternative Sources of the Lag Dynamics of Inflation." In *Price Adjustment and Monetary Policy*, 3–47. Ottawa: Bank of Canada. Available at http://www.bankofcanada.ca/en/conference/2002/Kozicki_Tinsley-final.pdf.
- Kryvtsov, Oleksiy, and Virgiliu Midrigan. 2009. "Inventories, Markups, and Real Rigidities in Menu Cost Models." Working Paper 14651. Cambridge, MA: National Bureau of Economic Research.
- Lucas, Robert E., Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63(3): 326–334.
- Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1: 9–46.
- Maćkowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt. 2008. "Sectoral Price Data and Models of Price Setting." Mimeo. Available at <http://www2.wiwi.hu-berlin.de/bartosz/MMW081205.pdf>.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99(3): 769–803.
- Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117(4): 1295–1328.
- McCallum, Andrew, and Frank Smets. 2008. "Real Wages and Monetary Policy Transmission in the Euro Area." Mimeo. European Central Bank.
- Muth, John. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29(3): 315–335.
- Nakamura, Emi, and Jón Steinsson. 2008a. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415–1464.
- Nakamura, Emi, and Jón Steinsson. 2008b. "More Facts about Prices—Supplement to 'Five Facts about Prices.'" Available at <http://www.columbia.edu/~en2198/papers/fivefactssupplement.pdf>.
- Orphanides, Athanasios. 2003. "The Quest for Prosperity without Inflation." *Journal of Monetary Economics* 50(3): 633–663.
- Orphanides, Athanasios, and John C. Williams. 2005a. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation Targeting Debate*, ed. Ben S. Bernanke and Michael Woodford, 201–246. Chicago: University of Chicago Press.
- Orphanides, Athanasios, and John C. Williams. 2005b. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations." *Journal of Economic Dynamics and Control* 29(11): 1927–1950.

Paciello, Luigi. 2008. "Price Responsiveness to Aggregate Technology and Monetary Policy Shocks under Rational Inattention." Working Paper 08-16. Rome: Einaudi Institute for Economics and Finance. Available at <http://gigipaciello.googlepages.com/JMPpacielloOct08.pdf>.

Phelps, Edmund S. 1968. "Money-Wage Dynamics and Labor-Market Equilibrium." *Journal of Political Economy* 76(4, Part 2): 678–711.

Phelps, Edmund S., and John B. Taylor. 1977. "Stabilizing Powers of Monetary Policy Under Rational Expectations." *Journal of Political Economy* 85(1): 163–190.

Phillips, A.W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica* 25(100): 283–299.

Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policymakers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121(3): 867–901.

Reis, Ricardo. 2006. "Inattentive Producers." *Review of Economic Studies* 73(3): 793–821.

Roberts, John. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27(4, Part 1): 975–984.

Romer, Christina D., and David H. Romer. 2002. "The Evolution of Economic Understanding and Postwar Stabilization Policy." In *Rethinking Stabilization Policy*, 11–78. Kansas City, MO: Federal Reserve Bank of Kansas City. Available at <http://www.kansascityfed.org/publicat/sympos/2002/pdf/S02RomerandRomer.pdf>.

Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90(6): 1187–1211.

Rotemberg, Julio J. 1983. "Aggregate Consequences of Fixed Costs of Price Adjustment." *American Economic Review* 73(3): 433–436.

Rotemberg, Julio J., and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." In *NBER Macroeconomics Annual*, ed. Ben S. Bernanke and Julio Rotemberg, 297–346. Cambridge, MA: The MIT Press.

Rudd, Jeremy, and Karl Whelan. 2006. "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" *American Economic Review* 96 (1): 303–320.

Rudd, Jeremy, and Karl Whelan. 2007. "Modeling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit and Banking* 39(1, Supplement): 155–170.

Samuelson, Paul A., and Robert M. Solow. 1960. "Analytical Aspects of Anti-Inflation Policy." *American Economic Review* 50(2): 177–194.

Sargent, Thomas J., and Neil Wallace. 1975. "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *Journal of Political Economy* 83(2): 241–254.

- Sargent, Thomas, Noah Williams, and Tao Zha. 2006. "Shocks and Government Beliefs: The Rise and Fall of American Inflation." *American Economic Review* 96 (4): 1193–1224.
- Sbordone, Argia M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49(2): 265–292.
- Sims, Christopher A. 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49 (December): 317–356.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665–690.
- Sims, Christopher A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96(2): 158–163.
- Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1(5): 1123–1175.
- Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.
- Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment, and Monetary Policy." *Journal of Economic Perspectives* 11(1): 33–51.
- Stock, James H., and Mark W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39(1): 3–34.
- Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88(1): 1–23.
- Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39 (December): 195–214.
- Taylor, John B. 1999. "Staggered Price and Wage Setting in Macroeconomics." In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford, 1009–1050. New York: Elsevier.
- Woodford, Michael. 1994. "Nonstandard Indicators for Monetary Policy: Can Their Usefulness Be Judged From Forecasting Regressions?" In *Monetary Policy*, ed. N. Gregory Mankiw, 95–116. Chicago: University of Chicago Press.
- Woodford, Michael. 1996. "Control of the Public Debt: A Requirement for Price Stability?" Working Paper 5684. Cambridge, MA: National Bureau of Economic Research.
- Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.
- Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics* 86(2): 514–533.

2

**Fifty Years of the Phillips Curve:
A Dialog on What We Have Learned**

Fifty Years of the Phillips Curve: A Dialog on What We Have Learned

Robert M. Solow and John B. Taylor
N. Gregory Mankiw, Moderator

Greg Mankiw: It's a great honor to be here to moderate this dialog between two of the most important macroeconomists of the past half century. It is very intimidating for me. But I've come to realize that all the hours I've spent watching TV talk shows has now finally produced some useful human capital for me. I've been trying to figure out who to emulate. I don't think I'm urbane enough to be Dick Cavett. I don't think I'm funny enough to be Johnny Carson. I'm going to aim for Jerry Springer. So, Bob and John, if you want to throw a chair at each other, just let me know and I'll duck.

The topic for today is the Phillips curve. I remember thinking a lot about the Phillips curve as a student. I remember thinking at the time that this was an incredibly important macroeconomic relationship. It almost defined macroeconomics and explained why classical economic principles didn't exactly apply in the short run.

The Phillips curve also made no sense to me as a student. I remember being very puzzled by it. I thought this was a great topic to do research on. I thought that if I start working on it, maybe, I'll figure out what's really behind this concept. I think I've done that for 20 years now. I still find it very frustrating and puzzling. I'm glad that we are all here today to finally settle the issue.

Let me start with a little history. Bob Solow was one of the people to bring the Phillips curve to the United States. With Paul Samuelson he wrote one of the great macroeconomic papers, a classic article titled "Analytical Aspects of Anti-Inflation Policy" and published in the *American Economic Review* in 1960, that applied the Phillips curve idea. So, I'm curious to first hear from Bob what he thought about the Phillips

curve when he first saw it and what it was like then. Did you realize that we'd still be talking about it 50 years later?

Robert Solow: Before I get to answer that question directly, I would like to take a minute or two to say a word about Bill Phillips. This is a Phillips curve retrospective, after all. Am I the only person in the room who knew Phillips reasonably well, personally? One? Good; I'm not alone then.

Well, let me just remind everybody that Phillips was born and bred in New Zealand. He loved the place. When he spent the 1962–1963 academic year at the Massachusetts Institute of Technology, he was a neighbor of ours, as he lived close by in Concord, Massachusetts. Naturally I spent some time with him, including several evenings just poring over picture books and photographs of the New Zealand landscape. Phillips simply loved it; he was truly nostalgic about New Zealand. I can only confirm that it was very beautiful.

He was trained as an electrical engineer, and eventually wound up in the British armed forces as a technical officer with a unit of the Royal Air Force. He was taken prisoner by the Japanese in Indonesia, having escaped to that place during the Japanese conquest of Singapore. He then spent several years in a Japanese prison camp. He was not abused, as far as I know, but it cannot have been nice. According to all reports, Bill Phillips behaved extraordinarily well, was an important person to his co-prisoners. At great risk he jiggled up radios out of odd pieces of wire and stolen parts so that they could know what was going on in the world. Then, after the war, he found his way to the London School of Economics (LSE), thought about studying sociology, thought better of it, and turned to economics. The rest, I hope, you know.

I learned about the Phillips curve the old-fashioned way. In those days, in 1958, you went to the library and looked at journals when they came out. There were only six or seven journals that an English speaker would be looking at; *Economica*, published at the LSE, was one of them. When I picked up the November 1958 issue to look at the table of contents, presumably in early 1959, here was this article on the relation between unemployment and the rate of change of money wages in the United Kingdom. I read it out of plain curiosity and I thought it was remarkable,

because of this amazingly stable empirical relationship that he found. I took it out of the library and showed it to Paul Samuelson.

Now, it happens that we had already agreed to give a joint talk at the American Economic Association (AEA) meetings in December 1959—indeed, a talk about the analytical foundations of anti-inflation policy. You have to think about the setting. In 1958 and 1959 everyone was thinking about what was then called “creeping inflation.” During the recessions of 1954 and 1958 the price level in the United States had continued to rise slowly at a rate that would be lost in the noise today, a mere 1 or 2 percent a year. The big debate was whether this inflation could be explained by “demand-pull” or “cost-push.” Paul and I began to think about that argument and the various empirical tests that had been proposed as a way of distinguishing between them. We concluded that once you had a glimmer of general equilibrium thinking in your head, all those tests were wrong, and could not make the distinction that was wanted. We finally decided that a much more sensible idea to bring to bear was the Phillips curve. We thought that a more useful distinction was between movements along the Phillips curve and shifts of the Phillips curve. So that is what we did.

Greg Mankiw: Did you coin the term, “the Phillips curve,” in that paper?

Robert Solow: I believe that Paul and I coined the phrase, “the Phillips curve.” Obviously, this phrase doesn’t appear in Bill Phillips’s article—I think we invented the name. He invented the curve.

Greg Mankiw: There is one footnote on the New Zealand fact. In the first edition of my macroeconomics textbook, I referred to Phillips as a British economist because he was from the LSE, and I got quite a few annoyed letters from New Zealand. So that errata got quickly corrected in the second edition.

Now, John, you started grad school in the late 1960s. Can you give us some sense of what the academic thinking on the Phillips curve looked like at the time? That was also the time when the Phillips curve was still looking okay, but some people were starting to worry about it. Can you give us a sense of the intellectual climate back then?

John Taylor: I'd also like to preface my remarks, like Bob did, with some broader statements about Phillips. While I never met him personally, I felt like I knew him personally because I benefited from so much of his work outside of the Phillips curve early in my life. I graduated from college in 1968, the academic year that Milton Friedman gave his AEA presidential address in December 1967. So, in some sense I'm post-Friedman here, but my senior thesis in college used a macroeconomic model based on Phillips's work. The model had a Phillips curve in it, of course. But it also had an investment equation and a money demand equation. Why was I doing that? I was simulating different types of policy rules to see how they would work in this model. Actually, in this respect I was following very much the line of research that Phillips established long before based on his training in engineering. There were these various methods to control dynamic systems that came out of the engineering literature—thermostatic policies, actually. You had proportional control, derivative control, and integral control. Phillips had written some papers on these control policies, and what I wanted to do was look at those policy questions in a model that involved both cycles and economic growth, which was in a paper Phillips published in *Economica* in 1961.

That model is related to your question because, as I say, a Phillips curve is in the model. It is a macroeconomic model, a little dynamic model, with differential equations. But one of the equations was a Phillips curve in the sense that it related inflation to the output gap. Rather than the unemployment rate, he had the output gap, just like we frequently do in modern times. And he had potential output, natural output, moving around according to a simple production function. It was sophisticated, quite frankly. But the Phillips curve wasn't an *expectations-augmented* Phillips curve. It was the kind of Phillips curve that Bob just was referring to. What I remember so much is that in my work on this at the time, for some reason, I didn't think for a minute of trying to exploit the fact that there was a long-run trade-off. In other words, all the policy simulations were concerned with stabilizing GDP around the natural or potential rate. For some reason, I never exploited the long-run features, which I think to some extent reflected Phillips's own intuition here. He didn't think of the curve as something that should be exploited that way. Even though, if you took the algebra literally, without the expectations term,

you'd argue it should be exploited by bringing GDP above potential and having a higher inflation rate. So, anyway, I didn't exploit it. I don't think he ever did either, but we could debate that.

With respect to the developments at the time, as I said, Milton's presidential address was in December 1967. I think that it changed thinking pretty rapidly. By the time I got to graduate school, that was the way we were thinking about things. It was an expectations-augmented curve. Of course you had Ned Phelps's important parallel work, and the Phelps volume of research following on very quickly. My recollection is that we all knew about the curve shifting over time and there was no long-run trade-off between inflation and unemployment. The change in academic thinking came very rapidly.

What was most striking to me in a policy context, I guess, is how this change didn't seem to influence policy very rapidly. Here is the reason I think there was a delay. People had been thinking in the 1960s about the long-run trade-off and that we could get more—let's say, less unemployment—with a higher inflation rate. Of course, by the late 1960s, early 1970s, mid-1970s, we were getting the high inflation rate but we weren't getting the low unemployment rate—so the Friedman-Phelps critique was validated. What happened with policy, however, is that actions were not taken instantly to bring the inflation rate down, even though the economy wasn't benefiting with lower unemployment. The new policy dilemma was that it was going to be costly to *reduce* the inflation rate. The expectations-augmented Phillips curve was put into the model, as Milton suggested, but with adaptive expectations—slowly adaptive expectations. Suddenly the policy trade-off about the level of inflation became a policy trade-off about the cost of disinflation. I think that “cost of disinflation” concept influenced economists all through the 1970s. It suddenly became fashionable to think that we can't get rid of inflation because it is going to be too costly. Those were the days that you had President Gerald Ford with the “Whip Inflation Now” buttons and his speech to a joint session of Congress. “How to Whip Inflation Now” was the theme at a White House conference, I think, in 1975. If you look through the transcript of that meeting, the only one at the time who was aggressive in terms of inflation reduction was Milton. So, it is a very interesting development, the interaction of the theory and the policy.

Greg Mankiw: John raised Milton Friedman's classic 1967 AEA presidential address. Bob, I was wondering what your reaction was to that paper? Did you think that address provided a fair assessment of what had come before? Or did you think it was a caricature of things that were said before? I went back some years ago and reread Samuelson and Solow and a lot of the caveats regarding the Phillips curve that we now talk about were in that 1960 article published in the *American Economic Review*. We just talked about expectations. In the paper you also talked about effects which today we call hysteresis.

See, a lot of the caveats were in that paper. I was wondering how you reacted to Milton Friedman's presidential address.

Robert Solow: It is true that our 1960 paper made all those allowances. We said explicitly that it is unlikely that one could successfully exploit the Phillips curve in the long run. We even mentioned the possibility that it was inflationary expectations that would shift the curve adversely if one tried. But I think we had something more general in mind: that the mere experience, however you process it, whether through expectations or the development of norms or behavior, would have that effect. So when I read Milton's address, that part didn't come as much of a surprise, though Milton dwelt on that point much more than we had thought to do so.

What did come as a surprise, and still comes as a surprise, though I didn't realize it right away, was that Milton had done something much more subtle and important, without explicitly saying so at all.

Let me go back to Phillips for a moment. Phillips's 1958 paper is purely empirical. All the theory in it is contained in the first two sentences of the paper, and what they say this: we are all used to the idea that excess demand in a market will cause the price to rise, and excess supply in a market will cause the price to fall. So why shouldn't the same idea apply to the labor market? That's all the theory there is. Evidently in Phillips's mind the arrow of causality clearly runs from the unemployment rate to the rate of inflation. He is thinking of the unemployment rate as a measure of the supply-demand balance in the labor market, an indicator of disequilibrium, and it would push the rate of wage inflation in the same way that you would expect excess supply or demand for peanut butter to

push the price of peanut butter. (I don't remember any explicit discussion about why the rate of inflation rather than the price level.)

We, as sophisticated as we are, realize that all of this is slightly dangerous talk because those are both endogenous variables in whatever model we have in the back of our mind. So saying that A causes B rather than B causing A is tricky. On the other hand, in our less-sophisticated moments we know damn well what Phillips meant. He meant that the causality ran primarily from disequilibrium in the labor market to wages or wage inflation. What Milton did without ringing any bells to warn you, was simply to take it that the causality ran the other way, that it's the deviation of the rate of inflation from the expected rate of inflation that pushes the unemployment rate away from the "natural" rate. Phillips is about disequilibrium in the labor market. There is no question about that, just from the first couple of sentences from the paper that I paraphrased from memory. After Milton's address, everybody treated this as an equilibrium matter, looking in the reverse direction.

In this new story, the only way you can get the unemployment rate to depart from the "natural" rate is to create an inflation rate different from the expected rate. That kind of causality puzzled me, when I realized what had been done. Why do I say it "puzzled" me? What I mean is that I don't believe a word of it, and I find it strange that anyone does. When we come to talk about the so-called New Keynesian Phillips curve, we may find an intellectual advance, but that will arise in due course. Anyway, slowly, adaptively, I gradually realized that what Milton had done was to change the Phillips curve from what Phillips meant to this altogether contradictory kind of arrangement.

Greg Mankiw: When I teach the history of the Phillips curve, I start off with Phillips and then go on to the modern attempts to try to understand this trade-off between inflation and unemployment. One of the very important things that we go through in class is the Taylor model, which then evolves into the Calvo model and then it evolves into the New Keynesian Phillips curve. So, I was wondering if you would each say what you think about the current state of this work. I think of the New Keynesian Phillips curve as it comes down from the Taylor and Calvo versions as being the canonical model in the literature. This is where we

are, perhaps with some bells and whistles attached, like adding lags of inflation, or indexation, or something. Do you think we are doing well, meaning heading in the right direction, or did we make a wrong turn somewhere along the way?

John Taylor: I think the research on price and wage adjustment has been very helpful, but actually to answer your question I would go back to the period well after Milton's critique, after the introduction of rational expectations, to talk about this phenomenon because Milton's model was all adaptive expectations, right? That adaptive expectations assumption is where the costly trade-off issue came from. After Milton's work you have the introduction of rational expectations by Bob Lucas, and of course Bob was using equilibrium models where the surprise change in inflation or in the price level generated deviations of unemployment from the natural rate or of GDP from trend.

Of course the implication of that rational expectation assumption, which Tom Sargent, Neil Wallace, and others emphasized, was that monetary policy was not going to be effective unless it tried to surprise. If it is at all predictable—which we thought it should be—then it wasn't going to have any impact. So, there was a real puzzle here, or a real vacuum or void, if you like, in the literature. It seemed to me that it had to be addressed and a lot of people started looking at it—Stanley Fischer, Phelps, Jo Anna Gray, myself. Originally this research focused on trying to build in a kind of price adjustment equation with rational expectations and sticky prices, and that is where I start to answer your question of whether we are on the right track. Those original “sticky price” models really lacked a lot of persistence. The models just basically jump back to equilibrium after the one period during which the prices were set or wages were set. The models didn't look that much different from what we call the new classical model. So, you had to do something else to get an empirically realistic model, and I'd like to explain the way I thought about it. This may be too long of a story, but I'll try to be quick.

We began to look at what was causing the persistence of inflation or price change. First we came up with the ideas that prices don't change every instant or every quarter or even every half-quarter, but they last for a while. They last for more than one period in our discrete time models, two periods rather than one, three periods rather than one.

But as soon as you do that, you begin to recognize that there is going to be some built-in inertia because not everybody is setting prices at the same time. You have this staggering nonsynchronization of price setting or wage setting, whatever it happens to be. So, I think that is the nonsynchronization phenomenon that people tried to think about: how do we model it? My own perspective is that it is a very hard problem to model. To me, it required a much different type of economic theory than we are used to teaching our students. So to get my hands around that, what I did was make a simple assumption that wages last for two periods, and half the workers change their wage every other period. Then there was the Calvo version of that, and now there are many others. But I think that approach generated—and this is why I think this literature is important—some important implications for policy.

Sometimes I think the value of a theory should be measured by what we get out of it. So, let me just mention what I think are five or six things you get out of that theory, which is continuing to evolve. First of all, you get some simple equations. When you are sitting in an office for a while, staring at the ceiling and trying to model some macroeconomic phenomenon, coming up with an equation is a big deal. You can estimate it and analyze it and compare different time periods, so the first thing is getting an equation. Second, you get some expectations of the *future* inflation mattering for the first time. In the Friedman-Phillips model, the augmented expectations were not expectations of the future. Similarly, in the “Lucas surprise” model, expectations are contemporaneous, not of the future. Now, for the first time, expectations of the future matter for prices today because of that staggering in the price setting. When they are setting their prices today, firms have to look to the future, and so for the first time you are getting expectations of the future mattering for inflation. I think that has a lot to do with the rationale for inflation targeting. Third, you get inertia. Inertia is built in because firms have to look back at past decisions, so you get a phenomenon that inertia lasts longer than the length of the longest contract. Fourth, you get a prediction that economic policy will affect the inertia. The more aggressive policy is, the less inertia or the lower inflation persistence will be. Fifth, you get a trade-off not between the level of inflation and the level of unemployment, but between the volatility of those two things, which has been useful. Besides the Calvo model, Jeff Fuhrer in his work with George Moore developed a

model to explain persistence more fully. But the results and the things you take away from the Fuhrer-Moore model are very similar. More recently, people are looking at models where you have state-dependent pricing rather than just time-dependent pricing, and those models are more general and interesting, but what you are getting out of them is very similar. So, I think it is a very powerful way to think about the economy and, if you believe what I say, it has had a huge influence. So, I'll leave it at that and maybe Bob and I can talk a little more.

Greg Mankiw: You point out a couple of features of these classic models that some people view as bugs. For example, the fact that expectations of future inflation matter has been turned on its head. Larry Ball says well, gee, in these models, therefore a fully credible disinflation can cause an output boom rather than a recession. It is the same feature of the model that creates persistence problems. The model creates persistence in the price level, but you don't get persistence in the inflation rate, which is very different. I think when people look at the data, you have persistence in the inflation rate as well. That is why people go to alternative models like the Fuhrer-Moore model, and so on.

So, I think the literature is still a little inconclusive—I'm trying to figure out which of these features are features and which of them are bugs.

John Taylor: I will try to include all of those in my comment. It seems to me the main message is what I just said, but yes, there are little differences one needs to worry about. First is a notion that somehow you could announce a perfect path for the money supply (usually it is the money supply, not the interest rate) that slows down gradually and maybe by two years later reaches a new level. And you assume everybody believes that. Then, yes, you can get a very costless disinflation. I think that is not how you use a model like that. That would be naïve. So, to me that is not a disadvantage of the model. Second, if you try to fit models without some sort of exogenous persistence built into them, I think they fit pretty well, quite frankly. In recent years, Luca Guerrieri at the Federal Reserve Board has shown that they fit quite well. So, I think it's not a slam dunk to show that these specific things are a problem with the models.

Greg Mankiw: Bob, do you have any comments?

Robert Solow: I don't have much to add to that. It still bothers me, I guess, that the standard equation for a New Keynesian Phillips curve has the property that if inflation is and is expected to be constant, then output settles at the natural rate. This happens immediately in the pure rational expectations case, and then only white noise deviations can occur, or it happens gradually if there is some inertia built in. That does not strike me as plausible, not remotely.

What I like about the New Keynesian Phillips curve is that it gets rid of this elementary causality business. By the way, a careful person, like Jordi Galí for instance, never says that this is a Phillips curve; he says that it looks like a Phillips curve. That is exactly the point: it looks like a Phillips curve, but it isn't. It is something very different. The output gap, as I understand it, enters only as a proxy for economy-wide marginal cost, so it isn't intended to do the disequilibrium work Phillips intended. The main point, however, is that a careful person like Jordi Galí (or John Taylor) always embeds this thing in a model. There is something like an IS curve. There is a Taylor rule equation instead of an LM curve. And then there is this New Keynesian Phillips curve, if we continue to call it that. Then at least you can talk rigorously about causality, which you couldn't otherwise.

Whether or not the current version of the New Keynesian Phillips curve can recreate the right amount of persistence appears to be doubtful. I don't keep up with the literature as closely as I would if I would 30 years younger, but the most recent item I happened to see (in the Federal Reserve Bank of Richmond's *Quarterly Review*) came to the conclusion that the standard version of the model, even with some exogenously imposed backward-looking effects, couldn't account for the degree of persistence that is found in the data. This may be inherent in stories that want to build so much around rational expectations. My old codger feeling is that the technology has made it so easy to do calculations and then write papers that we euphorically forget how hard it is to get time series properties right. It is too easy to go on to the next paper.

John Taylor: Could I just comment on that three-equation model? It usually has a very simple equation which looks like a Phillips curve. That equation has a coefficient (beta) times the expectation of the future infla-

tion rate on the right-hand side. It seems to me that those models can't fit data. They shouldn't be meant to fit data, they are just little models we use to teach students and discuss theorems. But you have to have more than that for policy on empirical work. In some sense, I think I am sympathizing with Bob here about fitting these models to the data. I once built a complex multicountry model using these staggered wage-setting models. I had to have a different structure for Japan and a different structure for Europe. Sometimes, the contracts had to last two years. Sometimes they have different distributions because you need to do that to fit the data. How could you possibly think that one little price adjustment equation could fit all different sorts of data? It is useful for expository purposes and debating, but we shouldn't hold it to a tough empirical standard. Although, probably people do to some degree.

Greg Mankiw: Bob raised the question of the natural rate hypothesis, and he said that he was skeptical that the long-run unemployment rate was invariant to whatever inflation rate the Fed decides to peg. Do you have a view as to what the current state of play is on that question?

John Taylor: Yes, I'm very much of the view that the natural rate of employment, or whatever we want to call it, is invariant to monetary policy. Let's put it this way. The idea of the classical dichotomy is very useful. I like that principle. Obviously, it's an approximation, but it seems to me that it is an important idea for describing long-term economic growth. You have the Solow model. You have trend productivity growth that is separate from monetary policy. You are going to have deviations from that, and monetary policy is important for minimizing the deviations. Similarly with the real interest rate; similarly with the natural rate of unemployment. Those are things that are from the real economy, and the more we can think about it that way, and teach and convince people, the better, because otherwise it is so confusing. When you read about monetary policy in the press, it always gets it mixed up. The more we can emphasize the invariance of the natural rate to monetary policy, the better. I think it is a good approximation. I don't have any problems with it.

Robert Solow: Everything John says is sensible; this is an experience I have had before. Let me try to explain what nags at me in all this. I'll say

it first generally, and then come to something particular. We are left here with a theory whose two central concepts, the natural rate of unemployment or output and the expected rate of inflation, have three suspicious characteristics in common. They are not directly observable. They are not very well defined. And so far as we can tell, they move around too much for comfort—they are not stable.

I suspect this is an intrinsic difficulty. I have no wish to minimize the importance of, say, inflationary expectations. But we are faced with a real problem: here is a concept that seems in our minds to play an important role in macro behavior, and yet it's very difficult to deal with because it escapes observation and it even escapes clear definition. On the natural rate of unemployment, I think the behavior of the profession exhibits problems. In order to make sensible use of this kind of theory, you want the natural rate of unemployment to be a fairly stable quantity. It won't do its job if it jumps around violently from one year to the next. But that's what seems to happen. We, the profession, are driven to explaining events by inventing movements of the natural rate, which we have not observed and have not very well defined. The issue came up first in the passage of the big European economies from 2 percent unemployment, on average, to 8 or 9 percent unemployment, on average, within a few years. The only way to explain that within the standard model is to say that the natural rate of unemployment must have increased from something like 2 percent to something like 8 or 9 percent. The actual facts that could account for any such dynamics never seemed to me or to any critical person to be capable of explaining so big a change. So we are left with inventing changes in the natural rate of unemployment to explain the facts, and it is all done in our heads, not in any tested model. I regret to say that you often find this kind of reasoning: the inflation rate is increasing because the unemployment rate is below the natural rate. How do you know that the unemployment rate is below the natural rate? Because the inflation rate is increasing. I think we are all good enough logicians to realize that this is exactly equivalent to saying that the rate of inflation is increasing, and nothing more.

It seems to me that we ought to be thinking much more about the determinants of whatever you choose to call it. I hate to use the phrase "natural rate" but of course I do. It was a masterpiece of persuasive

definition by Milton. Who could ever want an unnatural rate of unemployment? That was beautiful. But that issue is where attention ought to be directed. Similarly on the inflation side: why, for instance, are macroeconomists not talking about the shift from goods to services, and whether this changes price behavior very much, in a way that should be built into calculations? How exactly does international competition have an effect on movements of the domestic price level?

Greg Mankiw: One of the papers that influenced my thinking on this topic is by Jim Stock and Mark Watson. They wrote a paper about ten years ago trying to estimate how big the confidence intervals were around the NAIRU, as I think they called it, or the natural rate, which I'll take as a synonym. The answer was really big. At any point in time, you really don't know what it is. When Ben Bernanke goes to his staff and says, is the current unemployment rate higher or lower than the natural rate, they really need to say, "Who the hell knows?" But that does raise the question: is this concept useful for policymakers if one of the key parameters is immeasurable; that is, when we try to measure it, it has a large standard error? Is it useful merely as a theoretical construct to finish the Keynesian model, or is it practically useful when the Fed sits down and sets monetary policy if we can't estimate a key parameter in the equation?

John Taylor: Well, there is uncertainty, but it is not like it is useless. Bob is making a good point about how do you explain unemployment going from 2 to 7 percent in France, but that said, you have got to get out and do the empirical work. You get the macro data and you get the micro data. I don't think there is any substitute for that. But I don't want to give up on the concept of the Phillips curve. It seems to me that the concept is useful, but it's hard to estimate. That is just life. It has always been hard to estimate. It was hard at the beginning. It's a difficult policy problem, a difficult statistical problem. But why would you want to give up on the concept?

Greg Mankiw: No, I'm not giving up the concept, but it must be the case that the usefulness of the concept for making policy diminishes as the uncertainty about a key parameter increases.

John Taylor: Well, I'll give an example, just to respond. Some people, including my colleague Bob Hall, say that we should ignore the unemployment rate in the Taylor rule because it is so uncertain. Just look at inflation and ignore unemployment. I don't agree with that. I think that's a real mistake. Whether you want to use the GDP gap or other measures to measure capacity, you need to have some measure of capacity, some measure of where the overall economy is to do policy. So, I wouldn't go to the extreme of getting rid of all these concepts that try to measure where we are relative to normal.

Greg Mankiw: One of the things that we've been observing lately is tremendous increases in commodity prices. Some people may have noticed that. There has been a big literature that has tried to build in things like oil prices into the Phillips curve. There are other economists, and I think that Milton Friedman is one of them, who argue that is nonsense. The increase in the price of oil is an increase in the relative price of oil, and inflation is a monetary phenomenon. According to Friedman, relative prices have nothing to do with the inflation rate. But I think most of the people here would probably disagree with Milton's assessment. I don't know if you will agree with me on that. Do we have a good way to think about how relative price changes like the relative price of oil should fit into the overall inflation rate? What do you think is a good framework for thinking about that?

Robert Solow: Well, I don't know that we have an elegant framework for thinking about that at all, but I believe the point that Greg just made does correspond to the common sense of us all. One possible story about the current time is that the world is trying to bring about an increase in the relative price of oil. Maybe also the relative price of food, for all we know. It is one thing to be high and mighty and say, inflation is always and everywhere a monetary phenomenon, so therefore there is no reason why this event should affect the general price level at all. Other prices will fall far enough to offset the effect on the general price level of an increase in the price of oil, so there we are. We all know, whether we care to admit it or not, that there is some asymmetry about rising and falling prices, at least for many prices. It is going to be very difficult to generate the

market-required rise in the price of oil without having the general price level rise too. Just to be paradoxical and annoy people, we could insist that this is not inflation if the price level rises just once to accommodate the change in the relative price of oil. A one-time rise in the price level is not inflation in the relevant sense. There is some truth to that view. If we had a clearer picture of the asymmetry of price behavior, if there is such asymmetry, then I think you could do as John described, and sit and stare at the ceiling and find an equation.

Greg Mankiw: So you think oil price increases and decreases are fundamentally very different? Your asymmetry story suggests that.

Robert Solow: This implies nothing theoretically special about oil, except that it is important. The general point is that important changes in relative prices cannot be brought about without changes in the price level because prices are not infinitely or equally flexible up and down.

Greg Mankiw: John, do you want to say anything about oil?

John Taylor: I think that, in principle, it is very hard to distinguish the price level shifts and the inflation changes, as Bob is mentioning. In fact, I think it's hard because there are these tendencies for prices to pass through to other prices, and that is the inflation dynamic that we are trying to study. We do have equations. I looked at the period from the late 1960s and compared different countries, and it seems to me that you got much less of a pass-through of oil prices to the general inflation rate in countries that had monetary policies that were focused on keeping the inflation rate low, and oil price shocks were much less costly in terms of output, too.

So, if you just simply let the oil prices just go through to the inflation rate, it can be quite dangerous in terms of leading to higher inflation. I like the theorem that showed that a more inflation-focused monetary policy leads to a smaller pass-through in oil prices. Plus it seems to hold empirically. So, that seems to me a very important message, and we shouldn't forget it.

Robert Solow: I agree with about two-thirds of what John just said, especially about the danger of just letting the oil price go. I agree that

there will be a smaller effect on the price level the more restrictive monetary policy is. But I don't see why it should be less costly. I thought that Larry Ball had found, some time ago, in a cross-country comparison, that disinflation through monetary restriction was in fact quite costly.

John Taylor: Well, I'd say think of the cost in terms of how much better the business cycle has been in terms of smaller volatility of both output and inflation, since central banks were more aggressive in terms of target inflation.

Greg Mankiw: Okay, I have one more question and then I want to open it up to the floor for questions for these two gentlemen. I started this discussion off by looking back, and now I want to look forward. What are the big unanswered questions that the next generation of macroeconomists should be focused on in this line of work?

Robert Solow: I've already shot my wad on that. I think that there are changes in the structure of the economy, like the shift from goods to services, like the growing importance of import competition, and, presumably, some decrease in the domestic degree of monopoly, and I can't believe that changes of that kind and magnitude do not affect price behavior. That would include the part of price behavior that we try to capture in one version or another of the Phillips curve. Those things, and they are only the obvious ones, ought to be studied.

John Taylor: I would like to see much more work along the lines of testing the theories on micro data, getting into the details like Pete Klenow and Mark Bils are beginning to do it. I think it is very productive. You are able to discriminate against the different models. In my view, we need more work in the wage area, where we have very little micro-testing of the wage equations. The more that can be done along those lines, the better, and I think that will help lead to better theories. My second idea is to think more about the macro equations that we are trying to explain by the micro theories. Right now, we've got lots of micro theories out there for just about every macro equation that you can think of, whether money demand or price adjustment. None of them work perfectly. Maybe we are searching for something that we are never going to find, like *the* micro-

economic theory of price adjustment. You know, there are lots of different models out there that explain different aspects of price setting, and I don't know if that is necessarily a bad thing. I tried to think of an analogy from physics or physical science, and here is my analogy. It is probably a terrible one, but one of the most famous macro equations in physics, if you like, is Newton's Law of Gravity. The force of gravity depends on the product of the masses of the two bodies divided by the distance between them squared. It's an amazing equation and it fits very well.

Yet, there is no microeconomic foundational equivalent for that equation. There is no overarching theory in the sense that we call microeconomics to explain that theory. So, in this sense, maybe what we should be looking for is better macro equations. By the way, physicists have tried to find micro theories for that macro equation, done lots of work, and they just haven't found it.

The bottom line: maybe we should be looking for robust macro price equations that incorporate many different types of price adjustment at the micro level.

Greg Mankiw: Okay, thank you very much. Let me start off with taking some questions from the floor now on anything we've said or anything we didn't say.

Jeff Fuhrer: Jeff Fuhrer with the Federal Reserve Bank of Boston. So, we talked a little bit about large relative price changes, such as the one we are witnessing in energy. We also talked about expectations and the importance of expectations. We also talked about the idea that we have a partial understanding of how expectations are formed, perhaps. Could you comment on, for current circumstances, how you would think about the ways in which large relative price changes like energy might or might not feed into expectations, notwithstanding the strong presumption we have that they shouldn't feed into expectations. Then, what would a monetary policymaker in real time do about that?

Robert Solow: First of all, I have some confidence that a large rise in the price of energy generates expectations of a large rise in the price of energy. This question comes back to a point that John was making. The

question is: how does that micro expectation feed into an expectation about general inflation, or, more precisely, the general price level? John is surely right that an atmosphere of fairly restrictive monetary policy will limit the extent to which the initial price increase would be generalized into further inflation, into a further general rise in the price level. But how exactly it works and how either the actions or the statements of the central bank can influence how the public translates relative price changes into expectations about the consumer price index or the PCE (personal consumption expenditures) deflator, that's not obvious. The question alerts us to further conceptual problems about expected inflation. There are various interest groups in the economy: bankers, investors, savers, lenders, borrowers, buyers and sellers and what not. There is no reason for them to react in the same way. How does one aggregate expectations?

John Taylor: I agree with that. Just another thought is that while we talk a lot about globalization, I think in the current environment you really have to think of monetary policy in the world and not just one central bank. After all, oil is a globally priced commodity, and the pass-through and the price-setting expectations take place not just in the United States, but all over the world.

So, I've been thinking more and more about international coordination. We used to think that you didn't have to worry too much about coordination across countries, but now I think that this is an example where the pass-through not only depends on what the Fed does but on what the European Central Bank does and what other countries do. So, I'd like to think of it more on those terms.

Robert Leeson: Phillips was a very sophisticated theoretician, and in 1954 on a theoretical level developed the Phillips curve relationship between prices and output. In that model, published in the *Economic Journal*, there is a very precise role for inflationary expectations, such that when inflationary expectations become embedded in the system, the system becomes unstable. Now in Milton Friedman's model, inflationary expectations are equilibrating, and in Phillips's model these elements are destabilizing. Now, Phillips did try to find some empirical data looking

at prices and output and couldn't find it. But Henry Phelps Brown gave him some other data and that is the origin of the empirical Phillips curve.

Now, my question is that Phillips never really properly related his empirical work to his theoretical work, but then nobody else did either. So, how come the sophisticated theoretical work on which John was working in the 1960s got disconnected from the empirical work into the seminal role that Phillips played in developing the theory of inflation expectations. In fact, he gave Milton the formula, the adaptive inflation expectations formula. How come that got washed out of the system as well, and then it appeared that Milton was kind of developing this new concept that in fact he derived from Bill in the first place?

Robert Solow: I can't answer the second half of that, but on the first half, I'm sure John will correctly deny that in his mind the theory ever got separated from the facts, and I will confirm in advance the truth of what he is about to say.

Greg Mankiw: There is one fact that connects with that, which is that high inflation countries tend to have a lot of inflation instability, and the Friedman story doesn't naturally lead you to that fact. If you already expect 25 percent annual inflation, there is no particular reason that it couldn't be stable at that rate, but we don't seem to observe that in practice. I don't think that we've completely got a handle on that phenomenon.

Robert Solow: I think that's a good point.

Greg Mankiw: Thank you.

Peter Hooper: Peter Hooper from Deutsche Bank. I have an empirical question for the panelists. Given the essential importance, I think, of inflation expectations and applying the augmented model, how would you judge the best way to try to measure this? Would you ask a small group of professional forecasters? Would you ask a somewhat larger group of households in a survey or would you depend on what financial markets are seeing? To add to that, how much weight might you put on lagged inflation in your measure?

John Taylor: Well, I think you can't look at just one group. Actually, Bob made this point already. Different consumers are going to have different ways to process the expectation of inflation, so I think you have to look at all of them. If you are talking about looking at it for policy reasons, it seems to me that we are probably better off not trying to become too dependent on these expectations. Look for policies that are robust. Milton Friedman had this fixed money growth rule. Where did expectations appear in that?

So, it's not necessarily the case that you need to take these expectations into account that way when formulating policy. And we are going to hear more tomorrow about how miserable we are in forecasting inflation. How can you expect all of those other people out there that you are surveying to be that much better? So it seems to me that you'd have to have some humility in using both expectations of and forecasts of inflation to formulate policy and try, the best you can, to find ways to do it without using those measures.

Greg Mankiw: The empirical literature on the Phillips curve tends to find a pretty big role for lagged inflation. It's a bit of embarrassment from a theoretical standpoint since lagged inflation shouldn't necessarily be there once you have better proxies for expected inflation, but it is. So, people start putting ad hoc things into the theory. For people who are working in this literature, I think, it is no question an embarrassment.

Robert Solow: Once you start down that line, I think you come quickly to what I find embarrassing and difficult. Even if you were to divide the population into 30 different groups, each of whom hold some expectations about inflation, how do you weight them together? What weights do you apply, since you were not going to have a model with 30 or even 10 different expected rates of inflation? I think the better part of wisdom here may be not to pretend that that there is a precise concept, not to pretend that there is a numerical variable that you can fit into your model and call "the expected rate of inflation." Then you go on from there. If you do go on from there, there might easily be a role for lagged inflation as an indicator, as a proxy, as a last resort.

Michael Kiley: I'd like to follow up on Greg's comment that the usefulness of a theory must in some way be inversely related to our degree of uncertainty about the concepts that enter it. That idea must be at least somewhat important when we think about the natural rate of unemployment and how it links to inflation. When I think about writing down a model to help inform monetary policy, not just for forecasting purposes, I ask myself, well then what would it take for it not to be useful? To move to another model, you have to have the other model. So, if I weren't using that framework, what could I go to? I think we are at a point where this is the core theory that we have used to think about the links between unemployment and inflation and, like John said, we have uncertainty and we have to take that into account—but if there is no other option then we use what is there.

Greg Mankiw: Is that question for me?

Michael Kiley: Yes.

Greg Mankiw: I'm somewhat sympathetic to the argument that in the Taylor rule, for example, you give some weight to the unemployment gap, but the worse you are at measuring it, the lower the weight should be. What I learned from Mark Watson and Jim Stock is that we are very bad at measuring it. That is just a conjecture—I haven't written down the model to establish that notion. From a theoretical perspective, I have no problems with the idea of a Phillips curve. The question is, if the key parameters are very, very badly measured, then one has to question the utility for formulating policy. I think the Phillips curve becomes less useful. You end up setting policy by the seat of the pants.

V.V. Chari: V.V. Chari, University of Minnesota. Professor Solow suggested that there was a lot of asymmetry, and price increases in some sense being much more likely than price decreases. I think that sounds plausible and sounds right, but one of the valuable things I learned from Bils and Klenow's work is that just isn't so in the data. So, Mark and Pete looked at data at the level of individual commodities and at level of individual stores for a period in the United States when inflation was very moderate. What they found was that whenever prices changed, about

half of the prices went up on average by about 9 percent, and half of the prices fell on average by about 8 percent. That came as a surprise to me. But I think it is worthwhile remembering that data.

My second point is that much has been made of the difficulties of getting inflation persistence. Recent work by Argia Sbordone, Tim Cogley, and Peter Ireland has argued that some of the findings of inadequate inflation persistence may come from misspecification of monetary policy. Conventional Taylor rule-type policies seem to not be using enough inflation, but formulations in which the monetary authority's inflation target is a random walk seem to do much better. That also seems to be much more consistent with data from the pricing of long- and short-term bonds. Those are two cautionary things worth keeping in mind.

Joseph Carson: Joe Carson from AllianceBernstein. I have a question about your relative price discussion. Now you are talking about relative prices moving because of oil, but before we used to talk about housing and stocks or whatever. Do you think the problem of relative prices versus absolute prices is that we do not have a broad price index to cover all of these price movements? Do we focus on too narrow a price measure?

Robert Solow: I don't know.

John Taylor: No.

Greg Mankiw: Well, this does raise a question of which price index we should be focusing on. For example, there is wages versus prices. Samuelson and Solow, if I remember, was prices; Phillips was wages. *The Economist* magazine suggests that we stick equity prices in the price index. I don't think a lot of people here propose that, but there are a lot of prices in the economy, so which prices should we be focusing on when we think about the Phillips curve?

Robert Solow: I don't think that that's an abstract question nor a theoretical question. I think that you are looking for robust relationships, and I would settle for any price index that did a good job.

Greg Mankiw: So, we should use a price index that maximizes the fit of the Phillips curve?

Robert Solow: No –

Greg Mankiw: That’s an interesting project. I’m not even sure what that would look like.

Robert Solow: In all such problems, there is an interplay between common sense and goodness of fit. One wants both of those characteristics. If the price of thumb tacks happens to work best, surely you don’t think I would go for it. (But I would sure ask myself why it works best.)

Greg Mankiw: The thumb tacks standard.

Zvi Bodie: Zvi Bodie, Boston University. I am curious as to why, particularly Bob, why you wouldn’t consider the spread between the conventional Treasury rate and TIPS (Treasury Inflation-Protected Securities) as a reasonable proxy for—

Robert Solow: I thought of that, and as I turned things over in my mind, that might have been the best thing I could come up with. But it’s still not clear that the ordinary TIPS spread would tell you anything about what is in the mind of those who are engaged in wage determination or in the setting of some other prices. The reason why I did not fix on that is because it is a fairly narrowly restricted market. We know from surveys that there is a wide spread of expectations if you ask individuals how much higher they expect the price level to be a year from now. Somehow that has to be weighted and averaged. Inside that spread, there are groups that may behave differently, and it’s not clear to me that the people whose purchases and sales, or intentions to buy and sell, actually determine the price level also determine the TIPS spread, even indirectly.

Greg Mankiw: There is one other problem, I think, which is there may be differences in risk and liquidity premiums associated with these different instruments. I believe that the Cleveland Fed has some corrected series where they can say that the spread that you are looking at has been corrected for fluctuations over time in some liquidity premium as they estimated from some model. Still, if you are thinking you have to do that, your answer is only as good as the model of liquidity premium. So, if you think that is an important correction that needs to be done, that raises the possibility of measurement error there.

Barry Bluestone: Barry Bluestone, Northeastern University. I know tomorrow morning we are going to have a session on this, but I'm wondering to what extent, particularly when we look at the shifts in the Phillips curve, it is possible to model some of the institutional factors that, Bob, you had started to talk about. I am thinking of changes in the strength of the trade union movement, union density, trade policy, environmental policy, industry deregulation, and so on. I wonder if you could just comment on how we might even bring that into a model and what the impact might be of this inclusion?

Robert Solow: I agree with the general point and in fact, 48 years ago, Samuelson and I mentioned those things as important things to think about. It's damned hard to do that sensibly. When you do try, it devolves too soon into storytelling, and that sort of turns you off. It may be, although I'm not optimistic, that international differences could play a role there. The small year-to-year changes in union density, for instance, are never going to show up as a variable with a significant coefficient in regressions of that kind. It is conceivable that you might be able to interpret country differences, where institutional differences are much larger as a clue. But then there are so many other international differences that could mess up the relationship, if any.

Greg Mankiw: Allan Meltzer.

Allan Meltzer: I think this has been a very useful discussion because it's brought up a lot of the major problems. If one is teaching macroeconomics, especially to undergraduates, then I would think you would surely want to use the Phillips curve to illustrate what was going on and why, and how they might think about macroeconomic issues. When it comes to the policy issues, I think it's important—it's not just my view, but the view of the two most successful chairmen of the Federal Reserve, Paul Volcker and Alan Greenspan—that the Phillips curve was not very useful. In fact, it was not useful at all. Both of them didn't use it. Their quotations on this score will be in my forthcoming history of the Federal Reserve.

Why is that? I think that among the things that are part of the problem is the one that Bob Solow talked about. There isn't enough made of the distinction between large relative price changes and changes in the rate

of price change that will persist through time given the policies that we have. Another is that we don't pay much attention, to the extent to which people inform you of their expectations, trying to decide whether what they are observing is a temporary change or a persistent change—that is, whether it's one that is going to go on for a while. Surely in the history, one finds that 1973 and 1979 oil price changes largely ended by coming down; that is, the price of oil eventually came down. Much of the cost of those changes was born by wages; real wages fell or rose very slowly during that period.

The present change in the price of energy appears to be much more of a persistent change. Wouldn't you want to expect that this change is going to have a different effect than what the history would tell us about the 1970s? I think it should. I don't see where the models that we use do a very good job of trying to make this distinction between changes in relative prices and changes in the absolute price level and between persistent and relative changes. I think that affects not only what happens to prices, but it also affects what happens to output. For example, in the change that we observed in output in the 1990s, how long and how persistent would the change in productivity be that came at that point? Most of the models of the Phillips curve that we have do not have very explicit attention to details of that kind that would be important. Finally, I would say that much of the problem in translating the theory to the policy is it's not clear to me, it's never been clear to me, that the theory applies quarterly, but much of the policy is aimed at doing things quarterly, and that is a source, I believe, of errors.

I think economists probably have better ideas about what is going to happen over time than they do of what is going to happen from quarter to quarter. But that's not a view that my other colleagues necessarily share.

Greg Mankiw: Can I ask you a question about the two Fed chairmen that you mentioned? Did they think the Phillips curve wasn't very useful or that it was just a dumb idea? Did they think about it as theoretically useful, but not useful for me as a policymaker, or that it's kind of a goofy idea for you academics to think about?

Allan Meltzer: It's hard for me to pin down exactly what Volcker thought because he seemed to change his mind a lot. But with Greenspan, he was explicit. He said that this is a fine theoretical construct. It doesn't have much to do with what I have to do.

Robert Solow: On that issue I want to point out two things. First of all, Allan, it may be true that two very successful Fed chairmen had no use for the Phillips curve. The same is true of the unsuccessful chairmen. So there's not a lot of power in that test.

The second point is that I'm not sure that it is the right question to ask. Playing with the Phillips curve algebraically or graphically is something that some people have a temperament to do and others not. But I wonder if you had asked either Paul Volcker or Alan Greenspan this question—"if you strongly suspected that excess demand was emerging in the labor market, would you think wages were likely to rise?"—how do you think they would have replied?

Allan Meltzer: That's not a theoretical question, that's an empirical question. It's a question about whether that theory applies to the labor market, and the answer they would give is yes. The question then becomes for them, does that explain to me what is going to happen in the near term to the rate of price change and the rate of wage change and the rate of output change and the unemployment rate, and the answer they gave to that question was no.

Robert Solow: That I don't understand.

John Taylor: Maybe I could comment here. I think people have different ways of talking about the Phillips curve. Sometimes people say that the Phillips curve means there is a permanent trade-off between inflation and employment, others that there is only a short-run such Phillips curve. Some people refer to one simple equation. Others contend that you have to bring in wages too, at least. Still others say that you have to bring in an interaction between wages and prices, and I tend to agree with that viewpoint. One equation just isn't going to make it. So, when you read in the press that somebody believes in the Phillips curve, you think he must

be out of his mind. But that article is probably referring to one particular caricature of the Phillips curve or one particular aspect of it. Wage and price dynamics have been part of economics since the days of Hume, and it seems to me that's what people want to understand more about, whether you say they use the Phillips curve or not.

Greg Mankiw: With that, let me bring the session to a close. We thank both of you and the audience for the questions. Thank you very much.

Phillips Curve Inflation Forecasts

James H. Stock and Mark W. Watson

1. Introduction

Inflation is hard to forecast. There is now considerable evidence that Phillips curve forecasts do not improve upon good univariate benchmark models. Yet the backward-looking Phillips curve remains a workhorse of many macroeconomic forecasting models and continues to be the best way to understand policy discussions about the rates of unemployment and inflation.

After some preliminaries set forth in section 2, this paper begins its analysis in section 3 by surveying the past fifteen years of literature (since 1993) on inflation forecasting, focusing on papers that conduct a pseudo out-of-sample forecast evaluation.¹ A milestone in this literature is Atkeson and Ohanian (2001), who considered a number of standard Phillips curve forecasting models and showed that none improve upon a four-quarter random walk benchmark over the period 1984–1999. As we observe in this survey, Atkeson and Ohanian deserve the credit for forcefully making this point; however, their finding has precursors dating back at least to 1994. The literature after Atkeson and Ohanian finds that their specific result depends rather delicately on the sample period and the forecast horizon. If, however, one uses other univariate benchmarks (in particular, the unobserved components-stochastic volatility model of Stock and Watson (2007)), the broader point of Atkeson and Ohanian—that, at least since 1985, Phillips curve forecasts do not outperform univariate benchmarks on average—has been confirmed by several studies. The development of this literature is illustrated empirically using six prototype inflation forecasting models: three univariate

models, two backward-looking Phillips curve models—Gordon’s (1990) “triangle” model and an autoregressive-distributed lag model using the unemployment rate—and a model using the term spread, specifically the yield spread between one-year Treasury bonds and 90-day Treasury bills.

It is difficult to make comparisons across papers in this literature because the papers use different sample periods, different inflation series, and different benchmark models, and the quantitative results in the literature are curiously dependent upon these details. In section 4, we therefore undertake an empirical study that aims to unify and to assess the results in the literature using quarterly U.S. data from 1953:Q1–2008:Q1. This study examines the pseudo out-of-sample performance of a total of 192 forecasting procedures (157 distinct models and 35 combination forecasts), including the six prototype models of section 3, applied to forecasting five different inflation measures (CPI-all, CPI-core, PCE-all, PCE-core, and the GDP deflator). This study confirms the main qualitative results of the literature, although some specific results are found not to be robust. Our study also suggests an interpretation of why the literature’s conclusions strongly depend on the sample period. Specifically, one of our key findings is that the performance of Phillips curve forecasts is episodic: there are times, such as the late 1990s, when Phillips curve forecasts improved upon using univariate forecasts, but there are other times (such as the mid-1990s) when a forecaster would have been better off using a univariate forecast. This finding provides a rather more nuanced interpretation of Atkeson and Ohanian’s (2001) conclusion concerning Phillips curve forecasts, one that is consistent with the sensitivity of findings in the literature to the sample period.

A question that is both difficult and important is what this episodic performance implies for an inflation forecaster today. On average, over the past 15 years, it has been very hard to beat the best univariate model using any multivariate inflation forecasting model (Phillips curve or otherwise). But suppose you are told that next quarter the economy would plunge into recession, with the unemployment rate jumping by 2 percentage points. Would you change your inflation forecast? The literature is now full of formal statistical evidence suggesting that this information should be ignored, but we suspect that an applied forecaster would nevertheless revise downward his or her forecast of inflation over the one- to

two-year horizon. In the final section, we suggest some reasons why this revision might be justified.

2. Notation, Terminology, Families of Models, and Data

This section provides preliminary details concerning the empirical analysis and gives the six prototype inflation forecasting models that will be used in section 3 as a guide to the literature. We begin by reviewing some forecasting terminology.

Terminology

***b*-period inflation.** Inflation forecasting tends to focus on the one-year or two-year horizons. We denote *b*-period inflation by $\pi_t^b = b^{-1} \sum_{i=0}^{b-1} \pi_{t-i}$, where π_t is the quarterly rate of inflation at an annual rate; that is, $\pi_t = 400 \ln(P_t/P_{t-1})$ (using the log approximation), where P_t is the price index in quarter *t*. Four-quarter inflation at date *t* is $\pi_t^4 = 100 \ln(P_t/P_{t-4})$, the log approximation to the percentage growth in prices over the previous four quarters.

Direct and iterated forecasts. There are two ways to make an *b*-period ahead model-based forecast. A direct forecast has π_{t+b}^b as the dependent variable and *t*-dated variables (variables observed at date *t*) as regressors; for example, π_{t+b}^b could be regressed on π_t^b and the date-*t* unemployment rate (u_t). At the end of the sample (date *T*), the forecast of π_{T+b}^b is computed “directly” using the estimated forecasting equation. In contrast, an iterated forecast is based on a one-step ahead model; for example, π_{t+1} could be regressed on π_t , which is then iterated forward to compute future conditional means of π_s , $s > T + 1$, given data through time *t*. If predictors other than past π_t are used, then this requires constructing a subsidiary model for the predictor, or alternatively, modeling π_t and the predictor jointly—for example, as a vector autoregression (VAR)—and iterating the joint model forward.

Pseudo out-of-sample forecasts; rolling and recursive estimation. Pseudo out-of-sample forecasting simulates the experience of a real-time forecaster by performing all model specification and estimation using data through date *t*, making a *b*-step ahead forecast for date *t* + *b*, then moving forward to date *t* + 1 and repeating this through the sample.²

Pseudo out-of-sample forecast evaluation captures model specification uncertainty, model instability, and estimation uncertainty, in addition to the usual uncertainty of future events.

Model estimation can either be rolling (using a moving data window of fixed size) or recursive (using an increasing data window, always starting with the same observation). In this paper, rolling estimation is based on a window of ten years, and recursive estimation starts in 1953:Q1 or, for series starting after 1953:Q1, the earliest possible quarter.

Root mean squared error and rolling RMSE. The root mean squared forecast error (RMSE) of b -period ahead forecasts made over the period t_1 to t_2 is

$$(1) \quad RMSE_{t_1, t_2} = \sqrt{\frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} (\pi_{t+b}^b - \pi_{t+bt}^b)^2},$$

where π_{t+bt}^b is the pseudo out-of-sample forecast of π_{t+b}^b made using data through date t . This paper uses rolling estimates of the RMSE, which are computed using a weighted centered 15-quarter window:

$$(2) \quad \text{rolling RMSE}(t) = \sqrt{\frac{\sum_{s=t-7}^{t+7} K(|s-t|/8) (\pi_{s+b}^b - \pi_{s+bs}^b)^2}{\sum_{s=t-7}^{t+7} K(|s-t|/8)}},$$

where K is the biweight kernel, $K(x) = (15/16)(1 - x^2)^2 \mathbf{1}(|x| \leq 1)$.

Prototypical Inflation Forecasting Models

Single-equation inflation forecasting models can be grouped into four families: (1) forecasts based solely on past inflation; (2) forecasts based on activity measures (“Phillips curve forecasts”); (3) forecasts based on the forecasts of others; and (4) forecasts based on other predictors. This section lays out these families and provides prototype examples of each.

(1) **Forecasts based on past inflation.** This family includes univariate time series models such as autoregressive integrated moving average (ARIMA) models and nonlinear or time-varying univariate models. We also include in this family of forecasts those in which one or more inflation measure, other than the series being forecasted, is used as a predictor; for example, past Consumer Price Index (CPI) core inflation or past growth in wages could be used to forecast CPI-all inflation.

Three of our prototype models come from this family and serve as forecasting benchmarks. The first is a direct autoregressive (AR) forecast, computed using the direct autoregressive model,

$$(3) \quad \pi_{t+h}^b - \pi_t = \mu^b + \alpha^b(L)\Delta\pi_t + v_{t+h}^b, \quad (\text{AR(AIC)})$$

where μ^b is a constant, $\alpha^b(L)$ is a lag polynomial written in terms of the lag operator L , v_{t+h}^b is the h -step ahead error term (we will use v generically to denote regression error terms), and the superscript b denotes the quantity for the h -step ahead direct regression. In this prototype AR model, the lag length is determined by the Akaike Information Criterion (AIC) over the range of 1 to 6 lags. This specification imposes a unit autoregressive root.

The second prototype model is the Atkeson-Ohanian (2001) random walk model, in which the forecast of the four-quarter rate of inflation, π_{t+4}^4 , is the average rate of inflation over the previous four quarters, π_t^4 (Atkeson and Ohanian only considered four-quarter ahead forecasting). The Atkeson-Ohanian model thus is,

$$(4) \quad \pi_{t+4}^4 = \pi_t^4 + v_{t+4}^4 \quad (\text{AO}).$$

The third prototype model is the Stock-Watson (2007) unobserved components-stochastic volatility (UC-SV) model, in which π_t has a stochastic trend τ_t , a serially uncorrelated disturbance η_t , and stochastic volatility:

$$(5) \quad \pi_t = \tau_t + \eta_t, \text{ where } \eta_t = \sigma_{\eta,t} \zeta_{\eta,t}, \quad (\text{UC-SV})$$

$$(6) \quad \tau_t = \tau_{t-1} + \varepsilon_t, \text{ where } \varepsilon_t = \sigma_{\varepsilon,t} \zeta_{\varepsilon,t}$$

$$(7) \quad \ln \sigma_{\eta,t}^2 = \ln \sigma_{\eta,t-1}^2 + v_{\eta,t}$$

$$(8) \quad \ln \sigma_{\varepsilon,t}^2 = \ln \sigma_{\varepsilon,t-1}^2 + v_{\varepsilon,t}$$

where $\zeta_t = (\zeta_{\eta,t}, \zeta_{\varepsilon,t})$ is independent and identically distributed (i.i.d.) $N(0, I_2)$, $v_t = (v_{\eta,t}, v_{\varepsilon,t})$ is i.i.d. $N(0, \gamma I_2)$, ζ_t , and v_t are independently distributed, and γ is a scalar parameter. Although η_t and ε_t are conditionally normal given $\sigma_{\eta,t}$ and $\sigma_{\varepsilon,t}$, unconditionally these are random mixtures of normal random variables and can have heavy tails. This is a one-step ahead model and forecasts are iterated. The UC-SV model has only one parameter, γ , which controls the smoothness of the stochastic volatility

process. Throughout, we follow Stock and Watson (2007) and set $\gamma = 0.04$.

(2) *Phillips curve forecasts.* We interpret Phillips curve forecasts broadly to include forecasts produced using an activity variable, such as the unemployment rate, an output gap, or output growth, perhaps in conjunction with other variables, to forecast inflation or the change in inflation. This family includes both backward-looking Phillips curves and New Keynesian Phillips curves, although the latter appear infrequently (and only recently) in the inflation forecasting literature.

We consider two prototype Phillips curve forecasts. The first is Gordon's (1990) "triangle model," which in turn is essentially the model in Gordon (1982) with minor modifications.³ In the triangle model, inflation depends on lagged inflation, the unemployment rate u_t , and supply shock variables z_t :

$$(9) \quad \pi_{t+1} = \mu + \alpha^G(L)\pi_t + \beta(L)u_{t+1} + \gamma(L)z_t + v_{t+1}. \quad (\text{triangle})$$

The prototype triangle model used here is that in Gordon (1990), in which (9) is specified using the contemporaneous value plus 4 lags of u_t (total civilian unemployment rate ages 16+ years, seasonally adjusted), contemporaneous value plus 4 lags of the rate of inflation of food and energy prices (computed as the difference between the inflation rates in the deflator for "all-items" personal consumption expenditure (PCE) and the deflator for PCE less food and energy), lags 1 through 4 of the relative price of imports (computed as the difference of the rates of inflation of the GDP deflator for imports and the overall GDP deflator), two dummy variables for the Nixon wage-price control period, and 24 lags of inflation, where $\alpha^G(L)$ imposes the step-function restriction that the coefficients are equal within the groups of lags 1–4, 5–8, ..., 21–24, and also that the coefficients sum to one (a unit root is imposed).

Following Gordon (1998), forecasts based on the triangle model (9) are iterated using forecasted values of the predictors, where those forecasts are made using subsidiary univariate AR(8) models of u_t , food and energy inflation, and import inflation.

The second prototype Phillips curve model is direct version of (9) without the supply shock variables and without the step-function restriction

on the coefficients. This model is an autoregressive distributed lag (ADL) model in which forecasts are computed using the direct regression,

$$(10) \quad \pi_{t+h}^b - \pi_t = \mu^b + \alpha^b(L)\Delta\pi_t + \beta^b(L)u_t + v_{t+h}^b, \quad (\text{ADL-}u)$$

where $\alpha^b(L)$ and $\beta^b(L)$ are unrestricted with degrees chosen separately by AIC (maximum lag of 4), and (like the triangle model) the ADL- u specification imposes a unit root in the autoregressive dynamics for π_t .

(3) *Forecasts based on forecasts of others.* The third family computes inflation forecasts from explicit or implicit inflationary expectations or forecasts of others. These forecasts include regressions based on implicit expectations derived from asset prices, such as forecasts extracted from the term structure of nominal Treasury debt (which by the Fisher relation should embody future inflation expectations) and forecasts extracted from the Treasury Inflation-Protected Securities (TIPS) yield curve. This family also includes forecasts based on explicit forecasts of others, such as median forecasts from surveys such as the Survey of Professional Forecasters.

Our prototypical example of forecasts in this family is a modification of the Mishkin (1990a) specification, in which the future change in inflation is predicted by a matched-maturity spread between the interest rates on comparable government debt instruments, with no lags of inflation. Here we consider direct 4-quarter ahead forecasts based on an ADL model using as a predictor the interest spread, $spread1_90_t$, between one-year Treasury bonds and 90-day Treasury bills:

$$(11) \quad \pi_{t+4}^4 - \pi_t = \mu + \alpha(L)\Delta\pi_t + \beta(L)spread1_90_t + v_{t+4}^4. \quad (\text{ADL-spread})$$

We emphasize that Mishkin's (1990a) regressions appropriately use term spread maturities matched to the change in inflation being forecasted, which for (11) would be the change in inflation over quarters $t + 2$ to $t + 4$, relative to $t + 1$. (A matched maturity alternative to $spread1_90_t$ in (11) would be the spread between one-year Treasuries and the federal funds rate, however those instruments have different risks.) Because the focus of this paper is Phillips curve regressions we treat this regression simply as an example of this family and provide references to recent studies of this family in section 3.3.

(4) *Forecasts based on other predictors.* The fourth family consists of inflation forecasts that are based on variables other than activity or expectations variables. An example is a 1970s-vintage monetarist model in which M1 growth is used to forecast inflation. Forecasts in this fourth family perform sufficiently poorly relative to the three other approaches that these play negligible roles both in the literature and in current practice, so to avoid distraction we do not track a model in this family as a running example.

Data and Transformations

The data set is quarterly for the United States from 1953:Q1–2008:Q1. Monthly data are converted to quarterly data by computing the average value for the three months in the quarter prior to any other transformations; for example, quarterly CPI is the average of the three monthly CPI values, and quarterly CPI inflation is the percentage growth (at an annual rate, using the log approximation) of this quarterly CPI.

We examine forecasts of five measures of price inflation: the GDP price deflator (PGDP), the CPI for all items (CPI-all), CPI excluding food and energy (CPI-core), the personal consumption expenditure deflator (PCE-all), and the personal consumption expenditure deflator excluding food and energy (PCE-core).

In addition to the six prototype models, in section 4 we consider forecasts made using a total of 15 predictors, most of which are activity variables (GDP, industrial production, housing starts, the capacity utilization rate, etc.). The full list of variables and transformations is given in the appendix.

Gap variables. Consistent with the pseudo out-of-sample forecasting philosophy, the activity gaps used in the forecasting models in this paper are all one-sided. Following Stock and Watson (2007), gaps are computed as the deviation of the series (for example, log GDP) from a symmetric two-sided moving average (MA(80)) approximation to the optimal lowpass filter with pass band corresponding to periodicities of at least 60 quarters. The one-sided gap at date t is computed by padding observations at dates $s > t$ and $s < 1$ with iterated forecasts and backcasts based on an AR(4), estimated recursively through date t .

3. An Illustrated Survey of the Literature on Phillips Curve Forecasts, 1993–2008

This section surveys the literature during the past fifteen years (since 1993) on inflation forecasting in the United States. The criterion for inclusion in this survey is providing empirical evidence on inflation forecasts (model- and/or survey-based) in the form of a true or pseudo out-of-sample forecast evaluation exercise. Such an evaluation can use rolling or recursive forecasting methods based on final data, it can use rolling or recursive methods using real-time data, or it can use forecasts actually produced and recorded in real time such as survey forecasts. Most of the papers discussed here focus on forecasting at horizons of policy relevance, one or two years. The primary interest is in forecasting overall consumer price inflation (PCE, CPI), core inflation, or economy-wide inflation (GDP deflator). There is little work on forecasting producer prices, although a few papers consider producer prices as a predictor of headline inflation.

This survey also discusses some papers in related literatures; however, we do not attempt a comprehensive review of those related literatures. One such literature concerns the large amount of interesting work that has been done on inflation forecasting in countries other than the United States: see Rünstler (2002), Hubrich (2005), Canova (2007), and Diron and Mojon (2008) for recent contributions and references. Another closely related literature concerns in-sample statistical characterizations of changes in the univariate and multivariate inflation process in the United States (e.g., Taylor 2000; Brainard and Perry 2000; Cogley and Sargent 2002, 2005; Levin and Piger 2004; and Pivetta and Reis 2007) and outside the United States (e.g., the papers associated with the European Central Bank Inflation Persistence Network 2007). There is in turn a literature that asks whether these changes in the inflation process can be attributed, in a quantitative (in-sample) way, to changes in monetary policy; papers in this vein include Estrella and Fuhrer (2003), Roberts (2004), Sims and Zha (2006), and Primiceri (2006). A major theme of this survey is time-variation in the Phillips curve from a forecasting perspective, most notably at the end of the disinflation of the early 1980s but more subtly throughout the post-1984 period. This time-variation

is taken up in a great many papers; for example, papers estimating a time-varying NAIRU and time variation in the slope of the Phillips curve. In addition, there is a massive theoretical and empirical literature that develops and analyzes the New Keynesian Phillips curve (Roberts 1995). Papers in these literatures, however, are only discussed in passing unless they have a pseudo out-of-sample forecasting component.

The 1990s: Warning Signs

The great inflation and disinflation of the 1970s and the 1980s was the formative experience that dominated the minds and models of inflation forecasters through the 1980s and early 1990s, both because of the forecasting failures of 1960s-vintage (“non-accelerationist”) Phillips curves and, more mechanically, because most of the variation in the data comes from that period. The dominance of this episode is evident in figure 3.1,

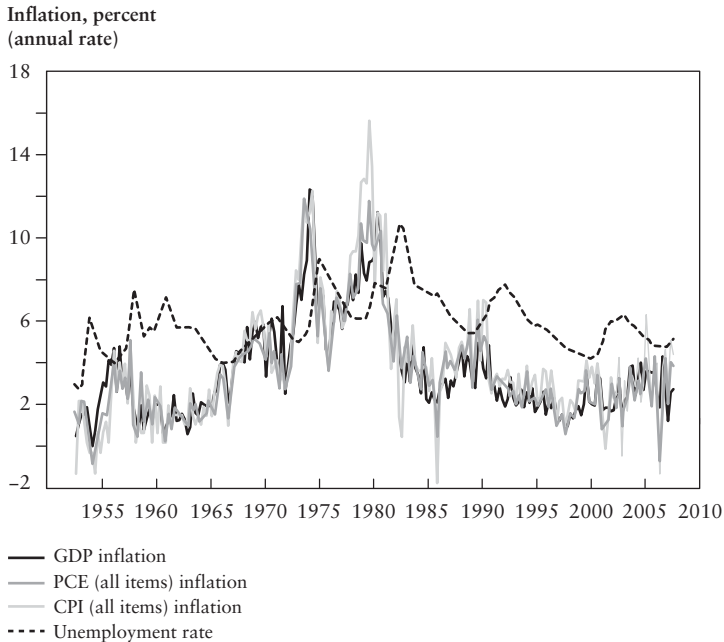


Figure 3.1

Quarterly U.S. Price Inflation at an Annual Rate as Measured by the GDP Deflator, PCE-all and CPI-all, and the Rate of Unemployment, 1953:Q1–2008:Q1

which plots the three measures of headline inflation (GDP, PCE-all, and CPI-all) from 1953:Q1 to 2007:Q4, along with the unemployment rate.

By the early 1980s, despite theoretical attacks on the backward-looking Phillips curve, Phillips curve forecasting specifications had coalesced around the Gordon (1982) triangle model (9) and variants. Figure 3.2 plots the rolling RMSE of the four-quarter ahead pseudo out-of-sample forecast of CPI-all inflation, computed using (2), for the recursively estimated AR(AIC) benchmark (3), the triangle model (9), and the ADL- u model (10). As can be seen in figure 3.2, these “accelerationist” Phillips curve specifications (unlike their non-accelerationist ancestors) did in fact outperform the AR(AIC) benchmark during the 1970s and 1980s.

While the greatest success of the triangle model and the ADL- u model was forecasting the fall in inflation during the early 1980s subsequent to the spike in the unemployment rate in 1980, in fact the triangle and ADL- u models improved upon the AR benchmark nearly uniformly from 1965 through 1990. The main exception occurred around 1986, when there was a temporary decline in oil prices. The four-quarter ahead pseudo out-of-sample forecasts produced by the AR(AIC), triangle, and




Figure 3.2
Rolling Root Mean Squared Errors for CPI-all Inflation Forecasts: AR(AIC), Triangle Model (constant NAIRU), and ADL- u Model

ADL- u models are shown respectively in panels (a)-(c) of figure 3.3. As can be seen in figure 3.3, the triangle model predicted too much too late: it initially failed to forecast the decline in inflation in 1986, then predicted inflation to fall further than it actually did. Interestingly, unlike the AR(AIC) and ADL- u models, triangle model forecasts did not over-extrapolate the decline in inflation in the early 1980s.


Stockton and Glassman (1987) documented the good performance of a triangle model based on the Gordon (1982) specification of the triangle model over the 1977–1984 period (they used the Council of Economic Advisors output gap instead of the unemployment rate and a 16-quarter, not 24-quarter, polynomial distributed lag). They reported a pseudo out-of-sample relative RMSE of the triangle model, relative to an AR(4) model of the change in inflation, of 0.80 (eight-quarter ahead iterated forecasts of inflation measured by the Gross Domestic Business Product fixed-weight deflator).⁴ Notably, Stockton and Glassman (1987) also emphasized that there seem to be few good competitors to this model: a variety of monetarist models, including some that incorporate expectations of money growth, all performed worse—in some cases, much worse—than the AR(4) benchmark. This said, the gains from using a Phillips curve forecast over the second half of the 1980s were slimmer than during the 1970s and early 1980s.

The earliest documentation of this relative deterioration of Phillips curve forecasts of which we are aware is a little-known (two Google Scholar cites) working paper by Jaditz and Sayers (1994). They undertook a pseudo out-of-sample forecasting exercise of CPI-all inflation using industrial production growth, the PPI, and the 90-day Treasury Bill rate in a VAR and in a vector error correction model (VECM), with a forecast period of 1986-1991 and a forecast horizon of one month. They reported a relative RMSE of .985 for the VAR and a relative mean squared error (MSE) in excess of one for the VECM, relative to an AR(1) benchmark.

Cecchetti (1995) also provided early evidence of instability in Phillips curve forecasts. However, that instability was apparent only using in-sample break tests and did not come through in his pseudo out-of-sample forecasting evaluation because of his forecast sample period. Cecchetti considered forecasts of CPI-all at horizons of one to four years based on

Inflation, percent
(annual rate)

(a) AR(AIC)



Inflation, percent
(annual rate)

(b) Triangle model




Figure 3.3
CPI-all Inflation and Pseudo Out-of-Sample Forecasts




Figure 3.3 (continued)

18 predictors, entered separately, for two forecast periods, 1977–1994 (ten-year rolling window) and 1987–1994 (five-year rolling window). Inspection of figure 3.2 indicates that Phillips curve forecasts did well on average over both of these samples, but that the 1987–1994 period was atypical of the post-1984 experience in that it is dominated by the relatively good performance of Phillips curve forecasts during the 1990 recession. Despite the good performance of Phillips curve forecasts over this period, using in-sample break tests Cecchetti (1995) found multiple breaks in the relationship between inflation and (separately) unemployment, the employment/population ratio, and the capacity utilization rate. He also found that good in-sample fit is essentially unrelated to subsequent pseudo out-of-sample forecasting performance.

Stock and Watson (1999) undertook a pseudo out-of-sample forecasting assessment of CPI-all and PCE-all forecasts at the one-year horizon using (separately) 168 economic indicators, of which 85 were measures of real economic activity (industrial production growth, unemployment, and so on). They considered recursive forecasts computed over two subsamples, 1970–1983 and 1984–1996. The split sample evidence indicated major changes in the relative performance of predictors in the two

subsamples; for example, the RMSE of the forecast based on the unemployment rate, relative to the AR benchmark, was .89 in the 1970–1983 sample but 1.01 in the 1984–1996 sample. Using in-sample test statistics, they also found structural breaks in the inflation-unemployment relationship, although interestingly these breaks were more detectable in the coefficients on lagged inflation in the Phillips curve specifications than on the activity variables.

Cecchetti, Chu, and Steindel (2000) examined CPI inflation forecasts at the two-year horizon using (separately) 19 predictors, including activity indicators. They reported dynamic forecasts in which future values of the predictors are used to make multiperiod ahead forecasts (future employment is treated as known at the time the forecast is made, so these are not pseudo out-of-sample). Strikingly, they found that activity-based dynamic forecasts (unemployment, employment-population ratio, and capacity utilization rate) typically underperformed the AR benchmark over this period at the one-year horizon.

Brayton, Roberts, and Williams (1999) considered long-lag Phillips curve specifications. In their pseudo out-of-sample results (six inflation measures, four- and eight-quarters ahead, forecast period of 1975–1998), standard Phillips curve forecasts are outperformed by longer-lag versions (25-quarter polynomial distributed lag specifications). Using in-sample statistics, they reject coefficient stability; they attribute the instability to a shift in the NAIRU in the 1990s, not to a change in the slope coefficients in the long-lag specification.

A final paper documenting poor Phillips curve forecasting performance, contemporaneous with Atkeson and Ohanian (2001), is Camba-Mendez and Rodriguez-Palenzuela (2003; originally published as a 2001 European Central Bank working paper). They showed that inflation forecasts at the one-year horizon based on realizable (that is, backward-looking) output gap measures, for the forecast period 1980–1999, underperform the AR benchmark.

In short, during the 1990s a number of papers provided results that activity-based inflation forecasts provided a smaller advantage relative to an AR benchmark after the mid-1980s than these forecasts had before. Ambiguities remained, however, because this conclusion seemed to depend on the sample period and specification, and in any event

one could find predictors which were exceptions in the sense that they appeared to provide improvements in the later sample, even if their performance was lackluster in the earlier sample.

Atkeson and Ohanian (2001)

Atkeson and Ohanian (2001) (AO) resolved the ambiguities in this literature from the 1990s by adopting a new, simple univariate benchmark: the forecast of inflation over the next four quarters is the value of four-quarter inflation today.⁵ Atkeson and Ohanian showed that this four-quarter random walk forecast improved substantially upon the AR benchmark over the 1984–1999 period. Figure 3.4 plots the moving RMSE of four-quarter ahead forecasts of CPI-all inflation for three univariate forecasts: the AR(AIC) forecast (3), the AO forecast (4), and the UC-SV forecast (5)–(8). Because the AO benchmark improved upon the AR forecast over the 1984–1999 period, and because the AR forecast had more or less the same performance as the unemployment-based Phillips curve on average over this period (see figure 3.2), it is not surprising that the AO forecast




Figure 3.4

Rolling Root Mean Squared Errors for Univariate CPI-all Inflation Forecasts: AR(AIC), Atkeson-Ohanian (AO), and Unobserved Components-Stochastic Volatility (UC-SV) Models

outperformed the Phillips curve forecast over the 1984–1999 period. As Atkeson and Ohanian dramatically showed, across 264 specifications (three inflation measures, CPI-all, CPI-core, and PCE-all, two predictors, the unemployment rate and the Chicago Fed National Activity Index [CFNAI], and various lag specifications), the relative RMSEs of a Phillips curve forecast to the AO benchmark ranged from 0.99 to 1.94: gains from using a Phillips curve forecast were negligible at best, and some Phillips curve forecasts went badly wrong. Atkeson and Ohanian went one step further and demonstrated that, over the 1984–1999 period, Greenbook forecasts of inflation also underperformed their four-quarter random walk forecast.

As figures 3.2 and 3.4 demonstrate, one important source of the problem with Phillips curve forecasts was their poor performance in the second half of the 1990s, a period of strong, but at the time unmeasured, productivity growth that held down inflation. The apparent quiescence of inflation in the face of strong economic growth was puzzling at the time (for example, see Lown and Rich 1997).

An initial response to Atkeson and Ohanian's result was to check whether their claims were accurate; with a few caveats, by and large these were, although only for the post-1984 period they considered. Sims (2002) confirmed Atkeson and Ohanian's results post-1984, but stressed that the AO model performs poorly over the 1979–1983 sample period. Bernanke (2003) cited unpublished work by Board of Governors staff that Atkeson and Ohanian's conclusions do not extend to periods of greater macroeconomic and inflation volatility. Fisher, Liu, and Zhou (2002) used rolling regressions with a 15-year window and showed that Phillips curve models outperformed the AO benchmark in 1977–1984, and also showed that for some inflation measures and some periods the Phillips curve forecasts outperform the AO benchmark post-1984 (for example, Phillips curve forecasts improve upon AO forecasts of PCE-all over 1993–2000). Fisher, Liu, and Zhou (2002) also pointed out that Phillips curve forecasts based on the CFNAI achieve 60–70 percent accuracy in directional forecasting of the change of inflation, compared with 50 percent for the AO coin flip. Fisher, Liu, and Zhou suggested that Phillips curve forecasts do relatively poorly in periods of low inflation volatility and after a regime shift.

Stock and Watson (2003) extended Atkeson and Ohanian's analysis to additional activity predictors (as well as other predictors) and confirmed the dominance of the AO forecast over 1985–1999 at the one-year horizon. Brave and Fisher (2004) extended Atkeson and Ohanian's and Fisher, Liu, and Zhou's (2002) analyses by examining additional predictors and combination forecasts. Brave and Fisher's (2004) findings are broadly consistent with Fisher, Liu, and Zhou (2002) in the sense that they found some individual and combination forecasts that outperform AO over 1993–2000, although not over 1985–1992. Orphanides and van Norden (2005) focused on Phillips curve forecasts using real-time gap measures, and they concluded that although *ex post* gap Phillips curves fit well using in-sample statistics, when real-time gaps and pseudo out-of-sample methods are used these too improve upon the AR benchmark prior to 1983, but fail to do so over the 1984–2002 sample period.

There are three notable recent studies that confirm Atkeson and Ohanian's basic finding and extend it, with qualifications. First, Stock and Watson (2007) focused on univariate models of inflation and confirmed that the good performance of the AO random walk forecast, relative to other univariate models, is specific to the four-quarter horizon and to Atkeson and Ohanian's sample period. At any point in time, the UC-SV model implies an IMA(1,1) model for inflation, with time-varying coefficients. The forecast function of this IMA(1,1) closely matches the implicit AO forecast function over the 1984–1999 sample—however the models diverge over other subsamples. Moreover, the rolling IMA(1,1) is in turn well approximated by a ARMA(1,1) because the estimated AR coefficient is nearly one.⁶ Stock and Watson (2007) also reported some (limited) results for bivariate forecasts using activity indicators (unemployment, one-sided gaps, and output growth) and confirmed Atkeson and Ohanian's finding that these Phillips curve forecasts fail to improve systematically on the AO benchmark or the UC-SV benchmark over the AO sample at the four-quarter horizon.

Second, Canova (2007) undertook a systematic evaluation of four- and eight-quarter ahead inflation forecasts for G7 countries using recursive forecasts over 1996–2000, using a variety of activity variables (unemployment, employment, output gaps, GDP growth) and other indicators (yield curve slope, money growth) as predictors. He found that, for the United States, bivariate direct regressions and trivariate VARs and BVARs did not

improve upon the univariate AO forecast, and that there was evidence of instability of forecasts based on individual predictors. Canova (2007) also considered combination forecasts and forecasts generated using a New Keynesian Phillips curve. Over the 1996–2000 U.S. sample, combination forecasts provided a small improvement over the AO forecast, and the New Keynesian Phillips curve forecasts were never the best and generally fared poorly. In the case of the United States, at least, these findings are not surprising in light of the poor performance of Phillips curve forecasts during the low-inflation boom of the second half of the 1990s.

Third, Ang, Bekaert, and Wei (2007) conducted a thorough assessment of forecasts of CPI, CPI-core, CPI excluding housing, and PCE inflation, using 10 variants of Phillips curve forecasts, 15 variants of term structure forecasts, combination forecasts, and ARMA(1,1) and AR(1)-regime switching univariate models in addition to AR and AO benchmarks. They too confirmed Atkeson and Ohanian's basic message that Phillips curve models fail to improve upon univariate models over forecast periods 1985–2002 and 1995–2002. Ang, Bekaert, and Wei's (2007) results constitute a careful summary of the current state of knowledge of inflation forecasting models (both Phillips curve and term structure) in the United States. One finding in their study is that combination forecasts do not systematically improve on individual indicator forecasts, a result that is puzzling in light of the success reported elsewhere of combination forecasts (we return to this puzzle below).⁷

Following Romer and Romer (2000),⁸ Sims (2002) and Ang, Bekaert, and Wei (2007) considered professional and survey forecasts, variously including the Federal Reserve Board's Greenbook, Data Resources, Inc., the Michigan Survey of Consumer Sentiment, the Philadelphia Fed's Livingston Survey, the Survey of Professional Forecasters, and Blue Chip surveys. Sims concluded that the Greenbook forecast outperformed the Atkeson-Ohanian forecast over the 1979–1995 period, but not over 1984–1995. Ang, Bekaert, and Wei (2007) found that, for the inflation measures that the survey respondents are asked to forecast, the survey forecasts nearly always beat the ARMA(1,1) benchmark, their best-performing univariate model over the 1985–2002 period; this finding is surprising in light of the literature that has postdated Atkeson and Ohanian (2001). Further study of rolling regressions led Ang, Bekaert, and Wei (2007) to suggest that the relatively good performance of the survey

forecasts might be due to the ability of professional forecasters to recognize structural change more quickly than automated regression-based forecasts.⁹

An alternative forecast, so far unmentioned, is that inflation is constant. This forecast works terribly over the full sample but Diron and Mojon (2008) found out that, for PCE-core from 1995:Q1–2007:Q4, a forecast of a constant 2.0 percent inflation rate outperforms AO and AR forecasts at the eight-quarter ahead horizon, although the AO forecast is best at the four-quarter horizon. Diron and Mojon choose 2.0 percent as representative of an implicit inflation target over this period; however, because the United States does not have an explicit *ex ante* inflation target, this value was chosen retrospectively and this choice does not constitute a pseudo out-of-sample forecast.

The evidence of forecast instability in the foregoing papers is based on changes in relative RMSEs, in some cases augmented by Diebold-Mariano (1995) or West (1996) tests using asymptotic critical values. As a logical matter, the apparent statistical significance of the changes in the relative RMSEs between sample periods could be a spurious consequence of using a poor approximation to the sampling distribution of the relevant statistics. Accordingly, Clark and McCracken (2006a) undertook a bootstrap evaluation of the relative RMSEs produced using real-time output gap Phillips curves for forecasting the GDP price deflator and CPI-core. They reached the more cautious conclusion that much of the relatively poor performance of forecasts using real-time gaps could simply be a statistical artifact that is consistent with a stable Phillips curve, although they did find evidence of instability in coefficients on the output gap. One interpretation of the Clark-McCracken (2006a) finding is that, over the 1990–2003 period, there are only 14 nonoverlapping observations on the four-quarter ahead forecast error, and estimates of ratios of variances with 14 observations inevitably have a great deal of sampling variability. Rossi and Sekhposyan (2007) also took a careful look at the statistical evidence for breaks using pseudo out-of-sample forecast statistics; theirs is one of the few studies also to use real-time data. Their formal tests for a one-time reversal of forecast performance find a sharp decline in the predictive ability of Phillips curve forecasts post-1984. Additional work is needed to reconcile the results in Clark and McCracken (2006a) and Rossi and Sekhposyan (2007).

Attempts to Resuscitate Multivariate Inflation Forecasts, 1999–2007

One response to Atkeson and Ohanian's findings has been to redouble efforts to find reliable multivariate forecasting models for inflation. Some of these efforts used statistical tools, including dynamic factor models, other methods for using a large number of predictors, time-varying parameter multivariate models, and nonlinear time series models. Other efforts exploited restrictions arising from economics, in particular from no-arbitrage models of the term structure. Unfortunately, these efforts have failed to produce substantial and sustained improvements over the AO or UC-SV univariate benchmarks.

Many-predictor forecasts I: dynamic factor models. The plethora of activity indicators used in Phillips curve forecasts indicates that there is no single, most natural measure; in fact, these indicators can all be thought of as different measures of overall economic activity. This suggests modeling the activity variables jointly using a dynamic factor model (Geweke 1977, Sargent-Sims 1977), estimating the common latent factor (underlying economic activity), and using that estimated factor as the activity variable in Phillips curve forecasts. Accordingly, Stock and Watson (1999) examined different activity measures as predictors of inflation, estimated (using principal components, as justified by Stock and Watson 2002) as the common factor among 85 monthly indicators of economic activity, and also as the first principal component of 165 series, including the activity indicators plus other series. In addition to using information in a very large number of series, Stock and Watson (2002) showed that principal components estimation of factors can be robust to certain types of instability in a dynamic factor model. Stock and Watson's (1999) empirical results indicated that these estimated factors registered improvements over the AR benchmark and over single-indicator Phillips curve specifications in both 1970–1983 and 1984–1996 subsamples.

A version of the Stock-Watson (1999) common factor, computed as the principal component of 85 monthly indicators of economic activity, has been published in real time since January 2001 as the Chicago Fed National Activity Index (CFNAI). Hansen (2005) confirmed the main findings in Stock and Watson (1999) about the predictive content of these estimated factors for inflation, relative to a random walk forecast over a forecast period of 1960–2000.

Recent studies, however, have raised questions about the marginal value of Phillips curve forecasts based on estimated factors, such as the CFNAI, for the post-1985 data. As discussed above, Atkeson and Ohanian showed that the AO forecast outperformed CFNAI-based Phillips curves over the 1984–1999 period; this is consistent with Stock and Watson (1999) finding a small improvement in dynamic factor model (DFM) forecasts over this period because Stock and Watson (1999) used an AR benchmark. Banerjee and Marcellino (2006) also found that Phillips curve forecasts using estimated factors perform relatively poorly for CPI-all inflation over a 1991–2001 forecast period. On the other hand, for the longer sample of 1983–2007, Gavin and Kliesen (2008) found that recursive factor forecasts improve upon both the direct AR(12) (monthly data) and AO benchmarks (relative RMSEs are between .88 and .95). In a finding that is inconsistent with Atkeson and Ohanian (2001) and with figure 3.4, Gavin and Kliesen (2008) also found that the AR(12) model outperforms AO at the 12-month horizon for three of the four inflation series; presumably this surprising result is either a consequence of using a slightly different sample than Atkeson and Ohanian (in particular, including 1983) or indicates some subtle differences between using quarterly data (as in Atkeson and Ohanian and in figure 3.4) and monthly data.

Additional papers which use estimated factors to forecast inflation include Watson (2003), Bernanke, Bovin, and Elias (2005), Boivin and Ng (2005, 2006), D’Agostino and Giannone (2006), and Giacomini and White (2006). In an interesting meta-analysis, Eichmeier and Ziegler (2008) considered a total of 52 studies of inflation and/or output forecasts using estimated factors, including 22,849 relative RMSEs for inflation forecasts in the United States and other countries. The dependent variable in their meta-regressions is the RMSE of a factor forecast relative to a benchmark. Eichmeier and Ziegler (2008) concluded that factor model inflation forecasts tend to outperform small model forecasts by a small margin. They also concluded that factor inflation forecasts tend to improve as the horizon increases, and that they improve as the number of series used to estimate the factors increases. Eichmeier and Ziegler’s (2008) meta-regressions do not control for sample period, a strategy that permits estimating the average performance of different methods but prevents examining the time-varying relative performance found in the other papers reviewed here. Although Eichmeier and Ziegler (2008) do

include indicator variables for the category of benchmark in their meta-regressions, the relative performance of those benchmarks changes over time and this too complicates the interpretation of their results for the purposes of this survey.

Many-predictor forecasts II: Forecast combination, Bayesian Model Averaging, Bagging, and other methods. Other statistical methods for using a large number of predictors are available and have been tried for forecasting inflation. One approach is to use leading index methods, in essence a model selection methodology. In the earliest high-dimensional inflation forecasting exercise of which we are aware, Webb and Rowe (1995) constructed a leading index of CPI-core inflation formed using 7 of 30 potential inflation predictors, selected recursively by selecting indicators with a maximal correlation with one-year ahead inflation over a 48-month window, thereby allowing for time variation. This produced a leading index with time-varying composition that improved upon an AR benchmark over the 1970–1994 period; however, Webb and Rowe (1995) did not provide sufficient information to assess the success of this index post-1983.

A second approach is to use forecast combination methods, in which forecasts from multiple bivariate models (each using a different predictor, lag length, or specification) are combined. Combination forecasts have a long history of success in economic applications—see the review in Timmermann (2006)—and are less susceptible to structural breaks in individual forecasting regressions because, in effect, these combination forecasts average out intercept shifts (Hendry and Clements 2004). Papers that include combination forecasts (pooled over models) include Stock and Watson (1999, 2003), Clark and McCracken (2006b), Canova (2007), Ang, Bekaert, and Wei (2007), and Inoue and Kilian (2008). Although combination forecasts often improve upon the individual forecasts, on average these do not substantially improve upon, and are often slightly worse than, factor-based forecasts.

A third approach is to apply model combination or model averaging tools, such as Bayesian Model Averaging (BMA), bagging, and LASSO, developed in the statistics literature for prediction using large data sets. Wright (2003) applied BMA to forecasts of CPI-all, CPI-core, PCE, and the GDP deflator, obtained from 30 predictors, and finds that BMA tended to improve upon simple averaging. Wright's (2003) relative

RMSEs are considerably less than one during the 1987–2003 sample; however this appears to be a consequence of a poor denominator model (an AR(1) benchmark) rather than good numerator models. Inoue and Kilian (2008) considered CPI-all forecasts with 30 predictors using bagging, LASSO, and factor-based forecasts (first principal component), along with BMA, pretest, shrinkage, and some other methods from the statistical literature. They reported a relative RMSE for the single-factor forecast of .80, relative to an AR(AIC) benchmark at the 12-month horizon over their 1983–2003 monthly sample. This is a surprisingly low value in light of Atkeson and Ohanian and subsequent literature, but (like Wright 2003) this low relative RMSE appears to be driven by the use of the AR (instead of AO or UC-SV) benchmark and by the sample period, which includes 1983. Inoue and Kilian (2008) found negligible gains from using the large dataset methods from the statistics literature: the single-factor forecasts beat almost all the other methods they examine, although in most cases the gains from the factor forecasts are slight (the relative RMSEs, relative to the single-factor model, range from .97, for LASSO, to 1.14).

A fourth approach is to model all series simultaneously using high-dimensional VARs with strong parameter restrictions. Bańbura, Gianonne, and Reichlin (2008) performed a pseudo out-of-sample experiment forecasting CPI-all inflation using Bayesian VARs with 3 to more than 100 variables. Over the 1970–2003 sample, they found substantial improvements of medium- to large-dimensional VARs relative to very low-dimensional VARs, but their results are hard to relate to the others in this literature because they do not report univariate benchmarks and do not examine split samples.

In summary, in some cases (some inflation series, some time periods, and some horizons) it appears to be possible to make gains using many predictor methods, either factor estimates or other methods. However, those gains are modest and not systematic and do not substantially overturn Atkeson and Ohanian's (2001) negative results.

Nonlinear models. If the conditional expectation of future inflation is a nonlinear function of the predictors, and if the predictors are persistent, then linear approximations to the conditional mean function can exhibit persistent time variation. Thus the time variation documented above could be a consequence of using linear models. Accordingly, one

approach to the apparent time variation in the inflation-output relation is to consider nonlinear Phillips curves and nonlinear univariate time series models. There is a substantial literature on nonlinear Phillips curves that reports only in-sample measures of fit, not pseudo out-of-sample forecasts; see Dupasquier and Ricketts (1998) and Barnes and Olivei (2003) for references. Barnes and Olivei (2003) is a noteworthy paper in that literature; they consider a piecewise linear specification and use dynamic simulations to argue that this specification (with a time-varying NAIRU) provides a better description of the late 1990s and early 2000s than does a linear specification. Their specification is capable of producing the episodically effective Phillips curve forecasts seen in the forecasting literature. Yet their use of only in-sample statistics makes it difficult to compare their findings to the forecasting literature that is the focus of this survey. Papers that evaluate nonlinear inflation forecasting models using pseudo out-of-sample methods include Dupasquier and Ricketts (1998), Moshiri and Cameron (2000), Tkacz (2000), Ascari and Marrocu (2003), and Marcellino (2008).

We read the conclusions of this literature on nonlinear Phillips curves and nonlinear univariate time series models as negative. Although nonlinearities are found using in-sample statistics, the pseudo out-of-sample literature fails to confirm any benefits of nonlinear models for forecasting inflation. Marcellino (2008) examined univariate rolling and recursive CPI-all forecasts (over 1980–2004 and 1984–2004) using logistic smooth transition autoregressions and neural networks (a total of 28 nonlinear models) and found little or no improvement from using nonlinear models. He also documented that nonlinear models can produce outlier forecasts, presumably because of overfitting. Ascari and Marrocu (2003) and Moshiri and Cameron (2000), who apply artificial neural networks to Canadian data, also provided negative conclusions. These negative results in the pseudo out-of-sample literature mean that exploitable nonlinearities have not been found, but not that they do not exist. Indeed, the in-sample results of Barnes and Olivei (2005) presage findings reported below in section 5.

Structural term structure models. Until now, this survey has concentrated on forecasts from the first two families of inflation forecasts (prices-only and Phillips curve forecasts). One way to construct inflation forecasts in the third family—forecasts based on forecasts made by

others—is to make inflation forecasts using the term structure of interest rates, as in (11). Starting with Barsky (1987), Mishkin (1990a, 1990b, 1991), and Jorion and Mishkin (1991), there is a large literature that studies such forecasting regressions. The findings of this literature, which are reviewed in Stock and Watson (2003), are generally negative; that is, term spread forecasts do not improve over Phillips curve forecasts in the pre-1983 period, and they do not improve over a good univariate benchmark in the post-1984 period.

This poor performance of first-generation term spread forecasts is evident in figure 3.5, which plots the rolling RMSE of the pseudo out-of-sample forecast based on the recursively estimated term spread model (11), along with the RMSEs of the AR(AIC) and AO univariate benchmarks. Term spreads are typically one of the variables included in the forecast comparison studies discussed earlier (Fisher, Liu, and Zhou 2002, Canova 2007, and Ang, Bekaert, and Wei 2007) and these recent studies also reach the same negative conclusion about unrestricted term spread forecasting regressions, either as the sole predictor or when used in addition to an activity indicator.




Figure 3.5

Rolling Root Mean Squared Errors for CPI-all Inflation Forecasts: AR(AIC), Atkeson-Ohanian (AO), and Term Spread Model (ADL-*spread*)

Recent attempts to forecast inflation using term spreads have focused on employing economic theory, in the form of no-arbitrage models of the term structure, to improve upon the reduced-form regressions, such as (11). Most of this literature uses full-sample estimation and measures of fit; see Ang, Bekaert, and Wei (2007), DeWachter and Lyrio (2006), and Berardi (2007) for references. The one paper of which we are aware that produces pseudo out-of-sample forecasts of inflation is Ang, Bekaert, and Wei (2007), who considered four-quarter ahead forecasts of CPI-all, CPI-core, CPI-excluding housing, and PCE inflation using two no-arbitrage term structure models, one with constant coefficients and one with regime switches. Neither model forecasted well, with relative RMSEs (relative to an ARMA(1,1)) ranging from 1.05 to 1.59 for the four inflation series and two forecast periods (1985–2002 and 1995–2002).

We are not aware of any papers that evaluate the performance of inflation forecasts backed out of the TIPS yield curve, and such a study would be of considerable interest.

Forecasting using the cross-section of prices. Another approach is to try to exploit information in the cross-section of inflation indexes (percentage growth of sectoral or commodity group price indexes) for forecasting headline inflation. Hendry and Hubrich (2007) used four high-level subaggregates to forecast CPI-all inflation. They explored several approaches, including combining disaggregated univariate forecasts and using factor models. Hendry and Hubrich (2007) found that exploiting the disaggregated information directly to forecast the aggregate improves modestly over an AR benchmark in their pseudo out-of-sample forecasts of CPI-all over 1970–1983 but negligibly over the AO benchmark over 1984–2004 at the 12-month horizon; however they also found that no single method for using the subaggregates works best. If one uses heavily disaggregated inflation measures, then some method must be used to control parameter proliferation, such as the methods used in the many-predictor applications discussed above. In this vein, Hubrich (2005) presented negative results concerning the aggregation of components forecasts for forecasting the Harmonized Index of Consumer Prices in the eurozone. Reis and Watson (2007) estimated a dynamic factor using a large cross-section of inflation rates but did not conduct any pseudo out-of-sample forecasting.

Rethinking the notion of core inflation suggests different approaches to using the inflation subaggregates. Building on the work of Bryan and Cecchetti (1994), Bryan, Cecchetti, and Wiggins (1997) suggested constructing core inflation as a trimmed mean of the cross-section of prices, where the trimming was chosen to provide the best (in-sample) estimate of underlying trend inflation (measured variously as a 24- to 60-month centered moving average). Smith (2004) investigated the pseudo out-of-sample forecasting properties of trimmed mean and median measures of core inflation (forecast period 1990–2000). Smith (2004) reported that the inflation forecasts based on weighted-median core measures have relative RMSEs of .85 for CPI-all and .80 for PCE-all, relative to an exponentially-declining AR benchmark (she does not consider the AO benchmark), although oddly she found that the trimmed mean performed worse than the AR benchmark.

4. A Quantitative Recapitulation: Changes in Univariate and Phillips Curve Inflation Forecast Performance

This section undertakes a quantitative summary of the literature review in the previous section by considering the pseudo out-of-sample performance of a range of inflation forecasting models using a single consistent data set. The focus is on activity-based inflation forecasting models, although some other predictors are considered. We do not consider survey forecasts or inflation expectations implicit in the TIPS yield curve. As Romer and Romer (2000), Sims (2002), and Ang, Bekaert, and Wei (2007) showed, Greenbook and some median survey forecasts perform quite well and thus are useful for policy work—but our task is to understand how to improve upon forecasting systems, not to delegate this work to others.

Forecasting Models

Univariate models. The univariate models consist of the AR(AIC), Atkeson-Ohanian, and UC-SV models in section 2.2; direct AR models with a fixed lag length of four lags, (AR(4)) and Bayes Information Criterion lag selection (AR(BIC)); and iterated AR(AIC), MA(1), and AR(24) models, where the AR(24) model imposes the Gordon (1990) step function lag

restriction and the unit root in π_t . AIC and BIC model selection used a minimum of 0 and a maximum of six lags. Both rolling and recursively estimated versions of these models are considered. In addition some fixed-parameter models were considered: MA(1) models with fixed MA coefficients of 0.25 and 0.65 (these are taken from Stock and Watson 2007), and the monthly MA model estimated by Nelson and Schwert (1977), temporally aggregated to quarterly data (see Stock and Watson 2007, equation (7)).

Triangle and time-varying NAIRU models. Four triangle models are considered: specification (9), the results of which were examined in section 3; specification (9) without the supply shock variables (relative price of food and energy, import prices, and Nixon dummies); and these two versions with a time-varying (TV) NAIRU. The TV-NAIRU specification introduces random walk intercept drift into (9) following Staiger, Stock, and Watson (1997) and Gordon (1998); specifically, the TV-NAIRU version of (9) is

$$(12) \quad \pi_{t+1} = \alpha^G(L)\pi_t + \beta(L)(u_{t+1} - \bar{u}_t) + \gamma(L)z_t + v_{t+1},$$

$$(13) \quad \bar{u}_{t+1} = \bar{u}_t + \eta_{t+1},$$

where v_t and η_t are modeled as independent i.i.d. normal errors with relative variance σ_η^2/σ_v^2 (recall that $\alpha^G(1) = 1$ so a unit root is imposed in (12)). For the calculations here, σ_η^2/σ_v^2 is set to 0.1.

ADL Phillips curve models. The ADL Phillips curve models are direct models of the form,

$$(14) \quad \pi_{t+h}^b - \pi_t = \mu^b + \alpha^b(L)\Delta\pi_t + \beta^b(L)x_t + v_{t+h}^b,$$

where x_t is an activity variable (an output gap, growth rate, or level, depending on the series). Lag lengths for π_t and x_t are chosen separately by AIC and, alternatively, BIC.

ADL models using other predictors. ADL models are specified and estimated the same way as the ADL Phillips curve model (14), but the activity variable x_t is replaced by another predictor (term spreads, core inflation, and so on).

Combination forecasts. Let $\{\hat{\pi}_{i,t+h|t}^b\}$ denote a set of n forecasts of π_{t+h}^b , made using data through date t . Combined forecasts are computed in three ways: by “averaging” (mean, median, and trimmed mean); by

a MSE-based weighting scheme; or by using the forecast that is most recently best. The MSE-based combined forecasts f_t are of the form $f_t = \sum_{i=1}^n \lambda_{it} \hat{\pi}_{i,t+bt}^b$, where six methods are used to compute the weights $\{\lambda_{it}\}$:

$$(15) \quad (A) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2) / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2), \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} 0.9^j e_{i,t-j}^2,$$

$$(16) \quad (B) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2) / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2), \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} 0.95^j e_{i,t-j}^2,$$

$$(17) \quad (C) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2) / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2), \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} e_{i,t-j}^2,$$

$$(18) \quad (D) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2)^2 / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2)^2, \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} 0.9^j e_{i,t-j}^2,$$

$$(19) \quad (E) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2)^2 / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2)^2, \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} 0.95^j e_{i,t-j}^2,$$

$$(20) \quad (F) \quad \lambda_{it} = (1 / \hat{\sigma}_{it}^2)^2 / \sum_{j=1}^n (1 / \hat{\sigma}_{jt}^2)^2, \text{ with } \hat{\sigma}_{it}^2 = \sum_{j=0}^{39} e_{i,t-j}^2,$$

where $e_{i,t} = \pi_t^b - \hat{\pi}_{i,t|t-b}^b$ is the pseudo out-of-sample forecast error for the i^{th} b -step ahead forecast and the MSEs are estimated using a 10-year rolling window and, for methods (A), (B), (D), and (E), discounting.

Inverse MSE weighting (based on population MSEs) is optimal if the individual forecasts are uncorrelated, and methods (A) – (C) are different ways to implement inverse MSE weighting. Methods (D) – (F) give greater weight to better-performing forecasts than does inverse MSE weighting. Optimal forecast combination using regression weights as in Bates and Granger (1969) is not feasible with the large number of forecasts under consideration. As Timmerman (2006) notes, equal-weighting (mean combining) often performs well and Timmerman (2006) provides a discussion of when mean combining is optimal under squared error loss.

The “recent best” forecasts are the forecasts from the model that has the lowest cumulative MSE over the past four (or, alternatively, eight) quarters.

Finally, in an attempt to exploit the time-varying virtues of the UC-SV and triangle models, the recent best is also computed using only the UC-SV and triangle model (with time-varying NAIRU and z variables).

The complete description of models considered is given in the notes to table 3.1.

Results

The pseudo out-of-sample forecasting performance of each forecasting procedure (model and combining method) is summarized in tabular and graphical form.

The tabular summary consists of relative RMSEs of four-quarter ahead inflation forecasts, relative to the UC-SV benchmark, for six forecast periods; these are tabulated in tables 3.1–3.5 for the five inflation series. The minimum model estimation sample was 40 quarters, and blank cells in the table indicate that for at least one quarter in the forecast period there were fewer than 40 observations available for estimation.

The graphical summary of each model's performance is given in figures 3.6–3.11 for the five inflation series. Figure 3.6 presents the rolling RMSE for the UC-SV benchmark for the five inflation series, and figures 3.7–3.11 show the RMSE of the various forecasts relative to the UC-SV benchmark. Part (a) of figure 3.7–3.11 displays the rolling relative RMSE

RMSE in percentage
points of inflation




Figure 3.6
Rolling Root Mean Squared Errors for Inflation Forecasts, Unobserved Components-Stochastic Volatility Model, for All Five Inflation Series


for the prototype models, where the rolling RMSE for each model is computed using (2). Parts (b) – (d) plot the ratio of the rolling RMSE for each category of models, relative to the UC-SV model: univariate models in part (b), Phillips curve forecasts (ADL and triangle) in part (c), and combination forecasts in part (d). In each of parts (b) – (d), leading case models or forecasts are highlighted. The unlabeled relative RMSE paths, which are presented using small dots in panels (b)–(d) of figures 3.7–3.11, portray the rolling RMSEs of all other forecasting models in tables 3.1–3.5 for the relevant inflation series and the indicated category of forecast. For example, figure 3.7(c) represents the relative rolling RMSEs for all the Phillips curve forecasts listed in table 3.1, three of which are labeled in the figure while the rest remain unlabeled.

These tables and figures present a great many numbers and facts. Inspection of these results leads us to the following conclusions:

1. There is strong evidence of time variation in the inflation process, in predictive relations, and in Phillips curve forecasts. This is consistent with the literature review, in which different authors reach different conclusions about Phillips curve forecasts depending on the sample period.
2. The performance of Phillips curve forecasts, relative to the UC-SV benchmark, has a considerable systematic component (part (c) of figures 3.7–3.11): during periods in which the ADL- μ prototype model is forecasting well, reasonably good forecasts can be made using a host of other activity variables. In this sense, the choice of activity variable is secondary to the choice of whether one should use an activity-based forecast.
3. Among the univariate models considered here, with and without time-varying coefficients, there is no single model, or combination of univariate models, that has uniformly better performance than the UC-SV model. Of the 82 cells in table 3.1 that give relative RMSEs for univariate CPI-all forecasts in different subsamples, only four cells have RMSEs less than 1.00, the lowest of which is .95, and these instances are for fixed-parameter MA models in the 1960s and in the 1985–1992 period. Similar results are found for the other four inflation measures. In some cases, the AR models do quite poorly relative to UC-SV. For example, in the 2001–2007 sample the AR forecasts of CPI-all and PCE-all inflation have very large relative MSEs (typically exceeding 1.3). In general, the performance of the AR model, relative to the UC-SV or AO benchmarks, is series- and period-specific. This reinforces the remarks in the literature

RMSE in percentage points of inflation

(a) Prototype model forecasts



RMSE in percentage points of inflation

(b) Univariate forecasts





Figure 3.7

Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model: CPI-all

RMSE in percentage points of inflation

(c) Phillips curve forecasts



RMSE in percentage points of inflation

(d) Combination forecasts




Figure 3.7 (continued)





Figure 3.8
Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model: CPI-core

RMSE in percentage points of inflation

(c) Phillips curve forecasts



RMSE in percentage points of inflation

(d) Combination forecasts





Figure 3.8 (continued)

RMSE in percentage points of inflation

(a) Prototype model forecasts



RMSE in percentage points of inflation

(b) Univariate forecasts

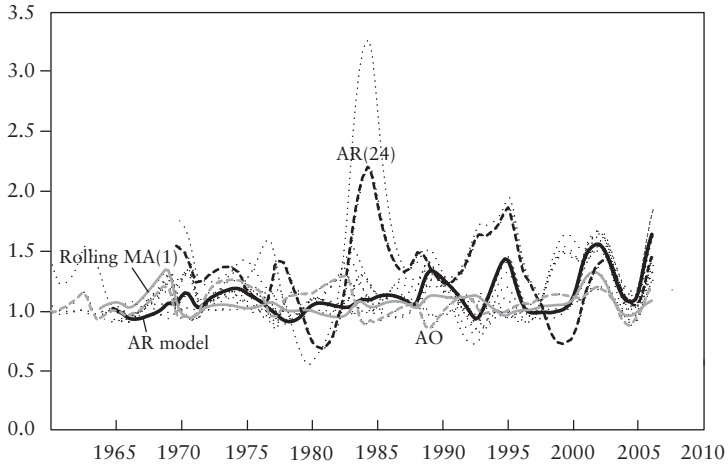
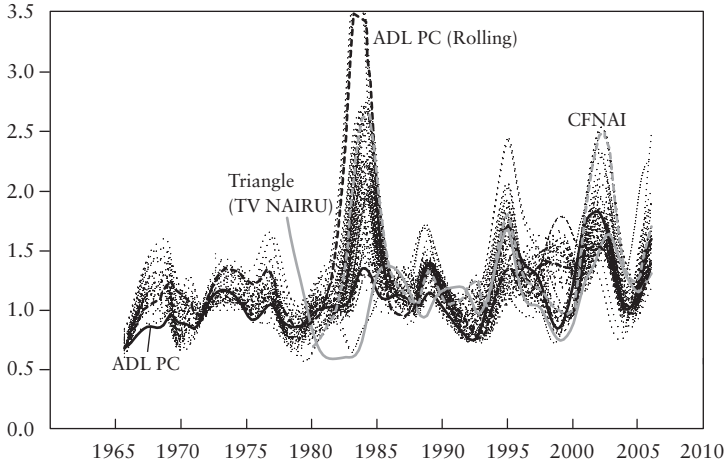


Figure 3.9
Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model: PCE-all

RMSE in percentage points of inflation

(c) Phillips curve forecasts



RMSE in percentage points of inflation

(d) Combination forecasts




Figure 3.9 (continued)


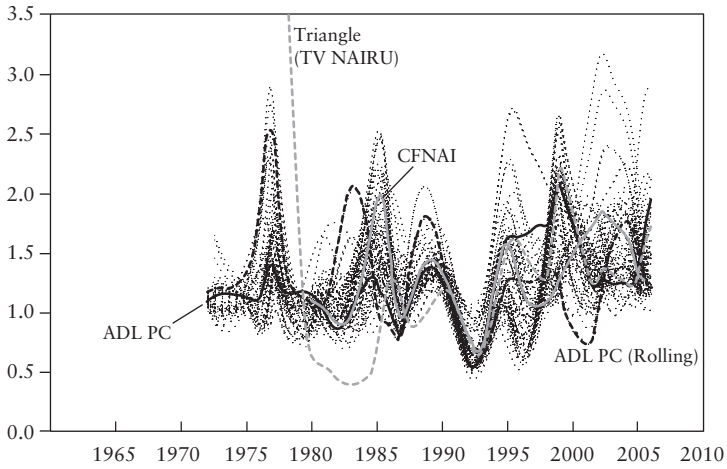


Figure 3.10
 Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model: PCE-core

RMSE in percentage points of inflation

(c) Phillips curve forecasts



RMSE in percentage points of inflation

(d) Combination forecasts




Figure 3.10 (continued)



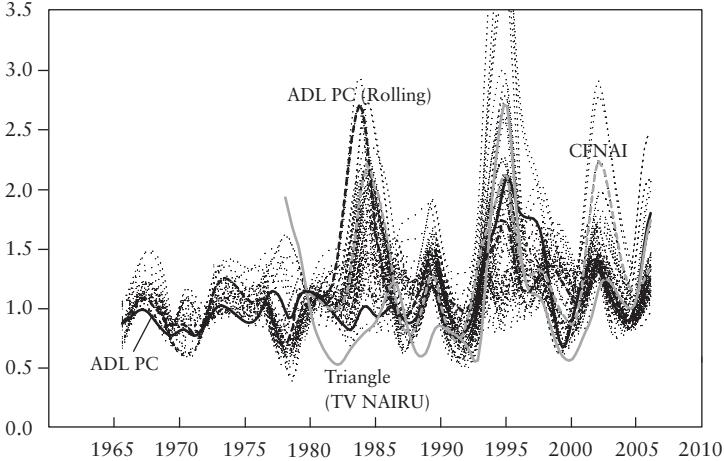


Figure 3.11
Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model: GDP Deflator

RMSE in percentage points of inflation

(c) Phillips curve forecasts



RMSE in percentage points of inflation

(d) Combination forecasts




Figure 3.11 (continued)

review about the importance of using a consistently good benchmark: the apparently good performance of a predictor for a particular inflation series over a particular period can be the result of a large denominator, not a small numerator.

4. Although some of the Phillips curve forecasts improved substantially on the UC-SV model during the 1970s and early 1980s, there is little or no evidence that it is possible to improve upon the UC-SV model on average over the full later samples. Nevertheless, there are notable periods and inflation measures for which Phillips curve models do quite well. The triangle model does particularly well during the high unemployment disinflation of the early 1980s for all five inflation measures. For CPI-all, PCE-all, and the GDP deflator, it also does well in the late 1990s, while for CPI-core and PCE-core the triangle model does well emerging from the 1990 recession. This episodically good behavior of the triangle model, and of Phillips-curve forecasts more generally, provides a more nuanced interpretation of the history of inflation forecasting models than the blanket Atkeson-Ohanian (2001) finding which, as stated in their paper's abstract, concluded that "none of the NAIRU forecasts is more accurate than the naïve forecast."

5. Forecast combining, which has worked so well in other applications (Timmerman 2006), generally improves upon the individual Phillips curve forecasts; however, the combination forecasts generally do not improve upon the UC-SV benchmark in the post-1993 periods. For example, for CPI-all, the mean-combined ADL-activity forecasts have a relative RMSE of .86 over 1977–1982 and .96 over 1985–1992; these mean-combined forecasts compare favorably to individual activity forecasts and to the triangle model. In the later periods, however, the forecasts being combined have relative RMSEs exceeding 1.0; combining them works no magic and fails to improve upon the UC-SV benchmark. Although some of the combining methods improve upon equal weighting, these improvements are neither large nor systematic. In addition, consistent with the results in Fisher, Liu, and Zhou (2002), factor forecasts (using the CFNAI) fail to improve upon the UC-SV benchmark on average over the later periods. These results are consistent with the lack of success found by attempts in the literature (before and after Atkeson and Ohanian 2001) to obtain large gains by using many predictors and/or model combinations.

6. Forecasts using predictors other than activities variables, while not the main focus of this paper, generally fare poorly, especially during the post-

1992 period. For example, the relative RMSE of the mean-combined forecast using nonactivity variables is at least 0.99 in each subsample in tables 3.1–3.5 (23 cases). We did not find substantial improvements using alternative measures of core (median and trimmed mean CPI) as predictors.¹⁰ Although our treatment of nonactivity variables is not comprehensive, these results largely mirror those in the literature.

5. When Were Phillips Curve Forecasts Successful, and Why?

If the relative performance of Phillips curve forecasts has been episodic, is it possible to characterize what makes for a successful or unsuccessful episode?

The relative RMSEs of the triangle and ADL- u model forecasts for headline inflation (CPI-all, PCE-all, and GDP deflator), relative to the UC-SV benchmark, are plotted in figure 3.12, along with the unemployment rate. One immediately evident feature is that the triangle model has substantially larger swings in performance than the ADL- u model. This said, the dates of relative success of these Phillips curve forecasts bear considerable similarities across models and inflation series. Both models perform relatively well for all series in the early 1980s, in the early 1990s, and around 1999; both models perform relatively poorly around 1985 and in the mid-1990s. These dates of relative success correspond approximately to dates of different phases of U.S. business cycles.

Figure 3.13 is a scatterplot of the quarterly relative RMSE for the triangle (panel a) and ADL- u (panel b) prototype models, versus the two-sided unemployment gap (the two-sided gap was computed using the two-sided version of the lowpass filter described in section 2), along with kernel regression estimates. The most striking feature of these scatterplots is that the relative RMSE is minimized, and is considerably less than 1.0, at the extreme values of the unemployment gap, both positive and negative. (The kernel regression estimator exceeds 1.0 at the most negative values of the unemployment gap for the triangle model in panel (a), but there are few observations in that tail.) When the unemployment rate is near the NAIRU (as measured by the lowpass filter), both Phillips curve models do worse than the UC-SV model. But when the unemployment




Figure 3.12
Rolling Root Mean Squared Errors, Relative to Unobserved Components-Stochastic Volatility Model, and the Unemployment Rate




Figure 3.12 (continued)

gap exceeds 1.5 percentage points in absolute value, the Phillips curve forecasts improve substantially upon the UC-SV model. Because the gap is largest in absolute value around turning points, the Phillips curve models provide improvements over the UC-SV model around cyclical turning points, but not during normal times.

Figure 3.14 takes a different perspective on the link between performance of the Phillips curve forecasts and the state of the economy, by plotting the relative RMSE against the four-quarter change in the unemployment rate. The relative improvements in the Phillips curve forecasts do not seem as closely tied to the change in the unemployment rate as to the gap (the apparent improvement at very large changes of the unemployment rate is evident in only a few observations).

Figures 3.15–3.17 examine a conjecture in the literature—that Phillips curve forecasts are relatively more successful when inflation is volatile—by plotting the rolling relative RMSE against the four-quarter change in four-quarter inflation. These figures provide only limited support for this conjecture, as do similar scatterplots (not provided here) of the rolling RMSE against the UC-SV estimate of the instantaneous variance of the first difference of the inflation rate. It is true that the worst performance occurs when in fact inflation is changing very little but, other than for the

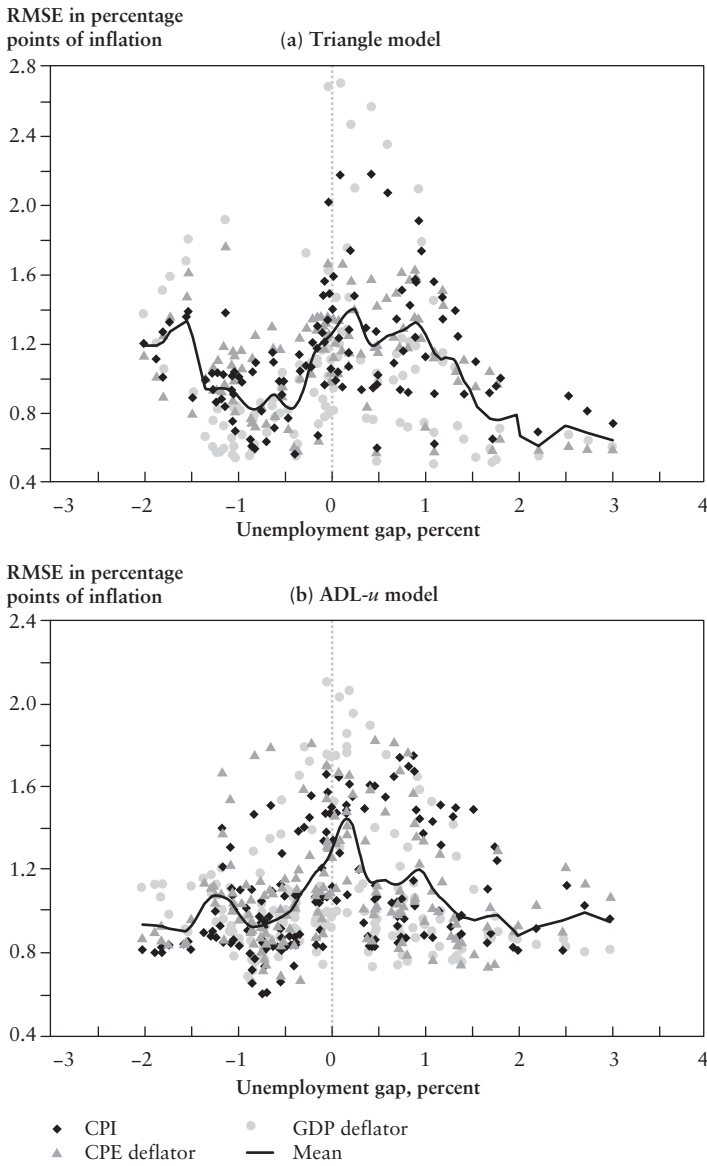


Figure 3.13

Scatterplot of Rolling Root Mean Square Errors of Headline Inflation Forecasts, Relative to Unobserved Components-Stochastic Volatility, versus the Unemployment Gap (two-sided bandpass)

Note: Mean is kernel regression estimate using data for all three series. Each point represents a quarter.

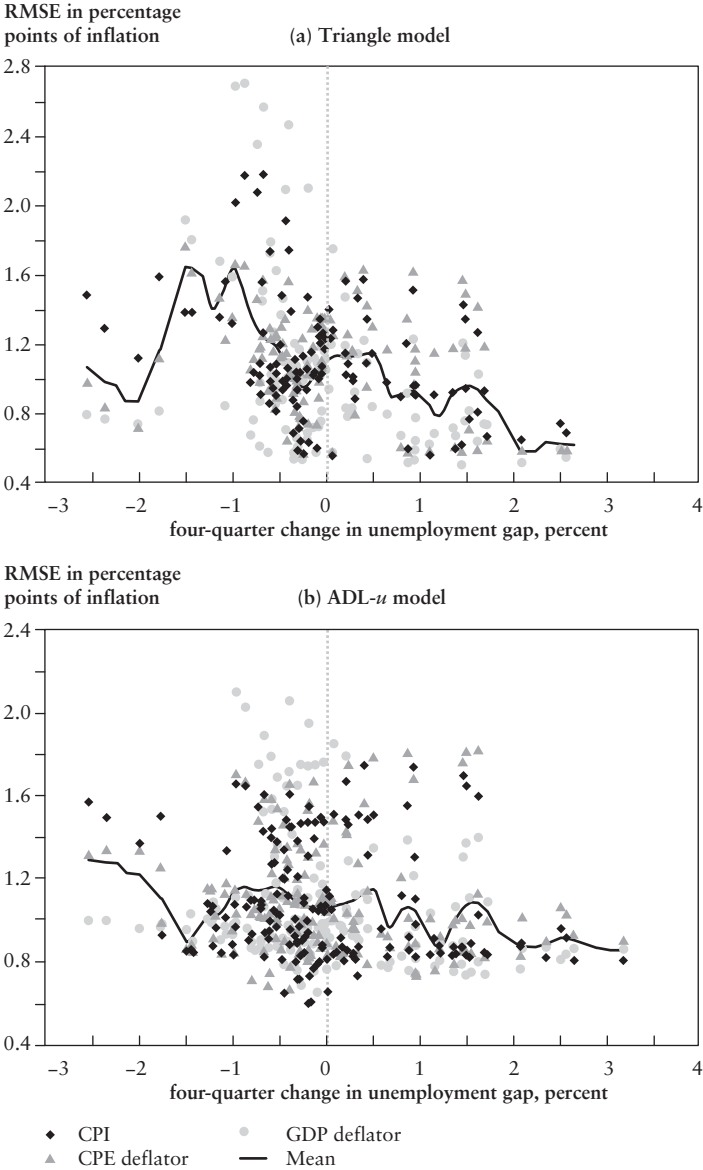


Figure 3.14 Scatterplot of Rolling Root Mean Square Errors of Headline Inflation Forecasts, Relative to Unobserved Components-Stochastic Volatility, versus the Four-quarter Change in the Unemployment Gap (two-sided bandpass)
Note: Mean is kernel regression estimate using data for all three series. Each point represents a quarter.





Figure 3.15

Scatterplot of Rolling Root Mean Square Errors of CPI-all Inflation Forecasts from (a) Triangle Model and (b) ADL- u Model, Relative to Unobserved Components-Stochastic Volatility Model, versus the Four-quarter Change in Four-quarter Inflation. Each point represents a quarter.





Figure 3.16
Scatterplot of Rolling Root Mean Square Errors of PCE-all Inflation Forecasts from (a) Triangle Model and (b) ADL- u Model, Relative to Unobserved Components-Stochastic Volatility Model, versus the Four-quarter Change in Four-quarter Inflation. Each point represents a quarter.




Figure 3.17
Scatterplot of Rolling Root Mean Square Errors of GDP Inflation Forecasts from (a) Triangle Model and (b) ADL- u Model, Relative to Unobserved Components-Stochastic Volatility Model, versus the Four-quarter Change in Four-quarter Inflation. Each point represents a quarter.

triangle model applied to the GDP deflator, the episodes of best performance are not associated with large changes in inflation.

As presented here, these patterns cannot yet be used to improve forecasts: the sharpest patterns are ones that appear using two-sided gaps. Still, these results point to a possible resolution of the Atkeson-Ohanian conundrum in which real economic activity seems to play no role in inflation forecasting. The results here suggest that, if times are quiet—if the unemployment rate is close to the NAIRU—then in fact one is better off using a univariate forecast than introducing additional estimation error by making a multivariate forecast. But if the unemployment rate is far from the NAIRU, then knowledge of that large unemployment gap is useful for inflation forecasting.

■ *We thank Ian Dew-Becker for research assistance, Michelle Barnes of the Federal Reserve Bank of Boston for data assistance, and an anonymous referee for some important references. This research was funded in part by National Science Foundation grant SBR-0617811. Data and replication files are available at <http://www.princeton.edu/~mwatson>.*

Table 3.1

Root Mean Squared Errors for Inflation Forecasting Models by Subperiod, Relative to the Unobserved Components-Stochastic Volatility Model: CPI-all

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
Number of observations	32	36	32	32	32	25
Root MSE of UC-SV forecast	0.82	1.99	2.35	1.39	0.68	1.05
Forecasting model and relative root mean square errors						
<i>Univariate forecasts</i>						
UC-SV	1.00	1.00	1.00	1.00	1.00	1.00
AR(AIC)_rec	.	1.09	1.05	1.12	1.03	1.39
AR(AIC)_iter_rec	.	1.06	1.00	1.12	1.02	1.43
AR(BIC)_rec	.	1.10	1.03	1.10	1.03	1.37
AO	1.01	1.23	1.12	1.00	1.10	1.14
MA(1)_rec	.	1.07	1.01	1.07	1.03	1.37
AR(4)_rec	.	1.12	1.02	1.13	1.02	1.42
AR(AIC)_roll	.	1.10	1.09	1.03	1.21	1.30
AR(AIC)_iter_roll	.	1.08	1.03	1.15	1.11	1.37
AR(BIC)_roll	.	1.09	1.08	1.02	1.14	1.32
AR(4)_roll	.	1.19	1.06	1.07	1.17	1.29
AR(24)_iter	.	.	.	1.30	1.04	1.33
AR(24)_iter_nocon	.	.	1.18	1.25	1.00	1.32
MA(1)_roll	.	1.04	1.02	1.07	1.05	1.13
MA(2) - NS	0.98	1.14	1.13	0.95	1.01	1.12
MA(1), $\theta=.25$	1.12	1.01	1.00	1.11	1.06	1.52
MA(1), $\theta=.65$	0.97	1.15	1.12	0.96	1.03	1.12
<i>Single-predictor ADL forecasts</i>						
UR(Level)_AIC_rec	.	0.96	0.92	0.98	1.28	1.36
UR(Dif)_AIC_rec	.	0.93	0.94	1.04	1.22	1.39
UR(1sdBP)_AIC_rec	.	0.96	0.95	1.00	1.22	1.38
GDP(Dif)_AIC_rec	.	0.88	0.93	1.00	1.09	1.36
GDP(1sdBP)_AIC_rec	.	1.03	0.90	1.00	1.08	1.34
IP(Dif)_AIC_rec	.	0.89	0.93	1.02	1.22	1.43
IP(1sdBP)_AIC_rec	.	0.95	0.93	1.01	1.17	1.40
Emp(Dif)_AIC_rec	.	0.93	0.86	1.01	1.06	1.53
Emp(1sdBP)_AIC_rec	.	0.95	0.87	1.02	1.14	1.49
CapU(Level)_AIC_rec	.	.	.	1.03	1.39	1.56
CapU((Dif)_AIC_rec	.	.	.	1.03	1.30	1.45
CapU(1sdBP)_AIC_rec	.	.	.	0.99	1.21	1.35
HPerm(Level)_AIC_rec	.	.	0.79	1.12	1.14	1.75
HPerm((Dif)_AIC_rec	.	.	0.91	1.29	0.97	1.67
HPerm(1sdBP)_AIC_rec	.	.	0.90	1.02	1.08	1.37
CFNAI(Dif)_AIC_rec	.	.	.	1.01	1.21	1.57
CFNAI(1sdBP)_AIC_rec	.	.	.	0.98	1.18	1.42
UR_5wk(Level)_AIC_rec	.	1.06	0.93	1.05	1.73	1.38
UR_5wk(Dif)_AIC_rec	.	0.94	0.91	1.07	1.34	1.40
UR_5wk(1sdBP)_AIC_rec	.	0.97	0.90	1.06	1.34	1.31
AHE(Dif)_AIC_rec	.	.	1.10	1.19	1.03	1.48

Table 3.1 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
AHE(1sdbP)_AIC_rec	.	.	1.12	1.20	1.01	1.46
RealAHE(Dif)_AIC_rec	.	.	1.10	1.19	1.03	1.48
RealAHE(1sdbP)_AIC_rec	.	.	1.12	1.20	1.01	1.46
LaborShare(Level)_AIC_rec	.	1.06	1.02	1.21	1.76	1.44
LaborShare(Dif)_AIC_rec	.	1.08	1.03	1.12	1.06	1.36
ULaborShare(1sdbP)_AIC_rec	.	1.10	1.01	1.09	1.30	1.36
CPI_Med(Level)_AIC_rec	.	.	.	1.34	1.39	1.54
CPI_Med(Dif)_AIC_rec	.	.	.	1.20	1.11	1.45
CPI_TrMn(Level)_AIC_rec	.	.	.	1.35	1.46	1.47
CPI_TrMn(Dif)_AIC_rec	.	.	.	1.10	1.07	1.45
ExRate(Dif)_AIC_rec	.	.	.	1.43	1.26	1.21
ExRate(1sdbP)_AIC_rec	.	.	.	1.82	1.04	1.28
tb_spr_AIC_rec	.	1.10	1.05	1.21	1.24	1.56
UR(Level)_AIC_rol	.	1.20	1.13	0.99	1.32	1.30
UR(Dif)_AIC_rol	.	1.07	1.00	1.04	1.23	1.28
UR(1sdbP)_AIC_rol	.	1.17	1.07	1.03	1.28	1.30
GDP(Dif)_AIC_rol	.	1.01	1.01	0.98	1.36	1.25
GDP(1sdbP)_AIC_rol	.	1.10	0.91	1.00	1.25	1.25
IP(Dif)_AIC_rol	.	0.95	0.99	1.05	1.26	1.33
IP(1sdbP)_AIC_rol	.	1.07	1.00	1.05	1.30	1.28
Emp(Dif)_AIC_rol	.	1.06	0.97	0.99	1.23	1.24
Emp(1sdbP)_AIC_rol	.	1.19	0.91	1.02	1.26	1.31
CapU(Level)_AIC_rol	.	.	.	0.98	1.38	1.33
CapU(Dif)_AIC_rol	.	.	.	1.02	1.27	1.29
CapU(1sdbP)_AIC_rol	.	.	.	0.97	1.35	1.22
HPerm(Level)_AIC_rol	.	.	0.75	1.27	1.23	1.41
HPerm(Dif)_AIC_rol	.	.	1.14	1.16	1.05	1.55
HPerm(1sdbP)_AIC_rol	.	.	0.94	1.21	1.20	1.32
CFNAI(Dif)_AIC_rol	.	.	.	0.97	1.28	1.25
CFNAI(1sdbP)_AIC_rol	.	.	.	1.02	1.28	1.25
UR_5wk(Level)_AIC_rol	.	1.19	0.93	1.18	1.60	1.34
UR_5wk(Dif)_AIC_rol	.	1.03	0.97	1.03	1.41	1.27
UR_5wk(1sdbP)_AIC_rol	.	1.08	0.85	1.06	1.45	1.31
AHE(Dif)_AIC_rol	.	.	1.10	1.08	1.38	1.24
AHE(1sdbP)_AIC_rol	.	.	1.12	1.07	1.33	1.19
RealAHE(Dif)_AIC_rol	.	.	1.10	1.08	1.38	1.24
RealAHE(1sdbP)_AIC_rol	.	.	1.12	1.07	1.33	1.19
LaborShare(Level)_AIC_rol	.	1.15	1.02	1.12	1.31	1.32
LaborShare(Dif)_AIC_rol	.	1.13	1.09	1.02	1.63	1.32
ULaborShare(1sdbP)_AIC_rol	.	1.18	1.09	1.04	1.31	1.30
CPI_Med(Level)_AIC_rol	.	.	.	1.15	1.34	1.18
CPI_Med(Dif)_AIC_rol	.	.	.	1.01	1.15	1.29
CPI_TrMn(Level)_AIC_rol	.	.	.	1.12	1.38	1.28
CPI_TrMn(Dif)_AIC_rol	.	.	.	1.05	1.15	1.31
ExRate(Dif)_AIC_rol	.	.	.	1.53	1.20	1.28
ExRate(1sdbP)_AIC_rol	.	.	.	1.91	1.16	1.34

Table 3.1 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
tb_spr_AIC_roll	.	1.13	1.23	1.33	1.42	1.37
UR(Level)_BIC_rec	.	0.92	0.91	0.98	1.28	1.36
UR(Dif)_BIC_rec	.	0.88	0.94	1.06	1.16	1.35
UR(1sdbP)_BIC_rec	.	0.91	0.93	0.96	1.17	1.35
GDP(Dif)_BIC_rec	.	0.95	0.99	1.00	1.09	1.36
GDP(1sdbP)_BIC_rec	.	0.99	0.95	0.99	1.05	1.33
IP(Dif)_BIC_rec	.	0.90	0.97	1.03	1.20	1.41
IP(1sdbP)_BIC_rec	.	0.95	0.99	1.00	1.11	1.39
Emp(Dif)_BIC_rec	.	0.90	0.92	0.98	1.05	1.51
Emp(1sdbP)_BIC_rec	.	0.93	0.93	0.99	1.09	1.45
CapU(Level)_BIC_rec	.	.	.	1.02	1.29	1.56
CapU((Dif)_BIC_rec	.	.	.	1.07	1.30	1.46
CapU(1sdbP)_BIC_rec	.	.	.	0.97	1.17	1.30
HPerm(Level)_BIC_rec	.	.	0.82	1.06	1.14	1.75
HPerm((Dif)_BIC_rec	.	.	1.05	1.32	0.97	1.65
HPerm(1sdbP)_BIC_rec	.	.	0.93	1.02	1.08	1.37
CFNAI(Dif)_BIC_rec	.	.	.	0.92	1.18	1.44
CFNAI(1sdbP)_BIC_rec	.	.	.	0.95	1.18	1.42
UR_5wk(Level)_BIC_rec	.	1.03	0.92	1.13	1.62	1.48
UR_5wk(Dif)_BIC_rec	.	0.94	0.96	1.15	1.17	1.49
UR_5wk(1sdbP)_BIC_rec	.	0.94	0.88	1.11	1.27	1.34
AHE(Dif)_BIC_rec	.	.	1.08	1.19	1.10	1.42
AHE(1sdbP)_BIC_rec	.	.	1.10	1.23	1.05	1.37
RealAHE(Dif)_BIC_rec	.	.	1.08	1.19	1.10	1.42
RealAHE(1sdbP)_BIC_rec	.	.	1.10	1.23	1.05	1.37
LaborShare(Level)_BIC_rec	.	1.02	0.99	1.20	1.61	1.44
LaborShare(Dif)_BIC_rec	.	1.08	1.03	1.13	1.07	1.36
ULaborShare(1sdbP)_BIC_rec	.	1.07	0.97	1.13	1.30	1.40
CPI_Med(Level)_BIC_rec	.	.	.	1.22	1.44	1.53
CPI_Med(Dif)_BIC_rec	.	.	.	1.21	1.14	1.51
CPI_TrMn(Level)_BIC_rec	.	.	.	1.23	1.43	1.49
CPI_TrMn(Dif)_BIC_rec	.	.	.	1.10	1.07	1.49
ExRate(Dif)_BIC_rec	.	.	.	1.53	1.19	1.32
ExRate(1sdbP)_BIC_rec	.	.	.	1.87	1.09	1.28
tb_spr_BIC_rec	.	1.09	1.09	1.17	1.06	1.40
UR(Level)_BIC_rol	.	1.16	1.05	0.99	1.33	1.31
UR(Dif)_BIC_rol	.	0.99	0.99	0.99	1.19	1.35
UR(1sdbP)_BIC_rol	.	1.14	0.98	0.96	1.24	1.28
GDP(Dif)_BIC_rol	.	0.96	1.01	0.98	1.28	1.31
GDP(1sdbP)_BIC_rol	.	1.04	0.92	0.98	1.18	1.30
IP(Dif)_BIC_rol	.	0.99	1.01	1.05	1.24	1.29
IP(1sdbP)_BIC_rol	.	1.08	0.96	0.97	1.31	1.32
Emp(Dif)_BIC_rol	.	1.06	0.95	1.02	1.22	1.27
Emp(1sdbP)_BIC_rol	.	1.12	0.92	1.05	1.24	1.27
CapU(Level)_BIC_rol	.	.	.	0.97	1.30	1.30
CapU(Dif)_BIC_rol	.	.	.	1.01	1.26	1.27

Table 3.1 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
CapU(1sdBP)_BIC_rol	.	.	.	0.93	1.23	1.25
HPerm(Level)_BIC_rol	.	.	0.77	1.25	1.22	1.43
HPerm((Dif)_BIC_rol	.	.	1.09	1.21	1.06	1.36
HPerm(1sdBP)_BIC_rol	.	.	0.92	1.21	1.19	1.33
CFNAI(Dif)_BIC_rol	.	.	.	0.96	1.26	1.32
CFNAI(1sdBP)_BIC_rol	.	.	.	1.02	1.22	1.24
UR_5wk(Level)_BIC_rol	.	1.18	0.95	1.19	1.35	1.43
UR_5wk(Dif)_BIC_rol	.	0.96	0.99	1.10	1.19	1.35
UR_5wk(1sdBP)_BIC_rol	.	1.04	0.93	1.09	1.32	1.37
AHE(Dif)_BIC_rol	.	.	1.09	1.03	1.19	1.38
AHE(1sdBP)_BIC_rol	.	.	1.10	1.16	1.12	1.23
RealAHE(Dif)_BIC_rol	.	.	1.09	1.03	1.19	1.38
RealAHE(1sdBP)_BIC_rol	.	.	1.10	1.16	1.12	1.23
LaborShare(Level)_BIC_rol	.	1.09	1.05	1.05	1.23	1.33
LaborShare(Dif)_BIC_rol	.	1.12	1.10	1.11	1.17	1.38
ULaborShare(1sdBP)_BIC_rol	.	1.11	1.06	1.02	1.29	1.36
CPI_Med(Level)_BIC_rol	.	.	.	1.07	1.28	1.24
CPI_Med(Dif)_BIC_rol	.	.	.	1.04	1.15	1.30
CPI_TrMn(Level)_BIC_rol	.	.	.	1.02	1.28	1.30
CPI_TrMn(Dif)_BIC_rol	.	.	.	1.07	1.14	1.36
ExRate(Dif)_BIC_rol	.	.	.	1.56	1.13	1.34
ExRate(1sdBP)_BIC_rol	.	.	.	1.95	1.11	1.36
tb_spr_BIC_rol	.	1.16	1.25	1.32	1.22	1.38
<i>Triangle model forecasts</i>						
Triangle Constant NAIRU	.	.	0.94	1.11	1.14	1.11
Triangle TV NAIRU	.	.	0.95	1.15	1.07	1.16
Triangle Constant NAIRU (no z)	.	.	1.02	1.19	1.34	1.34
Triangle TV NAIRU (no z)	.	.	1.12	1.23	1.10	1.52
<i>Combination forecasts</i>						
Activity Median Combining	.	0.96	0.88	0.96	1.13	1.30
Activity Mean Combining	.	0.97	0.86	0.96	1.11	1.30
Activity Tr. Mean Combining	.	0.97	0.87	0.96	1.11	1.30
Activity MSE(A) Combining	.	.	0.86	0.97	1.12	1.31
Activity MSE(B) Combining	.	.	0.86	0.96	1.12	1.31
Activity MSE(C) Combining	.	.	0.86	0.96	1.11	1.30
Activity MSE(D) Combining	.	.	0.86	0.98	1.14	1.33
Activity MSE(E) Combining	.	.	0.87	0.97	1.13	1.32
Activity MSE(F) Combining	.	.	0.87	0.96	1.12	1.30
Activity Rec. Best(4q) Combining	.	1.12	0.74	0.99	1.38	1.56
Activity Rec. Best(8q) Combining	.	1.07	0.90	1.22	1.48	1.36
OtherADL Median Combining	.	1.07	1.06	1.03	1.11	1.29
OtherADL Mean Combining	.	1.08	1.01	1.06	1.09	1.30
OtherADL Tr. Mean Combining	.	1.08	1.03	1.05	1.09	1.31
OtherADL MSE(A) Combining	.	.	0.98	1.07	1.11	1.30
OtherADL MSE(B) Combining	.	.	0.98	1.07	1.12	1.30
OtherADL MSE(C) Combining	.	.	0.99	1.07	1.12	1.30

Table 3.1 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
OtherADL MSE(D) Combining .	.	.	0.98	1.08	1.13	1.31
OtherADL MSE(E) Combining .	.	.	0.99	1.07	1.13	1.30
OtherADL MSE(F) Combining .	.	.	0.99	1.07	1.14	1.30
OtherADL Rec. Best(4q) Combining .	1.13	1.13	1.05	1.12	1.36	1.37
OtherADL Rec. Best(8q) Combining .	1.14	1.14	1.09	1.21	1.30	1.42
All Median Combining .	0.98	0.98	0.92	0.98	1.10	1.31
All Mean Combining .	0.99	0.99	0.89	0.98	1.07	1.29
All Tr. Mean Combining .	0.99	0.99	0.90	0.98	1.08	1.30
All MSE(A) Combining .	.	.	0.87	0.99	1.10	1.30
All MSE(B) Combining .	.	.	0.87	0.98	1.10	1.30
All MSE(C) Combining .	.	.	0.87	0.98	1.09	1.29
All MSE(D) Combining .	.	.	0.87	1.00	1.12	1.31
All MSE(E) Combining .	.	.	0.88	0.99	1.12	1.30
All MSE(F) Combining .	.	.	0.88	0.98	1.10	1.29
All Rec. Best(4q) Combining .	1.12	1.12	0.74	1.11	1.47	1.63
All Rec. Best(8q) Combining .	1.08	1.08	0.92	1.19	1.51	1.43
UCSV and Triangle Rec. Best(4q) Combining	1.02	1.05	1.01
UCSV and Triangle Rec. Best(8q) Combining	1.06	1.05	1.11

Notes to Table 3.1: Entries are Root Mean Squared Errors, relative to the Root Mean Squared Errors of the Unobserved Components-Stochastic Volatility model, over the indicated sample period. Blanks indicate insufficient data to compute forecasts over the indicated subsample. The abbreviations denote:

_AIC: AIC lag selection, up to six lags (for ADL models, AIC over the two lag lengths separately)

_BIC: BIC lag selection, up to six lags (for ADL models, AIC over the two lag lengths separately)

_rec: recursive estimation

_roll: rolling estimation

Level: indicated predictor appears in levels

Dif: indicated predictor appears in log differences

1sdbP: indicated predictor appears in gap form, computed using 1-sided bandpass filter as discussed in the text

Triangle: Triangle model or TV-triangle model, with or without supply shock (“z”) variables

mean, median, trimmed mean: forecast combining methods, for the indicated group of forecasts

MSE(A) – MSE(F): MSE-based combining as indicated in equations (15)–(20).

Best (four-quarter) and Best (eight-quarter): recently best forecast based on cumulative MSE over past four (or eight) quarters

UCSV and Triangle Rec. Best (four-quarter) and (eight-quarter) Combining: best of UC-SV and triangle models (constant NAIRU) based on cumulative MSE over past four (or eight) quarters

nocon: constant term is suppressed

Source: Authors’ calculations.

Table 3.2

Root Mean Squared Errors for Inflation Forecasting Models by Subperiod, Relative to the Unobserved Components-Stochastic Volatility Model: CPI-core

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
Number of observations	32	36	32	32	32	25
Root MSE of UC-SV forecast	0.82	2.15	2.30	0.58	0.31	0.53
Forecasting model and relative root mean square errors						
<i>Univariate forecasts</i>						
UC-SV	1.00	1.00	1.00	1.00	1.00	1.00
AR(AIC)_rec	.	.	1.07	1.07	1.04	1.05
AR(AIC)_iter_rec	.	.	1.08	1.06	1.04	1.06
AR(BIC)_rec	.	.	1.07	1.01	1.06	1.05
AO	1.03	1.14	1.01	1.08	1.04	1.06
MA(1)_rec	.	1.05	1.04	1.00	1.01	1.04
AR(4)_rec	.	.	1.09	1.09	1.03	1.04
AR(AIC)_roll	.	.	1.12	1.05	1.12	1.09
AR(AIC)_iter_roll	.	.	1.21	1.15	1.15	1.11
AR(BIC)_roll	.	.	1.11	1.03	1.11	1.09
AR(4)_roll	.	.	1.15	1.06	1.12	1.10
AR(24)_iter	.	.	1.24	1.57	1.51	0.93
AR(24)_iter_nocon	.	.	1.10	1.23	1.32	0.91
MA(1)_roll	.	1.04	1.03	1.12	1.00	1.07
MA(2) - NS	1.04	1.04	1.01	1.04	1.07	0.98
MA(1), $\theta = .25$	1.05	1.01	0.98	1.02	1.04	1.07
MA(1), $\theta = .65$	1.03	1.06	1.00	1.03	1.04	1.00
<i>Single-predictor ADL forecasts</i>						
UR(Level)_AIC_rec	.	.	0.89	0.83	1.92	1.11
UR(Dif)_AIC_rec	.	.	0.95	0.91	1.43	1.01
UR(1sdBP)_AIC_rec	.	.	0.91	1.01	1.63	1.05
GDP(Dif)_AIC_rec	.	.	1.02	0.91	1.02	0.92
GDP(1sdBP)_AIC_rec	.	.	0.95	1.00	1.17	1.09
IP(Dif)_AIC_rec	.	.	1.03	1.00	1.25	1.16
IP(1sdBP)_AIC_rec	.	.	0.97	0.85	1.54	1.39
Emp(Dif)_AIC_rec	.	.	0.92	0.90	1.18	1.32
Emp(1sdBP)_AIC_rec	.	.	0.91	0.90	1.28	1.27
CapU(Level)_AIC_rec	.	.	.	1.24	2.00	2.19
CapU((Dif)_AIC_rec	.	.	.	1.12	1.21	1.25
CapU(1sdBP)_AIC_rec	.	.	.	1.34	1.26	1.20
HPerm(Level)_AIC_rec	.	.	0.91	1.29	1.46	1.91
HPerm((Dif)_AIC_rec	.	.	1.06	1.10	1.21	1.04
HPerm(1sdBP)_AIC_rec	.	.	0.99	1.07	1.48	0.98
CFNAI(Dif)_AIC_rec	.	.	.	1.16	1.27	1.39
CFNAI(1sdBP)_AIC_rec	.	.	.	1.17	1.29	1.37
UR_5wk(Level)_AIC_rec	.	.	0.86	1.10	3.09	1.32
UR_5wk(Dif)_AIC_rec	.	.	0.91	1.19	1.91	1.08
UR_5wk(1sdBP)_AIC_rec	.	.	0.90	1.22	2.32	1.06
AHE(Dif)_AIC_rec	.	.	1.12	1.08	1.03	1.06

Table 3.2 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
AHE(1sdbp)_AIC_rec	.	.	1.16	1.36	1.10	1.10
RealAHE(Dif)_AIC_rec	.	.	1.12	1.08	1.03	1.06
RealAHE(1sdbp)_AIC_rec	.	.	1.16	1.36	1.10	1.10
LaborShare(Level)_AIC_rec	.	.	1.11	1.37	2.18	1.12
LaborShare(Dif)_AIC_rec	.	.	1.14	1.05	1.27	1.06
ULaborShare(1sdbp)_AIC_rec	.	.	1.14	1.15	1.58	1.07
CPI_Med(Level)_AIC_rec	.	.	.	1.30	2.06	1.10
CPI_Med(Dif)_AIC_rec	.	.	.	1.14	1.81	1.25
CPI_TrMn(Level)_AIC_rec	.	.	.	1.28	1.69	1.23
CPI_TrMn(Dif)_AIC_rec	.	.	.	1.11	1.38	1.20
ExRate(Dif)_AIC_rec	.	.	.	2.97	1.43	0.93
ExRate(1sdbp)_AIC_rec	.	.	.	3.14	1.25	1.24
tb_spr_AIC_rec	.	.	1.10	1.52	2.53	1.32
UR(Level)_AIC_rol	.	.	1.27	1.52	1.34	1.19
UR(Dif)_AIC_rol	.	.	1.02	1.26	1.07	1.07
UR(1sdbp)_AIC_rol	.	.	0.87	1.40	1.12	1.21
GDP(Dif)_AIC_rol	.	.	1.12	1.37	1.12	1.10
GDP(1sdbp)_AIC_rol	.	.	1.00	1.52	1.02	1.10
IP(Dif)_AIC_rol	.	.	1.08	1.48	1.15	1.26
IP(1sdbp)_AIC_rol	.	.	0.89	1.70	1.14	1.39
Emp(Dif)_AIC_rol	.	.	0.94	1.47	1.09	1.31
Emp(1sdbp)_AIC_rol	.	.	0.84	1.57	1.08	1.54
CapU(Level)_AIC_rol	.	.	.	1.59	1.31	1.39
CapU((Dif)_AIC_rol	.	.	.	1.48	1.11	1.17
CapU(1sdbp)_AIC_rol	.	.	.	1.48	1.26	1.26
HPerm(Level)_AIC_rol	.	.	0.89	2.35	1.04	1.24
HPerm((Dif)_AIC_rol	.	.	1.10	1.63	1.12	1.14
HPerm(1sdbp)_AIC_rol	.	.	1.04	1.95	1.05	1.12
CFNAI(Dif)_AIC_rol	.	.	.	1.44	1.03	1.17
CFNAI(1sdbp)_AIC_rol	.	.	.	1.50	0.92	1.25
UR_5wk(Level)_AIC_rol	.	.	1.05	2.12	1.28	1.13
UR_5wk(Dif)_AIC_rol	.	.	1.09	1.44	1.08	1.10
UR_5wk(1sdbp)_AIC_rol	.	.	0.85	1.32	1.33	1.12
AHE(Dif)_AIC_rol	.	.	1.23	1.23	1.13	1.16
AHE(1sdbp)_AIC_rol	.	.	1.22	1.53	1.12	1.15
RealAHE(Dif)_AIC_rol	.	.	1.23	1.23	1.13	1.16
RealAHE(1sdbp)_AIC_rol	.	.	1.22	1.53	1.12	1.15
LaborShare(Level)_AIC_rol	.	.	1.20	1.34	1.83	1.12
LaborShare(Dif)_AIC_rol	.	.	1.41	1.15	1.97	1.10
ULaborShare(1sdbp)_AIC_rol	.	.	1.30	1.16	1.69	1.10
CPI_Med(Level)_AIC_rol	.	.	.	1.37	1.69	0.80
CPI_Med(Dif)_AIC_rol	.	.	.	1.04	1.25	1.15
CPI_TrMn(Level)_AIC_rol	.	.	.	1.30	1.53	1.19
CPI_TrMn(Dif)_AIC_rol	.	.	.	1.08	1.14	1.19
ExRate(Dif)_AIC_rol	.	.	.	3.48	1.19	1.11
ExRate(1sdbp)_AIC_rol	.	.	.	3.58	1.05	1.10

Table 3.2 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
tb_spr_AIC_rol	.	.	1.15	1.61	1.22	1.12
UR(Level)_BIC_rec	.	.	0.97	0.83	1.92	1.11
UR(Dif)_BIC_rec	.	.	1.00	0.91	1.43	1.01
UR(1sdBP)_BIC_rec	.	.	0.88	1.01	1.62	1.05
GDP(Dif)_BIC_rec	.	.	1.04	0.98	1.06	0.91
GDP(1sdBP)_BIC_rec	.	.	0.96	0.96	1.14	1.07
IP(Dif)_BIC_rec	.	.	1.04	0.99	1.11	1.11
IP(1sdBP)_BIC_rec	.	.	1.00	0.90	1.41	1.34
Emp(Dif)_BIC_rec	.	.	0.97	0.90	1.18	1.32
Emp(1sdBP)_BIC_rec	.	.	0.83	0.90	1.28	1.27
CapU(Level)_BIC_rec	.	.	.	1.24	1.83	2.10
CapU((Dif)_BIC_rec	.	.	.	1.06	1.11	1.22
CapU(1sdBP)_BIC_rec	.	.	.	1.22	1.23	1.27
HPerm(Level)_BIC_rec	.	.	0.91	1.27	1.46	1.91
HPerm((Dif)_BIC_rec	.	.	1.05	1.16	1.21	1.04
HPerm(1sdBP)_BIC_rec	.	.	0.99	1.07	1.48	0.98
CFNAI(Dif)_BIC_rec	.	.	.	1.10	1.27	1.29
CFNAI(1sdBP)_BIC_rec	.	.	.	1.17	1.29	1.37
UR_5wk(Level)_BIC_rec	.	.	0.93	1.09	2.88	1.42
UR_5wk(Dif)_BIC_rec	.	.	1.01	1.19	1.94	1.15
UR_5wk(1sdBP)_BIC_rec	.	.	0.87	1.24	2.23	1.05
AHE(Dif)_BIC_rec	.	.	1.13	1.03	1.05	1.09
AHE(1sdBP)_BIC_rec	.	.	1.17	1.24	1.16	1.09
RealAHE(Dif)_BIC_rec	.	.	1.13	1.03	1.05	1.09
RealAHE(1sdBP)_BIC_rec	.	.	1.17	1.24	1.16	1.09
LaborShare(Level)_BIC_rec	.	.	1.08	1.22	1.68	1.09
LaborShare(Dif)_BIC_rec	.	.	1.09	1.02	1.34	1.08
ULaborShare(1sdBP)_BIC_rec	.	.	1.09	1.08	1.27	1.02
CPI_Med(Level)_BIC_rec	.	.	.	1.34	1.66	1.11
CPI_Med(Dif)_BIC_rec	.	.	.	1.02	1.25	1.10
CPI_TrMn(Level)_BIC_rec	.	.	.	1.28	1.71	1.23
CPI_TrMn(Dif)_BIC_rec	.	.	.	1.09	1.38	1.20
ExRate(Dif)_BIC_rec	.	.	.	1.73	1.34	1.09
ExRate(1sdBP)_BIC_rec	.	.	.	2.85	1.10	1.01
tb_spr_BIC_rec	.	.	1.07	1.50	2.40	1.32
UR(Level)_BIC_rol	.	.	1.16	1.49	1.26	1.16
UR(Dif)_BIC_rol	.	.	1.04	1.18	1.06	1.08
UR(1sdBP)_BIC_rol	.	.	0.88	1.33	1.15	1.10
GDP(Dif)_BIC_rol	.	.	1.01	1.30	1.11	1.04
GDP(1sdBP)_BIC_rol	.	.	0.95	1.47	1.10	1.11
IP(Dif)_BIC_rol	.	.	1.06	1.49	1.15	1.23
IP(1sdBP)_BIC_rol	.	.	0.86	1.67	1.13	1.31
Emp(Dif)_BIC_rol	.	.	1.00	1.45	1.15	1.16
Emp(1sdBP)_BIC_rol	.	.	0.84	1.56	1.15	1.50
CapU(Level)_BIC_rol	.	.	.	1.61	1.32	1.37
CapU(Dif)_BIC_rol	.	.	.	1.48	1.11	1.15

Table 3.2 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
CapU(1sdBP)_BIC_roll	.	.	.	1.46	1.30	1.29
HPerm(Level)_BIC_roll	.	.	0.89	2.29	1.05	1.17
HPerm(Dif)_BIC_roll	.	.	1.10	1.62	1.20	1.13
HPerm(1sdBP)_BIC_roll	.	.	1.05	1.93	1.08	1.11
CFNAI(Dif)_BIC_roll	.	.	.	1.47	1.03	1.13
CFNAI(1sdBP)_BIC_roll	.	.	.	1.53	0.91	1.20
UR_5wk(Level)_BIC_roll	.	.	1.08	1.96	1.19	1.12
UR_5wk(Dif)_BIC_roll	.	.	1.11	1.49	1.11	1.09
UR_5wk(1sdBP)_BIC_roll	.	.	0.85	1.36	1.31	1.08
AHE(Dif)_BIC_roll	.	.	1.12	1.15	1.12	1.13
AHE(1sdBP)_BIC_roll	.	.	1.16	1.32	1.13	1.16
RealAHE(Dif)_BIC_roll	.	.	1.12	1.15	1.12	1.13
RealAHE(1sdBP)_BIC_roll	.	.	1.16	1.32	1.13	1.16
LaborShare(Level)_BIC_roll	.	.	1.10	1.35	1.61	1.11
LaborShare(Dif)_BIC_roll	.	.	1.20	1.21	1.84	1.10
ULaborShare(1sdBP)_BIC_roll	.	.	1.13	1.12	1.56	1.10
CPI_Med(Level)_BIC_roll	.	.	.	1.41	1.47	0.84
CPI_Med(Dif)_BIC_roll	.	.	.	1.05	1.22	1.18
CPI_TrMn(Level)_BIC_roll	.	.	.	1.23	1.47	1.19
CPI_TrMn(Dif)_BIC_roll	.	.	.	1.08	1.15	1.18
ExRate(Dif)_BIC_roll	.	.	.	3.03	1.10	1.14
ExRate(1sdBP)_BIC_roll	.	.	.	3.35	1.06	1.06
tb_spr_BIC_roll	.	.	1.14	1.45	1.28	1.11
<i>Triangle model forecasts</i>						
Triangle Constant NAIRU	.	.	1.32	1.50	1.81	1.44
Triangle TV NAIRU	.	.	1.32	1.46	1.48	1.39
Triangle Constant NAIRU (no z)	.	.	1.05	1.11	2.34	1.17
Triangle TV NAIRU (no z)	.	.	1.07	1.22	1.63	1.23
<i>Combination forecasts</i>						
Activity Median Combining	.	.	0.86	0.86	1.00	1.07
Activity Mean Combining	.	.	0.86	0.89	1.02	1.02
Activity Tr. Mean Combining	.	.	0.86	0.88	1.01	1.04
Activity MSE(A) Combining	.	.	.	0.87	1.05	1.05
Activity MSE(B) Combining	.	.	.	0.88	1.06	1.04
Activity MSE(C) Combining	.	.	.	0.88	1.06	1.04
Activity MSE(D) Combining	.	.	.	0.87	1.09	1.07
Activity MSE(E) Combining	.	.	.	0.87	1.10	1.06
Activity MSE(F) Combining	.	.	.	0.88	1.11	1.05
Activity Rec. Best(4q) Combining	.	.	1.13	1.15	1.22	1.19
Activity Rec. Best(8q) Combining	.	.	0.96	1.40	1.50	1.40
OtherADL Median Combining	.	.	1.08	0.99	1.11	1.06
OtherADL Mean Combining	.	.	1.05	1.04	1.15	1.02
OtherADL Tr. Mean Combining	.	.	1.07	0.96	1.12	1.03
OtherADL MSE(A) Combining	.	.	.	1.02	1.16	1.04
OtherADL MSE(B) Combining	.	.	.	1.00	1.18	1.03
OtherADL MSE(C) Combining	.	.	.	0.99	1.18	1.03

Table 3.2 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
OtherADL MSE(D) Combining	1.07	1.16	1.06
OtherADL MSE(E) Combining	1.02	1.17	1.05
OtherADL MSE(F) Combining	1.00	1.18	1.04
OtherADL Rec. Best(4q) Combining .	.	.	1.08	2.06	1.18	1.06
OtherADL Rec. Best(8q) Combining .	.	.	1.07	1.54	1.10	1.38
All Median Combining .	.	.	0.94	0.83	1.01	1.04
All Mean Combining .	.	.	0.91	0.85	1.00	0.98
All Tr. Mean Combining .	.	.	0.92	0.82	1.00	1.00
All MSE(A) Combining	0.86	1.03	1.02
All MSE(B) Combining	0.84	1.03	1.01
All MSE(C) Combining	0.84	1.02	1.01
All MSE(D) Combining	0.88	1.05	1.04
All MSE(E) Combining	0.85	1.05	1.03
All MSE(F) Combining	0.84	1.04	1.02
All Rec. Best(4q) Combining .	.	.	1.19	1.53	1.17	1.08
All Rec. Best(8q) Combining .	.	.	0.96	1.67	1.45	1.50
UCSV and Triangle Rec. Best(4q) Combining	1.37	1.06	1.00
UCSV and Triangle Rec. Best(8q) Combining	1.02	1.16	1.09

Source: Authors' calculations.

Table 3.3

Root Mean Squared Errors for Inflation Forecasting Models by Subperiod, Relative to the Unobserved Components-Stochastic Volatility Model: PCE-all

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
Number of observations	32	36	32	32	32	25
Root MSE of UC-SV forecast	0.73	1.83	1.41	0.88	0.59	0.72
Forecasting model and relative root mean square errors						
<i>Univariate forecasts</i>						
UC-SV	1.00	1.00	1.00	1.00	1.00	1.00
AR(AIC)_rec	.	1.14	1.02	1.15	1.06	1.45
AR(AIC)_iter_rec	.	1.04	1.01	1.14	1.04	1.50
AR(BIC)_rec	.	1.12	1.02	1.15	1.07	1.58
AO	1.02	1.20	1.18	1.01	1.10	1.09
MA(1)_rec	.	1.08	1.00	1.08	1.04	1.42
AR(4)_rec	.	1.16	1.03	1.13	1.07	1.46
AR(AIC)_roll	.	1.13	1.06	1.09	1.25	1.28
AR(AIC)_iter_roll	.	1.26	1.06	1.20	1.31	1.35
AR(BIC)_roll	.	1.11	1.07	1.09	1.24	1.33
AR(4)_roll	.	1.23	1.06	1.15	1.22	1.26
AR(24)_iter	.	.	1.23	1.53	1.15	1.34
AR(24)_iter_nocon	.	.	1.10	1.42	1.12	1.33
MA(1)_roll	.	1.04	0.99	1.09	1.03	1.10
MA(2) - NS	1.01	1.09	1.20	1.00	1.01	1.15
MA(1), $\theta=.25$	1.10	1.01	0.99	1.12	1.07	1.59
MA(1), $\theta=.65$	0.99	1.12	1.18	0.99	1.02	1.14
<i>Single-predictor ADL forecasts</i>						
UR(Level)_AIC_rec	.	1.06	0.98	0.99	1.22	1.43
UR(Dif)_AIC_rec	.	1.02	1.04	1.10	1.15	1.47
UR(1sdbp)_AIC_rec	.	1.07	1.09	0.99	1.14	1.42
GDP(Dif)_AIC_rec	.	1.02	0.98	1.04	1.15	1.47
GDP(1sdbp)_AIC_rec	.	1.08	1.02	1.00	1.11	1.43
IP(Dif)_AIC_rec	.	1.01	1.00	1.04	1.24	1.51
IP(1sdbp)_AIC_rec	.	1.06	1.06	1.01	1.21	1.46
Emp(Dif)_AIC_rec	.	1.03	0.96	1.04	1.12	1.66
Emp(1sdbp)_AIC_rec	.	1.04	1.03	1.01	1.13	1.54
CapU(Level)_AIC_rec	.	.	.	1.13	1.31	1.75
CapU((Dif)_AIC_rec	.	.	.	1.21	1.31	1.70
CapU(1sdbp)_AIC_rec	.	.	.	1.13	1.22	1.50
HPerm(Level)_AIC_rec	.	.	0.96	1.05	1.07	1.74
HPerm((Dif)_AIC_rec	.	.	1.11	1.16	1.07	1.61
HPerm(1sdbp)_AIC_rec	.	.	0.99	1.03	1.07	1.43
CFNAI(Dif)_AIC_rec	.	.	.	1.15	1.18	1.76
CFNAI(1sdbp)_AIC_rec	.	.	.	1.12	1.19	1.64
UR_5wk(Level)_AIC_rec	.	1.19	0.95	1.13	1.47	1.44
UR_5wk(Dif)_AIC_rec	.	1.07	1.02	1.14	1.19	1.45
UR_5wk(1sdbp)_AIC_rec	.	1.09	0.98	1.12	1.19	1.38
AHE(Dif)_AIC_rec	.	.	1.10	1.27	1.08	1.59

Table 3.3 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
AHE(1sdbP)_AIC_rec	.	.	1.15	1.26	1.07	1.56
RealAHE(Dif)_AIC_rec	.	.	1.10	1.27	1.08	1.59
RealAHE(1sdbP)_AIC_rec	.	.	1.15	1.26	1.07	1.56
LaborShare(Level)_AIC_rec	.	1.15	0.96	1.28	1.65	1.45
LaborShare(Dif)_AIC_rec	.	1.15	1.01	1.14	1.09	1.40
ULaborShare(1sdbP)_AIC_rec	.	1.18	0.96	1.13	1.32	1.40
CPI_Med(Level)_AIC_rec	.	.	.	1.33	1.34	1.56
CPI_Med(Dif)_AIC_rec	.	.	.	1.22	1.13	1.64
CPI_TrMn(Level)_AIC_rec	.	.	.	1.28	1.34	1.50
CPI_TrMn(Dif)_AIC_rec	.	.	.	1.15	1.10	1.64
ExRate(Dif)_AIC_rec	.	.	.	1.28	1.33	1.31
ExRate(1sdbP)_AIC_rec	.	.	.	1.60	1.15	1.32
tb_spr_AIC_rec	.	1.18	1.31	1.25	1.01	1.51
UR(Level)_AIC_rol	.	1.24	1.53	1.04	1.32	1.25
UR(Dif)_AIC_rol	.	1.10	1.12	1.11	1.26	1.21
UR(1sdbP)_AIC_rol	.	1.29	1.30	1.06	1.29	1.25
GDP(Dif)_AIC_rol	.	1.07	1.14	1.16	1.33	1.19
GDP(1sdbP)_AIC_rol	.	1.24	1.12	1.02	1.33	1.21
IP(Dif)_AIC_rol	.	1.08	0.99	1.19	1.39	1.25
IP(1sdbP)_AIC_rol	.	1.20	1.18	1.14	1.43	1.37
Emp(Dif)_AIC_rol	.	1.10	1.20	1.12	1.34	1.22
Emp(1sdbP)_AIC_rol	.	1.24	1.12	1.15	1.34	1.32
CapU(Level)_AIC_rol	.	.	.	1.10	1.44	1.29
CapU(Dif)_AIC_rol	.	.	.	1.17	1.37	1.30
CapU(1sdbP)_AIC_rol	.	.	.	1.01	1.43	1.30
HPerm(Level)_AIC_rol	.	.	1.00	1.14	1.25	1.33
HPerm(Dif)_AIC_rol	.	.	1.19	1.08	1.11	1.88
HPerm(1sdbP)_AIC_rol	.	.	1.07	1.15	1.28	1.24
CFNAI(Dif)_AIC_rol	.	.	.	1.14	1.38	1.21
CFNAI(1sdbP)_AIC_rol	.	.	.	1.10	1.37	1.28
UR_5wk(Level)_AIC_rol	.	1.32	1.25	1.16	1.43	1.26
UR_5wk(Dif)_AIC_rol	.	1.12	1.06	1.17	1.34	1.28
UR_5wk(1sdbP)_AIC_rol	.	1.20	1.10	1.03	1.43	1.27
AHE(Dif)_AIC_rol	.	.	1.10	1.35	1.33	1.24
AHE(1sdbP)_AIC_rol	.	.	1.15	1.06	1.34	1.14
RealAHE(Dif)_AIC_rol	.	.	1.10	1.35	1.33	1.24
RealAHE(1sdbP)_AIC_rol	.	.	1.15	1.06	1.34	1.14
LaborShare(Level)_AIC_rol	.	1.26	1.01	1.21	1.45	1.27
LaborShare(Dif)_AIC_rol	.	1.16	1.04	1.11	1.32	1.48
ULaborShare(1sdbP)_AIC_rol	.	1.22	1.03	1.09	1.35	1.23
CPI_Med(Level)_AIC_rol	.	.	.	1.16	1.27	1.20
CPI_Med(Dif)_AIC_rol	.	.	.	1.04	1.25	1.32
CPI_TrMn(Level)_AIC_rol	.	.	.	1.23	1.31	1.13
CPI_TrMn(Dif)_AIC_rol	.	.	.	1.07	1.29	1.28
ExRate(Dif)_AIC_rol	.	.	.	1.33	1.25	1.24
ExRate(1sdbP)_AIC_rol	.	.	.	1.64	1.25	1.22

Table 3.3 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
tb_spr_AIC_rol	.	1.18	1.79	1.39	1.38	1.35
UR(Level)_BIC_rec	.	1.05	0.98	1.06	1.22	1.42
UR(Dif)_BIC_rec	.	1.00	1.08	1.11	1.15	1.48
UR(1sdBP)_BIC_rec	.	1.05	1.07	1.06	1.14	1.42
GDP(Dif)_BIC_rec	.	1.03	1.09	1.10	1.15	1.47
GDP(1sdBP)_BIC_rec	.	1.09	0.99	1.07	1.07	1.44
IP(Dif)_BIC_rec	.	1.04	1.03	1.10	1.23	1.50
IP(1sdBP)_BIC_rec	.	1.07	1.02	1.09	1.16	1.47
Emp(Dif)_BIC_rec	.	1.02	0.97	1.02	1.11	1.59
Emp(1sdBP)_BIC_rec	.	1.06	0.98	1.06	1.13	1.54
CapU(Level)_BIC_rec	.	.	.	1.16	1.28	1.75
CapU(Dif)_BIC_rec	.	.	.	1.22	1.31	1.79
CapU(1sdBP)_BIC_rec	.	.	.	1.10	1.20	1.47
HPerm(Level)_BIC_rec	.	.	0.95	1.12	1.07	1.74
HPerm(Dif)_BIC_rec	.	.	1.07	1.21	1.07	1.61
HPerm(1sdBP)_BIC_rec	.	.	0.97	1.09	1.07	1.43
CFNAI(Dif)_BIC_rec	.	.	.	1.16	1.17	1.81
CFNAI(1sdBP)_BIC_rec	.	.	.	1.09	1.18	1.54
UR_5wk(Level)_BIC_rec	.	1.16	0.94	1.15	1.43	1.57
UR_5wk(Dif)_BIC_rec	.	1.02	1.02	1.15	1.12	1.54
UR_5wk(1sdBP)_BIC_rec	.	1.08	0.97	1.13	1.20	1.43
AHE(Dif)_BIC_rec	.	.	1.07	1.29	1.11	1.62
AHE(1sdBP)_BIC_rec	.	.	1.15	1.31	1.08	1.57
RealAHE(Dif)_BIC_rec	.	.	1.07	1.29	1.11	1.62
RealAHE(1sdBP)_BIC_rec	.	.	1.15	1.31	1.08	1.57
LaborShare(Level)_BIC_rec	.	1.15	0.98	1.30	1.51	1.65
LaborShare(Dif)_BIC_rec	.	1.11	1.01	1.14	1.10	1.54
ULaborShare(1sdBP)_BIC_rec	.	1.18	0.99	1.14	1.32	1.64
CPI_Med(Level)_BIC_rec	.	.	.	1.37	1.26	1.63
CPI_Med(Dif)_BIC_rec	.	.	.	1.27	1.13	1.64
CPI_TrMn(Level)_BIC_rec	.	.	.	1.35	1.25	1.61
CPI_TrMn(Dif)_BIC_rec	.	.	.	1.27	1.10	1.64
ExRate(Dif)_BIC_rec	.	.	.	1.38	1.25	1.41
ExRate(1sdBP)_BIC_rec	.	.	.	1.53	1.16	1.46
tb_spr_BIC_rol	.	1.14	1.13	1.16	1.04	1.59
UR(Level)_BIC_rol	.	1.23	1.29	1.10	1.32	1.31
UR(Dif)_BIC_rol	.	1.06	1.13	1.08	1.33	1.30
UR(1sdBP)_BIC_rol	.	1.22	1.29	1.10	1.32	1.31
GDP(Dif)_BIC_rol	.	1.05	1.12	1.09	1.33	1.32
GDP(1sdBP)_BIC_rol	.	1.23	1.09	1.12	1.28	1.30
IP(Dif)_BIC_rol	.	1.05	1.01	1.18	1.43	1.27
IP(1sdBP)_BIC_rol	.	1.23	1.17	1.15	1.46	1.39
Emp(Dif)_BIC_rol	.	1.09	1.12	1.14	1.36	1.30
Emp(1sdBP)_BIC_rol	.	1.18	1.22	1.18	1.39	1.29
CapU(Level)_BIC_rol	.	.	.	1.15	1.48	1.34
CapU(Dif)_BIC_rol	.	.	.	1.16	1.40	1.32

Table 3.3 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
CapU(1sdBP)_BIC_rol	.	.	.	1.09	1.46	1.36
HPerm(Level)_BIC_rol	.	.	1.00	1.10	1.27	1.39
HPerm((Dif)_BIC_rol	.	.	1.09	1.07	1.10	1.33
HPerm(1sdBP)_BIC_rol	.	.	1.00	1.09	1.27	1.33
CFNAI(Dif)_BIC_rol	.	.	.	1.13	1.41	1.27
CFNAI(1sdBP)_BIC_rol	.	.	.	1.06	1.39	1.28
UR_5wk(Level)_BIC_rol	.	1.26	1.24	1.19	1.42	1.33
UR_5wk(Dif)_BIC_rol	.	1.04	1.07	1.12	1.32	1.30
UR_5wk(1sdBP)_BIC_rol	.	1.17	1.10	1.06	1.42	1.27
AHE(Dif)_BIC_rol	.	.	1.08	1.09	1.27	1.30
AHE(1sdBP)_BIC_rol	.	.	1.15	1.13	1.33	1.26
RealAHE(Dif)_BIC_rol	.	.	1.08	1.09	1.27	1.30
RealAHE(1sdBP)_BIC_rol	.	.	1.15	1.13	1.33	1.26
LaborShare(Level)_BIC_rol	.	1.19	0.97	1.20	1.38	1.31
LaborShare(Dif)_BIC_rol	.	1.13	1.04	1.24	1.34	1.44
ULaborShare(1sdBP)_BIC_rol	.	1.17	0.96	1.11	1.35	1.32
CPI_Med(Level)_BIC_rol	.	.	.	1.12	1.29	1.19
CPI_Med(Dif)_BIC_rol	.	.	.	1.09	1.27	1.29
CPI_TrMn(Level)_BIC_rol	.	.	.	1.15	1.27	1.16
CPI_TrMn(Dif)_BIC_rol	.	.	.	1.08	1.28	1.34
ExRate(Dif)_BIC_rol	.	.	.	1.28	1.25	1.40
ExRate(1sdBP)_BIC_rol	.	.	.	1.48	1.26	1.40
tb_spr_BIC_rol	.	1.15	1.76	1.15	1.33	1.37
<i>Triangle model forecasts</i>						
Triangle Constant NAIRU	.	.	1.14	1.18	1.25	1.20
Triangle TV NAIRU	.	.	1.07	1.20	1.04	1.30
Triangle Constant NAIRU (no z)	.	.	0.98	1.33	1.38	1.27
Triangle TV NAIRU (no z)	.	.	0.97	1.48	1.16	1.58
<i>Combination forecasts</i>						
Activity Median Combining	.	1.07	0.97	1.05	1.16	1.32
Activity Mean Combining	.	1.07	0.94	1.05	1.14	1.35
Activity Tr. Mean Combining	.	1.07	0.95	1.05	1.16	1.34
Activity MSE(A) Combining	.	.	0.95	1.04	1.15	1.35
Activity MSE(B) Combining	.	.	0.95	1.04	1.15	1.35
Activity MSE(C) Combining	.	.	0.95	1.04	1.14	1.35
Activity MSE(D) Combining	.	.	0.95	1.05	1.16	1.35
Activity MSE(E) Combining	.	.	0.95	1.04	1.16	1.35
Activity MSE(F) Combining	.	.	0.94	1.04	1.15	1.35
Activity Rec. Best(4q) Combining	.	1.20	1.10	1.25	1.44	1.42
Activity Rec. Best(8q) Combining	.	1.21	0.99	1.19	1.46	1.61
OtherADL Median Combining	.	1.14	1.04	1.15	1.17	1.33
OtherADL Mean Combining	.	1.14	1.02	1.12	1.14	1.35
OtherADL Tr. Mean Combining	.	1.14	1.03	1.13	1.14	1.35
OtherADL MSE(A) Combining	.	.	0.93	1.14	1.15	1.34
OtherADL MSE(B) Combining	.	.	0.93	1.14	1.15	1.34
OtherADL MSE(C) Combining	.	.	0.94	1.14	1.16	1.35

Table 3.3 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
OtherADL MSE(D) Combining .	.	.	0.94	1.14	1.16	1.32
OtherADL MSE(E) Combining .	.	.	0.93	1.13	1.17	1.34
OtherADL MSE(F) Combining .	.	.	0.94	1.13	1.17	1.36
OtherADL Rec. Best(4q) Combining .	1.29		1.13	1.18	1.32	1.11
OtherADL Rec. Best(8q) Combining .	1.32		0.98	1.32	1.33	1.26
All Median Combining .	1.07		0.94	1.07	1.15	1.31
All Mean Combining .	1.08		0.93	1.06	1.12	1.34
All Tr. Mean Combining .	1.08		0.93	1.07	1.14	1.33
All MSE(A) Combining .	.		0.92	1.06	1.14	1.34
All MSE(B) Combining .	.		0.92	1.06	1.14	1.34
All MSE(C) Combining .	.		0.92	1.06	1.13	1.34
All MSE(D) Combining .	.		0.93	1.07	1.16	1.33
All MSE(E) Combining .	.		0.92	1.06	1.15	1.34
All MSE(F) Combining .	.		0.92	1.06	1.14	1.34
All Rec. Best(4q) Combining .	1.25		1.18	1.34	1.44	1.32
All Rec. Best(8q) Combining .	1.22		1.01	1.30	1.52	1.44
UCSV and Triangle Rec. Best(4q) Combining .	.		.	1.07	1.18	1.07
UCSV and Triangle Rec. Best(8q) Combining .	.		.	1.15	1.16	1.07

Source: Authors' calculations.

Table 3.4

Root Mean Squared Errors for Inflation Forecasting Models by Subperiod, Relative to the Unobserved Components-Stochastic Volatility Model: PCE-core

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
Number of observations	32	36	32	32	32	25
Root MSE of UC-SV forecast	0.68	1.56	1.08	0.55	0.36	0.33
Forecasting model and relative root mean square errors						
<i>Univariate forecasts</i>						
UC-SV	1.00	1.00	1.00	1.00	1.00	1.00
AR(AIC)_rec	.	.	1.15	1.15	1.21	1.34
AR(AIC)_iter_rec	.	.	1.09	1.17	1.24	1.37
AR(BIC)_rec	.	.	1.13	1.22	1.21	1.34
AO	1.08	1.16	1.12	1.00	0.94	1.18
MA(1)_rec	.	1.03	1.03	1.09	1.16	1.27
AR(4)_rec	.	.	1.15	1.15	1.18	1.29
AR(AIC)_roll	.	.	1.15	1.24	1.18	1.30
AR(AIC)_iter_roll	.	.	1.15	1.31	1.21	1.27
AR(BIC)_roll	.	.	1.15	1.06	1.24	1.28
AR(4)_roll	.	.	1.19	1.25	1.16	1.27
AR(24)_iter	.	.	.	1.85	1.37	1.26
AR(24)_iter_nocon	.	.	1.14	1.32	1.26	1.24
MA(1)_roll	.	1.02	1.03	1.08	1.14	1.06
MA(2) - NS	1.09	1.05	1.14	0.99	1.07	1.06
MA(1), $\theta=.25$	1.01	1.01	0.99	1.11	1.19	1.33
MA(1), $\theta=.65$	1.08	1.07	1.12	0.98	1.03	1.07
<i>Single-predictor ADL forecasts</i>						
UR(Level)_AIC_rec	.	.	1.08	1.01	1.48	1.51
UR(Dif)_AIC_rec	.	.	1.14	1.16	1.42	1.30
UR(1sdbp)_AIC_rec	.	.	1.19	1.08	1.30	1.41
GDP(Dif)_AIC_rec	.	.	1.03	1.17	1.38	1.45
GDP(1sdbp)_AIC_rec	.	.	0.95	1.01	1.06	1.50
IP(Dif)_AIC_rec	.	.	1.01	1.15	1.39	1.29
IP(1sdbp)_AIC_rec	.	.	0.96	1.01	1.25	1.80
Emp(Dif)_AIC_rec	.	.	1.00	1.14	1.26	2.10
Emp(1sdbp)_AIC_rec	.	.	1.14	1.06	1.28	1.70
CapU(Level)_AIC_rec	.	.	.	1.23	1.53	2.81
CapU((Dif)_AIC_rec	.	.	.	1.23	1.39	1.27
CapU(1sdbp)_AIC_rec	.	.	.	1.19	1.17	1.52
HPerm(Level)_AIC_rec	.	.	1.13	1.17	1.30	2.11
HPerm((Dif)_AIC_rec	.	.	1.25	1.17	1.22	1.34
HPerm(1sdbp)_AIC_rec	.	.	1.15	1.04	1.38	1.32
CFNAI(Dif)_AIC_rec	.	.	.	1.15	1.25	1.68
CFNAI(1sdbp)_AIC_rec	.	.	.	1.00	1.12	1.63
UR_5wk(Level)_AIC_rec	.	.	0.96	1.18	1.99	1.73
UR_5wk(Dif)_AIC_rec	.	.	1.14	1.21	1.44	1.21
UR_5wk(1sdbp)_AIC_rec	.	.	1.06	1.22	1.57	1.24
AHE(Dif)_AIC_rec	.	.	1.18	1.24	1.22	1.34

Table 3.4 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
AHE(1sdbP)_AIC_rec	.	.	1.39	1.29	1.27	1.34
RealAHE(Dif)_AIC_rec	.	.	1.18	1.24	1.22	1.34
RealAHE(1sdbP)_AIC_rec	.	.	1.39	1.29	1.27	1.34
LaborShare(Level)_AIC_rec	.	.	0.99	1.39	1.91	1.57
LaborShare(Dif)_AIC_rec	.	.	1.09	1.18	1.40	1.50
ULaborShare(1sdbP)_AIC_rec	.	.	1.05	1.25	1.55	1.59
CPI_Med(Level)_AIC_rec	.	.	.	1.20	1.33	1.37
CPI_Med(Dif)_AIC_rec	.	.	.	1.12	1.17	1.49
CPI_TrMn(Level)_AIC_rec	.	.	.	1.20	1.19	1.32
CPI_TrMn(Dif)_AIC_rec	.	.	.	1.09	1.05	1.28
ExRate(Dif)_AIC_rec	.	.	.	1.29	1.35	1.27
ExRate(1sdbP)_AIC_rec	.	.	.	1.45	1.26	1.28
tb_spr_AIC_rec	.	.	1.40	1.33	1.22	1.49
UR(Level)_AIC_rol	.	.	1.44	1.20	1.03	1.49
UR(Dif)_AIC_rol	.	.	1.22	1.20	1.20	1.31
UR(1sdbP)_AIC_rol	.	.	1.20	0.93	1.01	1.60
GDP(Dif)_AIC_rol	.	.	1.17	1.24	1.23	1.33
GDP(1sdbP)_AIC_rol	.	.	1.00	0.98	1.01	1.55
IP(Dif)_AIC_rol	.	.	1.10	1.35	1.29	1.44
IP(1sdbP)_AIC_rol	.	.	0.98	1.11	1.20	1.73
Emp(Dif)_AIC_rol	.	.	1.08	1.14	1.18	1.58
Emp(1sdbP)_AIC_rol	.	.	1.23	1.17	1.12	2.07
CapU(Level)_AIC_rol	.	.	.	1.21	1.33	1.73
CapU(Dif)_AIC_rol	.	.	.	1.34	1.27	1.44
CapU(1sdbP)_AIC_rol	.	.	.	1.00	1.26	1.60
HPerm(Level)_AIC_rol	.	.	1.27	1.17	1.03	1.51
HPerm(Dif)_AIC_rol	.	.	1.37	1.17	1.26	1.23
HPerm(1sdbP)_AIC_rol	.	.	1.44	1.11	1.15	1.32
CFNAI(Dif)_AIC_rol	.	.	.	1.14	1.20	1.49
CFNAI(1sdbP)_AIC_rol	.	.	.	0.96	1.09	1.73
UR_5wk(Level)_AIC_rol	.	.	1.14	1.44	1.04	1.44
UR_5wk(Dif)_AIC_rol	.	.	1.17	1.06	1.16	1.28
UR_5wk(1sdbP)_AIC_rol	.	.	1.13	1.00	1.26	1.33
AHE(Dif)_AIC_rol	.	.	1.18	1.25	1.21	1.32
AHE(1sdbP)_AIC_rol	.	.	1.36	1.36	1.21	1.61
RealAHE(Dif)_AIC_rol	.	.	1.18	1.25	1.21	1.32
RealAHE(1sdbP)_AIC_rol	.	.	1.36	1.36	1.21	1.61
LaborShare(Level)_AIC_rol	.	.	1.20	1.35	1.56	1.58
LaborShare(Dif)_AIC_rol	.	.	1.32	1.12	1.16	1.63
ULaborShare(1sdbP)_AIC_rol	.	.	1.27	1.23	1.46	1.55
CPI_Med(Level)_AIC_rol	.	.	.	1.37	1.31	1.14
CPI_Med(Dif)_AIC_rol	.	.	.	1.20	1.21	1.31
CPI_TrMn(Level)_AIC_rol	.	.	.	1.34	1.28	1.33
CPI_TrMn(Dif)_AIC_rol	.	.	.	1.15	1.19	1.33
ExRate(Dif)_AIC_rol	.	.	.	1.39	1.21	1.44
ExRate(1sdbP)_AIC_rol	.	.	.	1.43	1.12	1.46

Table 3.4 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
tb_spr_AIC_rol	.	.	1.64	1.26	1.24	1.30
UR(Level)_BIC_rec	.	.	0.96	1.06	1.48	1.58
UR(Dif)_BIC_rec	.	.	1.17	1.16	1.42	1.34
UR(1sdbP)_BIC_rec	.	.	1.10	1.12	1.35	1.39
GDP(Dif)_BIC_rec	.	.	1.10	1.17	1.38	1.47
GDP(1sdbP)_BIC_rec	.	.	0.95	1.07	1.14	1.46
IP(Dif)_BIC_rec	.	.	1.14	1.15	1.40	1.33
IP(1sdbP)_BIC_rec	.	.	0.97	1.03	1.23	1.67
Emp(Dif)_BIC_rec	.	.	1.01	1.15	1.30	2.09
Emp(1sdbP)_BIC_rec	.	.	1.06	1.10	1.27	1.66
CapU(Level)_BIC_rec	.	.	.	1.25	1.45	2.53
CapU(Dif)_BIC_rec	.	.	.	1.23	1.39	1.27
CapU(1sdbP)_BIC_rec	.	.	.	1.20	1.18	1.48
HPerm(Level)_BIC_rec	.	.	1.10	1.13	1.20	1.74
HPerm(Dif)_BIC_rec	.	.	1.15	1.23	1.22	1.34
HPerm(1sdbP)_BIC_rec	.	.	1.11	1.11	1.38	1.38
CFNAI(Dif)_BIC_rec	.	.	.	1.16	1.30	1.68
CFNAI(1sdbP)_BIC_rec	.	.	.	1.09	1.16	1.63
UR_5wk(Level)_BIC_rec	.	.	1.03	1.18	1.98	1.76
UR_5wk(Dif)_BIC_rec	.	.	1.14	1.20	1.21	1.25
UR_5wk(1sdbP)_BIC_rec	.	.	1.04	1.23	1.58	1.31
AHE(Dif)_BIC_rec	.	.	1.19	1.26	1.22	1.34
AHE(1sdbP)_BIC_rec	.	.	1.31	1.37	1.31	1.28
RealAHE(Dif)_BIC_rec	.	.	1.19	1.26	1.22	1.34
RealAHE(1sdbP)_BIC_rec	.	.	1.31	1.37	1.31	1.28
LaborShare(Level)_BIC_rec	.	.	1.07	1.32	1.73	1.40
LaborShare(Dif)_BIC_rec	.	.	1.06	1.14	1.46	1.36
ULaborShare(1sdbP)_BIC_rec	.	.	1.08	1.22	1.66	1.46
CPI_Med(Level)_BIC_rec	.	.	.	1.29	1.36	1.46
CPI_Med(Dif)_BIC_rec	.	.	.	1.19	1.24	1.48
CPI_TrMn(Level)_BIC_rec	.	.	.	1.21	1.28	1.38
CPI_TrMn(Dif)_BIC_rec	.	.	.	1.14	1.19	1.33
ExRate(Dif)_BIC_rec	.	.	.	1.15	1.38	1.31
ExRate(1sdbP)_BIC_rec	.	.	.	1.19	1.26	1.32
tb_spr_BIC_rec	.	.	1.26	1.22	1.12	1.50
UR(Level)_BIC_rol	.	.	1.40	1.18	1.01	1.50
UR(Dif)_BIC_rol	.	.	1.25	1.02	1.21	1.30
UR(1sdbP)_BIC_rol	.	.	1.27	0.97	1.00	1.57
GDP(Dif)_BIC_rol	.	.	1.17	1.10	1.23	1.30
GDP(1sdbP)_BIC_rol	.	.	1.02	1.02	1.00	1.53
IP(Dif)_BIC_rol	.	.	1.24	1.11	1.31	1.34
IP(1sdbP)_BIC_rol	.	.	1.05	1.12	1.18	1.65
Emp(Dif)_BIC_rol	.	.	1.14	1.11	1.24	1.47
Emp(1sdbP)_BIC_rol	.	.	1.17	1.06	1.12	2.04
CapU(Level)_BIC_rol	.	.	.	1.21	1.33	1.44
CapU(Dif)_BIC_rol	.	.	.	1.07	1.28	1.39

Table 3.4 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
CapU(1sdBP)_BIC_rol	.	.	.	1.07	1.24	1.37
HPerm(Level)_BIC_rol	.	.	1.15	1.15	1.00	1.54
HPerm((Dif)_BIC_rol	.	.	1.19	1.08	1.26	1.21
HPerm(1sdBP)_BIC_rol	.	.	1.20	1.12	1.11	1.31
CFNAI(Dif)_BIC_rol	.	.	.	1.08	1.24	1.37
CFNAI(1sdBP)_BIC_rol	.	.	.	0.99	1.09	1.74
UR_5wk(Level)_BIC_rol	.	.	1.14	1.42	1.05	1.45
UR_5wk(Dif)_BIC_rol	.	.	1.20	1.05	1.18	1.28
UR_5wk(1sdBP)_BIC_rol	.	.	1.16	1.01	1.27	1.27
AHE(Dif)_BIC_rol	.	.	1.16	1.11	1.24	1.34
AHE(1sdBP)_BIC_rol	.	.	1.31	1.17	1.22	1.49
RealAHE(Dif)_BIC_rol	.	.	1.16	1.11	1.24	1.34
RealAHE(1sdBP)_BIC_rol	.	.	1.31	1.17	1.22	1.49
LaborShare(Level)_BIC_rol	.	.	1.14	1.22	1.41	1.45
LaborShare(Dif)_BIC_rol	.	.	1.12	1.09	1.26	1.58
ULaborShare(1sdBP)_BIC_rol	.	.	1.24	1.15	1.45	1.46
CPI_Med(Level)_BIC_rol	.	.	.	1.24	1.29	1.14
CPI_Med(Dif)_BIC_rol	.	.	.	1.05	1.22	1.27
CPI_TrMn(Level)_BIC_rol	.	.	.	1.23	1.29	1.30
CPI_TrMn(Dif)_BIC_rol	.	.	.	1.07	1.25	1.31
ExRate(Dif)_BIC_rol	.	.	.	1.14	1.29	1.43
ExRate(1sdBP)_BIC_rol	.	.	.	1.07	1.08	1.41
tb_spr_BIC_rol	.	.	1.62	1.10	1.25	1.27
<i>Triangle model forecasts</i>						
Triangle Constant NAIRU	.	.	1.69	1.06	1.80	1.55
Triangle TV NAIRU	.	.	1.94	0.99	1.37	1.44
Triangle Constant NAIRU (no z)	.	.	1.17	1.64	2.20	1.58
Triangle TV NAIRU (no z)	.	.	1.28	1.58	1.48	2.13
<i>Combination forecasts</i>						
Activity Median Combining	.	.	0.95	1.01	1.07	1.27
Activity Mean Combining	.	.	0.93	1.00	1.07	1.28
Activity Tr. Mean Combining	.	.	0.95	1.01	1.08	1.27
Activity MSE(A) Combining	.	.	.	0.99	1.08	1.31
Activity MSE(B) Combining	.	.	.	0.98	1.08	1.30
Activity MSE(C) Combining	.	.	.	0.98	1.07	1.30
Activity MSE(D) Combining	.	.	.	0.97	1.10	1.34
Activity MSE(E) Combining	.	.	.	0.95	1.08	1.33
Activity MSE(F) Combining	.	.	.	0.95	1.07	1.33
Activity Rec. Best(4q) Combining	.	.	1.01	1.39	1.37	1.95
Activity Rec. Best(8q) Combining	.	.	0.78	1.35	1.29	1.57
OtherADL Median Combining	.	.	1.10	1.11	1.16	1.24
OtherADL Mean Combining	.	.	1.08	1.12	1.14	1.23
OtherADL Tr. Mean Combining	.	.	1.10	1.11	1.14	1.22
OtherADL MSE(A) Combining	.	.	.	1.13	1.16	1.26
OtherADL MSE(B) Combining	.	.	.	1.13	1.16	1.25
OtherADL MSE(C) Combining	.	.	.	1.13	1.16	1.24

Table 3.4 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
OtherADL MSE(D) Combining	1.13	1.18	1.29
OtherADL MSE(E) Combining	1.13	1.17	1.26
OtherADL MSE(F) Combining	1.13	1.17	1.25
OtherADL Rec. Best(4q) Combining .	.	.	1.30	1.37	1.51	1.42
OtherADL Rec. Best(8q) Combining .	.	.	1.40	1.32	1.49	1.72
All Median Combining .	.	.	1.01	1.03	1.09	1.26
All Mean Combining .	.	.	0.96	1.03	1.08	1.24
All Tr. Mean Combining .	.	.	0.98	1.04	1.09	1.24
All MSE(A) Combining	1.01	1.10	1.27
All MSE(B) Combining	1.01	1.09	1.26
All MSE(C) Combining	1.00	1.08	1.26
All MSE(D) Combining	1.00	1.12	1.30
All MSE(E) Combining	0.98	1.10	1.29
All MSE(F) Combining	0.97	1.08	1.28
All Rec. Best(4q) Combining .	.	.	1.20	1.44	1.43	1.93
All Rec. Best(8q) Combining .	.	.	0.78	1.35	1.43	1.61
UCSV and Triangle Rec. Best(4q) Combining	1.04	1.13	1.03
UCSV and Triangle Rec. Best(8q) Combining	1.04	1.26	1.13

Source: Authors' calculations.

Table 3.5

Root Mean Squared Errors for Inflation Forecasting Models by Subperiod, Relative to the Unobserved Components-Stochastic Volatility Model: GDP deflator

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
Number of observations	32	36	32	32	32	25
Root MSE of UC-SV forecast	0.72	1.76	1.28	0.70	0.41	0.57
Forecasting model and relative root mean square errors						
<i>Univariate forecasts</i>						
UC-SV	1.00	1.00	1.00	1.00	1.00	1.00
AR(AIC)_rec	.	1.03	1.06	1.06	1.02	1.16
AR(AIC)_iter_rec	.	1.11	1.08	1.04	0.99	1.19
AR(BIC)_rec	.	1.03	1.04	1.08	1.07	1.24
AO	0.97	1.10	1.17	1.04	0.95	1.02
MA(1)_rec	.	1.02	1.00	1.04	1.02	1.16
AR(4)_rec	.	1.07	1.07	1.04	0.99	1.17
AR(AIC)_roll	.	1.11	1.05	1.15	1.19	1.16
AR(AIC)_iter_roll	.	1.10	1.06	1.02	1.12	1.11
AR(BIC)_roll	.	1.08	1.05	1.21	1.17	1.11
AR(4)_roll	.	1.15	1.08	1.08	1.14	1.15
AR(24)_iter	.	.	1.42	1.10	1.02	0.99
AR(24)_iter_nocon	.	.	1.34	1.02	0.99	0.99
MA(1)_roll	.	1.03	0.99	1.05	0.98	1.02
MA(2) - NS	0.97	1.02	1.19	1.03	1.02	1.02
MA(1), $\theta=.25$	1.03	1.00	1.00	1.07	1.08	1.25
MA(1), $\theta=.65$	0.96	1.03	1.17	1.02	0.99	1.01
<i>Single-predictor ADL forecasts</i>						
UR(Level)_AIC_rec	.	0.93	0.99	0.91	1.23	1.30
UR(Dif)_AIC_rec	.	0.93	1.11	0.96	1.25	1.22
UR(1sdBP)_AIC_rec	.	0.94	1.12	0.91	1.14	1.19
GDP(Dif)_AIC_rec	.	0.94	1.04	0.91	1.06	1.09
GDP(1sdBP)_AIC_rec	.	0.98	1.01	0.89	0.96	1.15
IP(Dif)_AIC_rec	.	0.90	1.05	0.89	1.25	1.15
IP(1sdBP)_AIC_rec	.	0.93	1.04	0.86	1.18	1.23
Emp(Dif)_AIC_rec	.	0.93	1.03	0.93	1.11	1.42
Emp(1sdBP)_AIC_rec	.	0.94	1.05	0.94	1.19	1.35
CapU(Level)_AIC_rec	.	.	.	1.03	1.54	1.87
CapU((Dif)_AIC_rec	.	.	.	0.96	1.39	1.22
CapU(1sdBP)_AIC_rec	.	.	.	0.91	1.23	1.23
HPerm(Level)_AIC_rec	.	.	1.07	1.05	0.89	1.60
HPerm((Dif)_AIC_rec	.	.	1.17	1.09	1.01	1.13
HPerm(1sdBP)_AIC_rec	.	.	1.18	1.04	1.10	1.09
CFNAI(Dif)_AIC_rec	.	.	.	1.00	1.16	1.54
CFNAI(1sdBP)_AIC_rec	.	.	.	0.89	1.12	1.46
UR_5wk(Level)_AIC_rec	.	1.08	1.01	0.92	1.83	1.32
UR_5wk(Dif)_AIC_rec	.	0.99	1.07	0.93	1.32	1.11
UR_5wk(1sdBP)_AIC_rec	.	0.99	1.04	0.95	1.30	1.14
AHE(Dif)_AIC_rec	.	.	1.09	1.09	1.05	1.16

Table 3.5 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
AHE(1sdBP)_AIC_rec	.	.	1.21	1.19	1.07	1.13
RealAHE(Dif)_AIC_rec	.	.	1.09	1.09	1.05	1.16
RealAHE(1sdBP)_AIC_rec	.	.	1.21	1.19	1.07	1.13
LaborShare(Level)_AIC_rec	.	1.10	0.99	1.18	1.85	1.27
LaborShare(Dif)_AIC_rec	.	1.09	1.06	1.07	1.12	1.14
ULaborShare(1sdBP)_AIC_rec	.	1.10	1.01	1.08	1.47	1.13
CPI_Med(Level)_AIC_rec	.	.	.	1.09	1.38	1.29
CPI_Med(Dif)_AIC_rec	.	.	.	1.02	1.15	1.34
CPI_TrMn(Level)_AIC_rec	.	.	.	1.12	1.17	1.21
CPI_TrMn(Dif)_AIC_rec	.	.	.	1.02	0.96	1.19
ExRate(Dif)_AIC_rec	.	.	.	1.23	1.44	1.06
ExRate(1sdBP)_AIC_rec	.	.	.	1.66	1.25	0.91
tb_spr_AIC_rec	.	1.08	1.23	1.19	0.95	1.27
UR(Level)_AIC_rol	.	1.09	1.43	1.06	1.16	1.11
UR(Dif)_AIC_rol	.	0.97	1.35	1.12	1.25	1.06
UR(1sdBP)_AIC_rol	.	1.03	1.33	1.04	1.25	1.13
GDP(Dif)_AIC_rol	.	1.13	1.21	1.08	1.26	1.05
GDP(1sdBP)_AIC_rol	.	1.15	0.96	1.05	1.21	1.07
IP(Dif)_AIC_rol	.	1.01	1.22	1.10	1.42	1.10
IP(1sdBP)_AIC_rol	.	1.05	1.10	1.12	1.48	1.20
Emp(Dif)_AIC_rol	.	1.04	1.24	1.15	1.23	1.08
Emp(1sdBP)_AIC_rol	.	1.06	1.23	1.17	1.27	1.23
CapU(Level)_AIC_rol	.	.	.	1.12	1.39	1.12
CapU(Dif)_AIC_rol	.	.	.	1.11	1.34	1.10
CapU(1sdBP)_AIC_rol	.	.	.	1.05	1.33	1.06
HPerm(Level)_AIC_rol	.	.	1.18	1.36	0.92	1.35
HPerm(Dif)_AIC_rol	.	.	1.20	1.21	1.12	1.08
HPerm(1sdBP)_AIC_rol	.	.	1.36	1.33	1.07	1.10
CFNAI(Dif)_AIC_rol	.	.	.	1.12	1.29	1.11
CFNAI(1sdBP)_AIC_rol	.	.	.	1.10	1.31	1.16
UR_5wk(Level)_AIC_rol	.	1.28	1.18	1.21	1.86	1.20
UR_5wk(Dif)_AIC_rol	.	1.09	1.11	1.17	1.45	1.14
UR_5wk(1sdBP)_AIC_rol	.	1.08	1.09	0.94	1.38	1.15
AHE(Dif)_AIC_rol	.	.	1.07	1.06	1.18	1.14
AHE(1sdBP)_AIC_rol	.	.	1.19	1.11	1.22	1.06
RealAHE(Dif)_AIC_rol	.	.	1.07	1.06	1.18	1.14
RealAHE(1sdBP)_AIC_rol	.	.	1.19	1.11	1.22	1.06
LaborShare(Level)_AIC_rol	.	1.22	1.39	1.18	1.25	1.14
LaborShare(Dif)_AIC_rol	.	1.11	1.53	1.11	1.60	1.26
ULaborShare(1sdBP)_AIC_rol	.	1.15	1.44	1.13	1.30	1.15
CPI_Med(Level)_AIC_rol	.	.	.	1.14	1.31	0.87
CPI_Med(Dif)_AIC_rol	.	.	.	1.06	1.20	1.14
CPI_TrMn(Level)_AIC_rol	.	.	.	1.16	1.21	0.72
CPI_TrMn(Dif)_AIC_rol	.	.	.	1.07	1.16	1.13
ExRate(Dif)_AIC_rol	.	.	.	1.42	1.15	1.19
ExRate(1sdBP)_AIC_rol	.	.	.	1.67	1.21	1.13

Table 3.5 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
tb_spr_AIC_rol	.	1.09	1.60	1.29	1.16	1.21
UR(Level)_BIC_rec	.	0.94	0.96	0.92	1.21	1.28
UR(Dif)_BIC_rec	.	0.94	1.08	0.97	1.20	1.20
UR(1sdbP)_BIC_rec	.	0.93	1.06	0.94	1.14	1.19
GDP(Dif)_BIC_rec	.	1.00	1.09	0.94	1.08	1.11
GDP(1sdbP)_BIC_rec	.	0.99	1.00	0.93	0.87	1.17
IP(Dif)_BIC_rec	.	0.96	1.06	0.89	1.22	1.11
IP(1sdbP)_BIC_rec	.	0.96	0.99	0.89	1.12	1.22
Emp(Dif)_BIC_rec	.	0.90	0.99	0.90	1.12	1.30
Emp(1sdbP)_BIC_rec	.	0.97	0.99	0.90	1.14	1.31
CapU(Level)_BIC_rec	.	.	.	1.06	1.47	1.83
CapU(Dif)_BIC_rec	.	.	.	1.02	1.39	1.22
CapU(1sdbP)_BIC_rec	.	.	.	0.94	1.19	1.21
HPerm(Level)_BIC_rec	.	.	0.99	1.05	0.89	1.60
HPerm(Dif)_BIC_rec	.	.	1.09	1.08	1.05	1.22
HPerm(1sdbP)_BIC_rec	.	.	1.08	0.99	1.09	1.05
CFNAI(Dif)_BIC_rec	.	.	.	1.00	1.16	1.51
CFNAI(1sdbP)_BIC_rec	.	.	.	0.86	1.08	1.32
UR_5wk(Level)_BIC_rec	.	1.08	1.00	1.01	1.73	1.43
UR_5wk(Dif)_BIC_rec	.	0.98	1.08	1.03	1.17	1.20
UR_5wk(1sdbP)_BIC_rec	.	1.00	1.03	0.98	1.28	1.18
AHE(Dif)_BIC_rec	.	.	1.07	1.17	1.07	1.24
AHE(1sdbP)_BIC_rec	.	.	1.21	1.31	1.11	1.20
RealAHE(Dif)_BIC_rec	.	.	1.07	1.17	1.07	1.24
RealAHE(1sdbP)_BIC_rec	.	.	1.21	1.31	1.11	1.20
LaborShare(Level)_BIC_rec	.	1.09	1.03	1.22	1.68	1.34
LaborShare(Dif)_BIC_rec	.	1.06	1.04	1.06	1.15	1.22
ULaborShare(1sdbP)_BIC_rec	.	1.11	1.03	1.09	1.49	1.23
CPI_Med(Level)_BIC_rec	.	.	.	1.14	1.38	1.29
CPI_Med(Dif)_BIC_rec	.	.	.	1.08	1.18	1.27
CPI_TrMn(Level)_BIC_rec	.	.	.	1.16	1.17	1.21
CPI_TrMn(Dif)_BIC_rec	.	.	.	1.08	0.99	1.22
ExRate(Dif)_BIC_rec	.	.	.	1.12	1.48	1.09
ExRate(1sdbP)_BIC_rec	.	.	.	1.28	1.33	1.11
tb_spr_BIC_rec	.	1.03	1.09	1.11	0.99	1.33
UR(Level)_BIC_rol	.	1.14	1.49	1.04	1.21	1.15
UR(Dif)_BIC_rol	.	1.03	1.33	1.14	1.21	1.10
UR(1sdbP)_BIC_rol	.	0.98	1.36	1.02	1.20	1.17
GDP(Dif)_BIC_rol	.	1.15	1.15	1.17	1.28	1.00
GDP(1sdbP)_BIC_rol	.	1.15	0.98	1.06	1.20	1.10
IP(Dif)_BIC_rol	.	1.05	1.16	1.10	1.34	1.08
IP(1sdbP)_BIC_rol	.	1.09	1.13	1.18	1.43	1.24
Emp(Dif)_BIC_rol	.	1.03	1.21	1.15	1.25	1.11
Emp(1sdbP)_BIC_rol	.	1.04	1.24	1.24	1.27	1.24
CapU(Level)_BIC_rol	.	.	.	1.01	1.37	1.14
CapU(Dif)_BIC_rol	.	.	.	1.10	1.27	1.05

Table 3.5 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
CapU(1sdBP)_BIC_rol	.	.	.	0.96	1.29	1.12
HPerm(Level)_BIC_rol	.	.	1.08	1.25	0.94	1.37
HPerm((Dif)_BIC_rol	.	.	1.08	1.25	1.14	1.07
HPerm(1sdBP)_BIC_rol	.	.	1.22	1.20	1.08	1.08
CFNAI(Dif)_BIC_rol	.	.	.	1.12	1.28	1.10
CFNAI(1sdBP)_BIC_rol	.	.	.	1.14	1.25	1.13
UR_5wk(Level)_BIC_rol	.	1.28	1.07	1.17	1.55	1.22
UR_5wk(Dif)_BIC_rol	.	1.07	1.08	1.14	1.27	1.08
UR_5wk(1sdBP)_BIC_rol	.	1.06	1.09	0.93	1.32	1.11
AHE(Dif)_BIC_rol	.	.	1.07	1.19	1.17	1.10
AHE(1sdBP)_BIC_rol	.	.	1.19	1.32	1.17	1.09
RealAHE(Dif)_BIC_rol	.	.	1.07	1.19	1.17	1.10
RealAHE(1sdBP)_BIC_rol	.	.	1.19	1.32	1.17	1.09
LaborShare(Level)_BIC_rol	.	1.21	1.24	1.22	1.24	1.06
LaborShare(Dif)_BIC_rol	.	1.06	1.32	1.23	1.31	1.16
ULaborShare(1sdBP)_BIC_rol	.	1.17	1.32	1.23	1.29	1.09
CPI_Med(Level)_BIC_rol	.	.	.	1.36	1.34	0.92
CPI_Med(Dif)_BIC_rol	.	.	.	1.13	1.18	1.26
CPI_TrMn(Level)_BIC_rol	.	.	.	1.26	1.14	0.73
CPI_TrMn(Dif)_BIC_rol	.	.	.	1.16	1.16	1.11
ExRate(Dif)_BIC_rol	.	.	.	1.26	1.22	1.16
ExRate(1sdBP)_BIC_rol	.	.	.	1.49	1.20	1.15
tb_spr_BIC_rol	.	1.13	1.50	1.21	1.14	1.17
<i>Triangle model forecasts</i>						
Triangle Constant NAIRU	.	.	1.08	0.78	1.22	1.20
Triangle TV NAIRU	.	.	0.98	0.81	1.07	1.23
Triangle Constant NAIRU (no z)	.	.	1.22	0.95	1.64	1.22
Triangle TV NAIRU (no z)	.	.	1.17	1.21	1.33	1.61
<i>Combination forecasts</i>						
Activity Median Combining	.	0.98	0.99	0.93	1.10	1.09
Activity Mean Combining	.	1.00	0.97	0.91	1.07	1.10
Activity Tr. Mean Combining	.	1.00	0.98	0.93	1.09	1.10
Activity MSE(A) Combining	.	.	0.98	0.91	1.09	1.11
Activity MSE(B) Combining	.	.	0.98	0.91	1.09	1.11
Activity MSE(C) Combining	.	.	0.98	0.91	1.08	1.10
Activity MSE(D) Combining	.	.	0.98	0.89	1.10	1.11
Activity MSE(E) Combining	.	.	0.98	0.89	1.09	1.11
Activity MSE(F) Combining	.	.	0.98	0.89	1.08	1.11
Activity Rec. Best(4q) Combining	.	1.09	1.24	1.02	1.33	1.07
Activity Rec. Best(8q) Combining	.	1.11	1.29	1.02	1.30	1.39
OtherADL Median Combining	.	1.09	1.02	1.06	1.12	1.08
OtherADL Mean Combining	.	1.10	0.99	1.02	1.09	1.07
OtherADL Tr. Mean Combining	.	1.09	0.98	1.05	1.09	1.07
OtherADL MSE(A) Combining	.	.	1.09	1.08	1.10	1.06
OtherADL MSE(B) Combining	.	.	1.07	1.08	1.11	1.06
OtherADL MSE(C) Combining	.	.	1.06	1.08	1.11	1.06

Table 3.5 (continued)

Forecast period	1960:Q1– 1967:Q4	1968:Q1– 1976:Q4	1977:Q1– 1984:Q4	1985:Q1– 1992:Q4	1993:Q1– 2000:Q4	2001:Q1– 2007:Q4
OtherADL MSE(D) Combining .	.	.	1.11	1.08	1.10	1.04
OtherADL MSE(E) Combining .	.	.	1.09	1.08	1.11	1.04
OtherADL MSE(F) Combining .	.	.	1.07	1.08	1.12	1.06
OtherADL Rec. Best(4q) Combining .	1.07	1.18	1.18	1.42	1.12	0.83
OtherADL Rec. Best(8q) Combining .	1.15	1.32	1.32	1.22	1.05	0.91
All Median Combining .	1.01	0.97	0.97	0.98	1.10	1.07
All Mean Combining .	1.02	0.94	0.94	0.94	1.05	1.08
All Tr. Mean Combining .	1.02	0.94	0.94	0.97	1.06	1.07
All MSE(A) Combining .	.	0.97	0.97	0.95	1.07	1.07
All MSE(B) Combining .	.	0.96	0.96	0.95	1.07	1.07
All MSE(C) Combining .	.	0.96	0.96	0.95	1.07	1.07
All MSE(D) Combining .	.	0.98	0.98	0.94	1.09	1.07
All MSE(E) Combining .	.	0.98	0.98	0.93	1.09	1.07
All MSE(F) Combining .	.	0.97	0.97	0.93	1.07	1.07
All Rec. Best(4q) Combining .	1.07	1.35	1.35	1.20	1.32	0.90
All Rec. Best(8q) Combining .	1.12	1.37	1.37	1.03	1.16	1.03
UCSV and Triangle Rec. Best(4q) Combining	0.91	1.14	1.17
UCSV and Triangle Rec. Best(8q) Combining	0.89	1.13	1.21

Source: Authors' calculations.

Notes

1. Experience has shown that the good in-sample fit of a forecasting model does not necessarily imply a good out-of-sample performance. The method of pseudo out-of-sample forecast evaluation aims to address this disjunction by simulating the experience a forecaster would have had using a forecasting model. In a pseudo out-of-sample forecasting exercise, one simulates standing at a given date t and performing all model specification and parameter estimation using only the data available at that date, then computing the b -period ahead forecast for date $t + b$; this is repeated for all dates in the forecast period.
2. A strict interpretation of pseudo out-of-sample forecasting would entail the use of real-time data (data of different vintages), but we interpret the term more generously to include the use of final data.
3. The specification in Gordon (1990), which is used here, differs from Gordon (1982, table 5, column 2) in three ways: (a) Gordon (1982) uses a polynomial distributed lag specification on lagged inflation, while Gordon (1990) uses a step function; (b) Gordon (1982) includes additional intercept shifts in 1970:Q3–1975:Q4 and 1976:Q1–1980:Q4, which are dropped in Gordon (1990); (c) Gordon (1982) uses Perry-weighted unemployment, whereas here we use overall unemployment.
4. Stockton and Glassman (1987), table 6, ratio of PHL(16,FE) to ARIMA RMSE for average of four intervals.
5. The random walk benchmark is a standard tool for forecast assessment, but it seems to have played at most a minor role in the inflation forecasting literature before Atkeson and Ohanian. The four-quarter random walk benchmark is nested in the AR(AIC) model, but evidently imposing the four-quarter random walk restriction matters considerably.
6. The UC-SV model imposes a unit root in inflation, so it is consistent with the Pivetta-Reis (2007) evidence that the largest AR root in inflation has been essentially one throughout the postwar sample. But the time-varying relative variances of the permanent and transitory innovation allow for persistence to change over the course of the sample and for spectral measures of persistence to decline over the sample, consistent with Cogley and Sargent (2002, 2005).
7. Koenig (2003, table 3) presented in-sample evidence that real-time markups (nonfinancial corporate GDP divided by nonfinancial corporate employee compensation), in conjunction with the unemployment rate, significantly contribute to a forecast combination regression for four-quarter CPI inflation over 1983–2001; however he did not present pseudo out-of-sample RMSEs. Two of Ang, Bekaert, and Wei's (2007) models (their PC9 and PC10) include the output gap and the labor income share, specifications similar to the Koenig's (2003), and the pseudo out-of-sample performance of these models is poor: over Ang, Bekaert, and Wei's (2007) two subsamples and four inflation measures, the RMSEs, relative to the ARMA(1,1) benchmark, range from 1.17 to 3.26. These results sug-

gest that markups are not a solution to the poor performance of Phillips curve forecasts over the post-85 samples.

8. Romer and Romer (2000) compared the performance of real-time professional inflation forecasts and found that Fed Greenbook forecasts outperform commercial forecasts (Data Resources, Inc., Blue Chip, Survey of Professional Forecasters) over the period starting 1968:M11–1980:M1 (the start date depends on the forecast source) through 1999:M11. Romer and Romer’s (2000) findings do not speak directly to the inflation forecasting literature discussed here, however, because they do not analyze performance relative to a univariate benchmark, nor do they report results for post-1984 subsamples.

9. Cecchetti et. al. (2007, section 7) provided in-sample evidence that survey inflation forecasts are correlated with future trend inflation, measured using the Stock-Watson (2007) UC-SV model. Thus a different explanation of why surveys perform well is that survey inflation expectations anticipate movements in trend inflation.

10. The exceptions are rolling forecasts for the 2001–2007 sample: for CPI-core inflation using median CPI as a predictor, and for GDP inflation using either median or trimmed-mean CPI as a predictor. However, the relative RMSEs exceed one (typically, they exceed 1.15) for other inflation series, other samples, and for recursive forecasts, and we view these three exceptional cases as outliers. Most likely, the difference between our negative results for median CPI and Smith’s (2004) positive results over 1990–2000 are differences in the benchmark model, which in her case is a univariate AR with exponential lag structure imposed.

References

- Ang, Andrew, Geert Bekaert, and Min Wei. 2007. “Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?” *Journal of Monetary Economics* 54(4): 1163–1212.
- Ascari, Guido, and Emanuela Marrocu. 2003. “Forecasting Inflation: A Comparison of Linear Phillips Curve Models and Nonlinear Time Series Models.” Working Paper CRENoS 200307, Centre for North South Economic Research, University of Cagliari, Italy. Available at <http://veprints.unica.it/273/1/03-07.pdf>.
- Atkeson, Andrew, and Lee E. Ohanian. 2001. “Are Phillips Curves Useful for Forecasting Inflation?” *Federal Reserve Bank of Minneapolis Quarterly Review* 25(1): 2–11. Available at <http://www.minneapolisfed.org/research/QR/QR2511.pdf>.
- Bañbura, Marta, Domenico Gianonne, and Lucrezia Reichlin. 2008. “Large Bayesian VARs.” Working Paper No. 966. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp966.pdf>.
- Banerjee, Anindya, and Massimiliano Marcellino. 2006. “Are There Any Reliable Leading Indicators for U.S. Inflation and GDP Growth?” *International Journal of Forecasting* 22(1): 137–151.

- Barnes, Michelle L., and Giovanni P. Olivei. 2003. "Inside and Outside the Bounds: Threshold Estimates of the Phillips Curve." *New England Economic Review*: 3–18. Available at <http://www.bos.frb.org/economic/neer/neer2003/neer03a.pdf>.
- Barsky, Robert B. 1987. "The Fisher Hypothesis and the Forecastability and Persistence of Inflation." *Journal of Monetary Economics* 19(1): 3–24.
- Bates, J.M., and Clive W.J. Granger. 1969. "The Combination of Forecasts." *Operations Research Quarterly* 20: 451–468.
- Berardi, Andrea. 2007. "Term Structure, Inflation, and Real Activity." Manuscript, Department of Economics, University of Verona. Available at <http://dse.univr.it/berardi/tsira.pdf>.
- Bernanke, Ben S. 2003. "An Unwelcome Fall in Inflation?" Speech delivered at the Economics Roundtable, University of California, San Diego, July 23. Available at <http://www.federalreserve.gov/boarddocs/speeches/2003/20030723/>.
- Bernanke, Ben S., Jean Bovian, and Piotr Elias. 2005. "Measuring the Effects of Monetary Policy: a Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120(1): 387–422.
- Boivin, Jean, and Serena Ng. 2005. "Understanding and Comparing Factor-Based Forecasts." *International Journal of Central Banking* 1(December): 117–151. Available at <http://www.ijcb.org/journal/ijcb05q4a4.pdf>.
- Boivin, Jean, and Serena Ng. 2006. "Are More Data Always Better for Factor Analysis?" *Journal of Econometrics* 132(1): 169–194.
- Brainard, William C., and George L. Perry. 2000. "Making Policy in a Changing World." In *Economic Events, Ideas, and Policies: The 1960s and After*, ed. George L. Perry and James Tobin, 43–68. Washington, DC: Brookings Institution Press.
- Brayton, Flint, John M. Roberts, and John C. Williams. 1999. "What's Happened to the Phillips Curve?" Finance and Economics Discussion Series 1999-49. Washington, DC: Board of Governors of the Federal Reserve. Available at <http://www.federalreserve.gov/pubs/feds/1999/199949/199949pap.pdf>.
- Brave, Scott, and Jonas D.M. Fisher. 2004. "In Search of a Robust Inflation Forecast." Federal Reserve Bank of Chicago *Economic Perspectives* 28 (4): 12–31. Available at http://www.chicagofed.org/publications/economicperspectives/ep_4qtr2004_part2_Brave_Fisher.pdf.
- Bryan, Michael F., and Stephen G. Cecchetti. 1994. "Measuring Core Inflation." In *Monetary Policy*, ed. N. Gregory Mankiw, 195–215. Chicago: University of Chicago Press.
- Bryan, Michael F., Stephen G. Cecchetti, and Rodney L. Wiggins II. 1997. "Efficient Inflation Estimation." Working Paper No. 6183. Cambridge, MA: National Bureau of Economic Research.
- Camba-Mendez, Gonzalo, and Diego Rodriguez-Palenzuela. 2003. "Assessment Criteria for Output Gap Estimates." *Economic Modeling* 20(3): 528–561.
- Canova, Fabio. 2007. "G-7 Inflation Forecasts: Random Walk, Phillips Curve or What Else?" *Macroeconomic Dynamics* 11(1): 1–30.

Cecchetti, Stephen G. 1995. "Inflation Indicators and Inflation Policy." In *NBER Macroeconomics Annual* 10, ed. Ben S. Bernanke and Julio J. Rotemberg, 189–219. Cambridge, MA: The MIT Press.

Cecchetti, Stephen G., Rita S. Chu, and Charles Steindel. 2000. "The Unreliability of Inflation Indicators." Federal Reserve Bank of New York *Current Issues in Economics and Finance* 6(4):1–6. Available at http://www.newyorkfed.org/research/current_issues/ci6-4.pdf.

Cecchetti, Stephen G., Peter Hooper, Bruce C. Kasman, Kermit L. Schoenholtz, and Mark W. Watson. 2007. "Understanding the Evolving Inflation Process." Report Prepared for the 2007 U.S. Monetary Policy Forum. Available at <http://research.chicagogsb.edu/igm/docs/2007USMPF-Report.pdf>.

"The Chicago Fed National Activity Index." Various dates. Chicago: Federal Reserve Bank of Chicago. Available at http://www.chicagofed.org/economic_research_and_data/cfnaicfm.

Clark, Todd E., and Michael W. McCracken. 2006a. "The Predictive Content of the Output Gap for Inflation: Resolving In-Sample and Out-of-Sample Evidence." *Journal of Money, Credit and Banking* 38(5): 1127–1148.

Clark, Todd E., and Michael W. McCracken. 2006b; revised 2008. "Combining Forecasts from Nested Models." Research Working Paper 06-02. Kansas City, MO: Federal Reserve Bank of Kansas City. Available at <http://www.kansascityfed.org/Publicat/Reswkpap/PDF/RWP06-02v2.pdf>.

Cogley, Timothy, and Thomas J. Sargent. 2002. "Evolving Post-World War II U.S. Inflation Dynamics." In *NBER Macroeconomics Annual* 2001, ed. Ben S. Bernanke and Kenneth Rogoff, 331–388. Cambridge, MA: The MIT Press.

Cogley, Timothy, and Thomas J. Sargent. 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post World War II U.S." *Review of Economic Dynamics* 8(2): 262–302.

D'Agostino, Antonello, and Domenico Giannone. 2006. "Comparing Alternative Predictors Based on Large-Panel Factor Models." Working Paper Series 680. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp680.pdf>.

Diebold, Francis X., and Roberto S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13(3): 253–263.

DeWachter, Hans, and Marco Lyrio. 2006. "Macro Factors and the Term Structure of Interest Rates." *Journal of Money, Credit, and Banking* 38(1): 119–140.

Diron, Marie, and Benoît Mojon. 2008. "Are Inflation Targets Good Inflation Forecasts?" Federal Reserve Bank of Chicago *Economic Perspectives* 32(2): 33–45. Available at http://www.chicagofed.org/publications/economic_perspectives/ep_2qtr2008_part3_diron_mojon.pdf.

Dupasquier, Chantal, and Nicholas Ricketts. 1998. "Nonlinearities in the Output-Inflation Relationship." In *Price Stability, Inflation Targets and Monetary Policy*, 131–167. Ottawa: Bank of Canada. Available at <http://www.bankofcanada.ca/en/conference/con97/cn97-8.pdf>.

- Eichmeier, Sandra, and Christina Ziegler. 2008. "How Good are Dynamic Factor Models at Forecasting Output and Inflation? A Meta-Analytic Approach." *Journal of Forecasting* 27(3): 237–265.
- Estrella, Arturo, and Jeffrey C. Fuhrer. 2003. "Monetary Policy Shifts and the Stability of Monetary Policy Models." *Review of Economics and Statistics* 85(1): 94–104.
- European Central Bank Inflation Persistence Network. Various Dates. Research papers available at http://www.ecb.eu/home/html/researcher_ipn_papers.en.html.
- Fisher, Jonas D.M., Chin Te Liu, and Ruilin Zhou. 2002. "When Can We Forecast Inflation?" Federal Reserve Bank of Chicago *Economic Perspectives* 26(1): 30–42. Available at http://www.chicagofed.org/publications/economic_perspectives/2002/1qepart4.pdf.
- Gavin, William T., and Kevin L. Kliesen. 2008. "Forecasting Inflation and Output: Comparing Data-Rich Models with Simple Rules." Federal Reserve Bank of St. Louis *Review* 90(3, Part 1): 175–192. Available at <http://research.stlouisfed.org/publications/review/08/05/GavinKliesen.pdf>.
- Geweke, John. 1977. "The Dynamic Factor Analysis of Economic Time Series." In *Latent Variables in Socio-Economic Models*, ed. Dennis J. Aigner and Arthur S. Goldberger, 365–383. Amsterdam: North-Holland.
- Giacomini, Raffaella, and Halbert White. 2006. "Tests of Conditional Predictive Ability." *Econometrica* 74(6): 1545–1578.
- Gordon, Robert J. 1982. "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." In *Workers, Jobs and Inflation*, ed. Martin N. Baily, 89–158. Washington, DC: The Brookings Institution.
- Gordon, Robert J. 1990. "U.S. Inflation, Labor's Share, and the Natural Rate of Unemployment." In *Economics of Wage Determination*, ed. Heinz Konig, 1–34. New York: Springer-Verlag.
- Gordon, Robert J. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 2: 297–333.
- Hansen, Peter Reinhard. 2005. "A Test for Superior Predictive Ability." *Journal of Business and Economic Statistics* 23(4): 365–380. Available at <http://pubs.amstat.org/doi/pdfplus/10.1198/073500105000000063>.
- Hendry, David F., and Michael P. Clements. 2004. "Pooling of Forecasts." *Econometrics Journal* 7(1): 1–31.
- Hendry, David F., and Kirstin Hubrich. 2007. "Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate." Paper presented at Oxford University. Conference in Honour of David F. Hendry, 23 August. Oxford: Oxford University. Available at http://www.economics.ox.ac.uk/hendryconference/Papers/Hubrich_DFHVol.pdf.
- Hubrich, Kirstin. 2005. "Forecasting Euro Area Inflation: Does Aggregating Forecasts by HICP Component Improve Forecast Accuracy?" *International Journal of Forecasting* 21(1): 119–136.

- Inoue, Atsushi, and Lutz Kilian. 2008. "How Useful is Bagging in Forecasting Economic Time Series? A Case Study of U.S. CPI Inflation." *Journal of the American Statistical Association* 103(482): 511–522.
- Jaditz, Ted, and Chera Sayers. 1994. "Predicting Inflation." Manuscript, Bureau of Labor Statistics, Washington, DC.
- Jorion, Philippe, and Frederic Mishkin. 1991. "A Multi-Country Comparison of Term Structure Forecasts at Long Horizons." *Journal of Financial Economics* 29(1): 59–80.
- König, Evan F. 2003. "Is the Markup a Useful Real-Time Predictor of Inflation?" *Economics Letters* 80(2): 261–267.
- Levin, Andrew T., and Jeremy M. Piger. 2004. "Is Inflation Persistence Intrinsic in Industrial Economics?" Working Paper Series 334. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp334.pdf>.
- Lown, Cara S., and Robert W. Rich. 1997. "Is There an Inflation Puzzle?" Federal Reserve Bank of New York *Economic Policy Review* 3(4): 51–69. Available at <http://www.newyorkfed.org/research/epr/97v03n4/9712lown.pdf>.
- Marcellino, Massimiliano. "A Linear Benchmark for Forecasting GDP Growth and Inflation?" *Journal of Forecasting* 27(4): 305–340.
- Mishkin, Frederic S. 1990a. "What Does the Term Structure Tell Us About Future Inflation?" *Journal of Monetary Economics* 25(1): 77–95.
- Mishkin, Frederic S. 1990b. "The Information in the Longer-Maturity Term Structure About Future Inflation." *Quarterly Journal of Economics* 105(3): 815–828.
- Mishkin, Frederic S. 1991. "A Multi-Country Study of the Information in the Term Structure About Future Inflation." *Journal of International Money and Finance* 10(1): 2–22.
- Mishkin, Frederic S. 2007. "Inflation Dynamics." Working Paper 13147. Cambridge, MA: National Bureau of Economic Research.
- Moshiri, Saeed, and Norman Cameron. 2000. "Neural Network Versus Econometric Models in Forecasting Inflation." *Journal of Forecasting* 19(3): 201–217.
- Nelson, Charles R., and G. William Schwert. 1977. "Short-Term Interest Rates as Predictors of Inflation: On Testing the Hypothesis that the Real Rate of Interest Is Constant." *American Economic Review* 67(3): 478–486.
- Orphanides, Athanasios, and Simon van Norden. 2005. "The Reliability of Inflation Forecast Based on Output Gap Estimates in Real Time." *Journal of Money, Credit and Banking* 37(3): 583–600.
- Pivetta, Frederic, and Ricardo Reis. 2007. "The Persistence of Inflation in the United States." *Journal of Economic Dynamics and Control* 31(4): 1326–1358.
- Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policy-makers' Beliefs and U.S. Postwar Stabilization Policy." *The Quarterly Journal of Economics* 121(3): 867–901.

- Reis, Ricardo, and Mark W. Watson. 2007. "Relative Goods' Prices and Pure Inflation." Working Paper No. 13615. Cambridge, MA: National Bureau of Economic Research.
- Roberts, John M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27(4): 975–984.
- Roberts, John M. 2004. "Monetary Policy and Inflation Dynamics." *International Journal of Central Banking* 2(3): 193–230. Available at <http://www.ijcb.org/journal/ijcb06q3a6.pdf>.
- Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90(3): 429–457.
- Rossi, Barbara, and Tatevik Sekhposyan. 2007. "Has Models' Forecasting Performance for U.S. Output Growth and Inflation Changed Over Time, and When?." Manuscript, Duke University. Available at <http://www.econ.duke.edu/~brossi>.
- Rünstler, Gerhard. 2002. "The Information Content of Real-Time Output Gap Estimates: An Application to the Euro Area." Working Paper Series 182. Frankfurt: European Central Bank.
- Sargent, Thomas J., and Christopher A. Sims. 1977. "Business Cycle Modeling Without Pretending to Have Too Much A Priori Economic Theory." Working Paper No. 55. Minneapolis: Federal Reserve Bank of Minneapolis. Available at <http://www.minneapolisfed.org/research/WP/WP55.pdf>.
- Sims, Christopher A. 2002. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity* 2: 1–40.
- Sims, Christopher A., and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review* 96(1): 54–81.
- Smith, Julie K. 2004. "Weighted Median Inflation: Is This Core Inflation?" *Journal of Money, Credit and Banking* 36(2): 253–263.
- Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment, and Monetary Policy." *Journal of Economic Perspectives* 11(1): 33–51.
- Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44(2): 293–335.
- Stock, James H., and Mark W. Watson. 2002. "Forecasting Using Principal Components from a Large Number of Predictors." *Journal of the American Statistical Association* 97(460): 1167–1179.
- Stock, James H., and Mark W. Watson. 2003. "Forecasting Output and Inflation: The Role of Asset Prices." *Journal of Economic Literature* 41(3): 788–829.
- Stock, James H., and Mark W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39(1): 3–34.
- Stockton, David J., and James E. Glassman. 1987. "An Evaluation of the Forecast Performance of Alternative Models of Inflation." *The Review of Economics and Statistics* 69(1): 108–117.

- Taylor, John B. 2000. "Low Inflation, Pass-Through, and the Pricing Power of Firms." *European Economic Review* 44(7): 1389–1408.
- Timmermann, Allan. 2006. "Forecast Combinations." In *Handbook of Economic Forecasting*, Vol. 1, ed. Graham Elliott, Clive W. J. Granger, and Allan Timmermann, 135–196. Amsterdam: North-Holland.
- Tkacz, Greg. 2000. "Non-Parametric and Neural Network Models of Inflation Changes." Working Paper 2000-7. Ottawa: Bank of Canada. Available at <http://www.bankofcanada.ca/en/res/wp/2000/wp00-7.pdf>.
- Watson, Mark W. 2003. "Macroeconomic Forecasting Using Many Predictors." In *Advances in Econometrics: Theory and Applications*, Eighth World Congress, ed. Mathias Dewatripont, Lars Peter Hansen, and Stephen Turnovsky, 87–115. New York: Cambridge University Press.
- Webb, Roy H., and Tazewell S. Rowe. 1995. "An Index of Leading Indicators for Inflation." Federal Reserve Bank of Richmond *Economic Quarterly* 81(2): 75–96. Available at http://www.richmondfed.org/publications/research/economic_quarterly/1995/spring/pdf/webb.pdf.
- West, Kenneth D. 1996. "Asymptotic Inference about Predictive Accuracy." *Econometrica* 64(5): 1067–1084.
- Wright, Jonathan. 2003. "Forecasting U.S. Inflation by Bayesian Model Averaging." International Finance Discussion Paper No. 780. Washington, DC: Board of Governors of the Federal Reserve System. Available at <http://www.federalreserve.gov/pubs/ifdp/2003/780/ifdp780.pdf>.

Data Appendix

The definitions and sources of the series used in this analysis are summarized in the following table. The "trans" column indicates the transformation applied to the series: logarithm (ln), first difference of logarithm $((1-L)\ln)$, accumulation $((1-L)^{-1})$, or no transformation (level). When the original series is monthly, quarterly data are constructed as the average of the monthly values in the quarter before any other transformation. Sources are Federal Reserve Bank of St. Louis FRED database (F), the Bureau of Economic Analysis (BEA), and other Federal Reserve banks as indicated.

Short name	Trans	Definition	Mnemonic (Source)
Inflation series			
CPI-all	(1-L)ln	CPI, all items	CPIAUCSL (F)
CPI-core	(1-L)ln	CPI less food and energy	CPILFESL (F)
PCE-all	(1-L)ln	PCE deflator, all items	PCECTPI (F)
PCE-core	(1-L)ln	PCE deflator, less food and energy	JCXFE (F)
GDP deflator	(1-L)ln	GDP deflator	GDPCTPI (F)
Predictors			
UR	level	Unemployment rate, total civilian 16+	UNRATE (F)
GDP	ln	Real GDP	GDP96 (F)
IP	ln	Index of Industrial Production (total)	INDPRO (F)
EMP	ln	Nonagricultural civilian employment (total)	PAYEMS (F)
CapU	level	Capacity utilization rate	TCU (F)
HPerm	ln	Housing permits (starts)	PERMIT (F)
CFNAI	(1-L) ⁻¹	Chicago Fed National Activity Index (accumulated)	FRB-Chicago
UR-5wk	level	Unemployment rate for unemployed < 5 week	UEMPLT5(F) / CLF160V(F)
AHE	(1-L)ln	Average hourly earnings	AHETPI (F)
Real AHE	(1-L)ln	real average hourly earnings	AHETPI (F)/ GDPCTPI (F)
Labor Share	ln	labor share	AHETPI (F)/ GDPCTPI (F)
CPI-Median	level	Cleveland Fed median CPI inflation “Original” CPI-Median through 2007:M7; “Revised” CPI-Median after 2007:M7)	FRB-Cleveland
CPI-TrMn	level	Cleveland Fed trimmed mean CPI inflation (“Original” CPI-Trimmed Mean through 2007:M7; “Revised” CPI-Trimmed Mean after 2007:M7)	FRB-Cleveland
ExRate	level	trade-weighted exchange rate	TWEXMMTH (F)
TB_sp	level	1 Year Treasury bond rate minus 3 Month Treasury bill rate (at annual rate)	Fed Board of Governors
RPFE	(1-L)ln	Relative Price of Food and Energy	PCECTPI (F)/ JCXFE (F)
RPImp	(1-L)ln	Relative Price of Imports	B021RG3(BEA)/ GDPCTPI(F)
Price Control Variable 1	level	0.8 for 1971:Q3 ≤ t ≤ 1972:Q2, 0 otherwise	Gordon (1982)
Price Control Variable 2	level	-0.4 for t = 1974:Q2 or 1975:Q1, -1.6 for 1974:Q3 ≤ t ≤ 1974:Q4, 0 otherwise.	Gordon (1982)

Comments on “Phillips Curve Inflation Forecasts” by James H. Stock and Mark W. Watson

Adrian Rodney Pagan

Stock and Watson provide a thorough (one might say exhaustive) review of the forecasting performance of many inflation models. These include models with some economics in them and others that are purely statistical. Overall, the best of the statistical models seems to be their unobserved components-stochastic volatility (UC-SV) model. This is a two-component model of the form

$$\pi_t = z_t^* + v_t$$

and

$$z_t^* = z_{t-1}^* + u_t.$$

As written, z_t^* represents a permanent component to inflation and v_t is a transitory one. The shocks to the components are taken as having time-varying stochastic volatilities (SV) σ_{jt}^2 of the form $\ln \sigma_{jt}^2 = \ln \sigma_{jt-1}^2 + \eta_{jt}$. Consequently these have a unit root. The variances of the shock terms η_{jt} are equal. This would mean that the true volatilities would also be equal if the initial conditions were the same. The estimated ones can however vary as they depend upon the data.

Statistical models are generally judged by how well they fit and forecast over limited sample sizes and horizons. That is probably just as well here since it would be hard to think about targeting an inflation rate that really behaved like the UC-SV model, since it has inflation following a unit root and with the variances of v_t and u_t being unbounded, meaning there is no second moment for the change in inflation. If you try to simulate a process like UC-SV, it blows up very quickly.

I guess I am rather doubtful about whether I want to use a model like this, even if it produces good forecasts, as it would be hard to believe

that we could keep inflation in a target range for very long if it behaved in such a way. It is only over short periods that this model makes much sense. Given this reality, I wonder why Stock and Watson set up the UC model with a permanent shock. Indeed, one could have instead chosen a very persistent process such as $z_t^* = .99z_{t-1}^* + u_t$. Since putting any number on the degree of persistence is arbitrary, one might as well use something like .99, because that ties in better with what we hope is the nature of the inflation process. In doing so, no extra coefficients would be estimated and it seems highly likely that over short forecast horizons the forecasts using both sets of parameters would be very close. It might be different if we looked at an eight-period horizon, since most central banks forecast both one and two years ahead.

Stock and Watson conclude that there are periods of time when the UC-SV model can be beaten by models featuring economic variables, principally when there are large departures from the NAIRU or when one is in an extreme recession. But for most of the time the economic variables don't contribute much to forecasts. Of course there are good reasons why we still believe that economic variables are influential, even if we cannot detect a precise role for these variables in forecasts. It is a well-known fact that relatively simple statistical models win forecasting competitions. But in a world in which we are increasingly forced to explain policy actions, any forecasts underlying them, and the risks associated with those forecasts, statistical model forecasts are clearly of limited use. In practice many central banks use a Phillips-curve type equation to give a central forecast—and even judgmental forecasts often have this as a base—relegating the statistical model forecasts to the role of checking and “tweaking” the central forecast. It is hard to imagine any central banker not putting some faith in the role of excess demand in accounting for inflation outcomes, even if measures of excess demand do little for forecasting inflation a year ahead. There are many reasons why we might see this play a role in forecasting inflation. The excess demand may need to be sustained for a long time, many measures of it are exceedingly volatile, inflation itself can have a lot of noise, and it may only be when a threshold is exceeded that there are substantial effects on inflation. These are all hard to measure precisely in a model given the length of the data sets we typically have to work with. But at some

point, economic variables need to enter the forecast—otherwise it is not going to be easy to explain any actions you take as a consequence of the forecasts.

So let us remind ourselves why the UC-SV model might win the Stock and Watson forecasting competition. To do this, assume that there is no SV in the UC model and that the variances of the two shocks u_t and v_t are in a fixed ratio q . Then, it is well known that the forecast from the UC model is (provided $E_t(v_{t+1}) = 0$)

$$E_t \pi_{t+1} = (1 - \phi) \sum_{j=0}^{\infty} \phi^j \pi_{t-j}.$$

The Kalman predictor is used on the UC model to find that ϕ solves $\phi + q\phi^2 - 1 = 0$, q can be found from $q = -2 - \frac{1}{\rho_1}$, and ρ_1 is the first-order serial correlation coefficient of $\Delta\pi_t$. This is the exponentially weighted moving average (EWMA) forecasting formula that is widely used in industry for forecasting the level of product demand. So perhaps it is not surprising that it produces a good forecast for inflation. This formula has also been used recently in the financial literature to forecast series that are random walks in which the drift term changes over time; in other words, the EWMA formula has some robustness to breaks in those series. So it may be a good vehicle for inflation forecasting as well.

To develop this theme further, note that $\rho_1 = \frac{\alpha}{1 + \alpha^2}$, where α is the moving average coefficient in the autoregressive integrated moving average (ARIMA) representation of the UC model, or $\Delta\pi_t = \varepsilon_t + \alpha\varepsilon_{t-1}$. Stock and Watson note that α has been varying a good deal over U.S. history and, recently, $\alpha \cong -.85$, and so $\rho_1 = -.493$, implying that q is close to zero (and ϕ is close to unity), meaning there is little weight placed on past inflation. Indeed a large negative value of α is really consistent with inflation not having a unit root, and the way this shows up in the UC model is for the variance of the transitory shocks to become large relative to the permanent ones.

Now one reason why the EWMA forecasting formula has been popular is that it represents a simple forecasting mechanism that is relatively robust to structural change, provided one modifies q at different times. A plot of the inflation rate suggests that such changes have occurred, and there is an extensive literature maintaining that such breaks have

occurred in many countries in the past three decades, as well as in the United States. The key question then becomes how to vary q in response to such developments, and this is what the SV part of Stock and Watson's UC-SV model does, since the estimated relative volatilities can change over time, leading to a change in q .

So this leads us to ask what the implications of the Stock and Watson paper are for forecasting. This may not be a fair question since their brief seems to have been to ask if Phillips curve-type economic variables are useful for forecasting rather than to ask what is the best forecasting method. So I am possibly being unfair when I ask if their UC-SV model is the best forecasting mechanism, but I think their results are sufficiently striking for me to make some comments on this.

As mentioned above, in Stock and Watson's case q adapts to the data to account for breaks through the relative size of the stochastic volatilities. Are there other ways of doing this? Pesaran and Timmermann (2007) point out that we need to detect when a break in the inflation process took place and also the size of the break; stated differently, we need to know when to vary q and by how much. There has been much research on methods to detect a break in inflation but less on determining the size. However there is now an emerging literature on techniques to gauge the size of the break. Pesaran and Timmermann (2007) have proposed the idea that one should average forecasts not across models (Stock and Watson show that this does not give much advantage) but across the windows over which parameter estimation is performed prior to making the forecast. Pesaran and Timmermann demonstrate that this can yield improvements in forecasting a random walk process in the presence of breaks. Pesaran and Pick (2008) show that there are theoretical reasons to expect that this procedure will improve forecasts in the presence of breaks. They also look at EWMA forecasts for different q values and then average the forecasts from these values. An advantage of focusing upon the EWMA approach is that no judgment is being made about the nature of the inflation process (one still uses a weighted average of inflation rates). If the inflation series was white noise, then one would simply put $q = 0$. So I think it would be interesting to compare the forecasts from this methodology with those from the UC-SV model.

Now let us think about introducing economic information into the forecasts. Traditionally this involved replacing z_t^* with some functions of lagged inflation, unemployment, supply-side effects, deviations of this from the NAIRU, expectations, and so on. In the U.S. literature this is often referred to as Gordon's triangle model, although models like it have been present in many countries since the 1970s. Whatever individual information is introduced needs to be combined together to produce a persistent component. The flaws in the approach are the need to estimate the parameters in any such relation when providing forecasts of these covariates. Consequently, it is possible that we will do worse than a model that ignores the effects, even if we believed that the economics in such a model tells us something about what might have driven an observed inflation path. That will almost certainly happen if the parameters are imprecisely determined, and one would have to say that this is indeed true of most Phillips curve models. Moreover many of the variables added into the relation can change quite dramatically as a result of data revisions. So even if the forecasts are good with data that has been finally revised, the need to use real-time data may result in the forecasts being quite poor—see Robinson, Stone, and van Zyl (2003) for an Australian example. Since Stock and Watson did not use real-time data it would seem that the Phillips curve-based forecasts they report would most likely be better than these would be in real time. It is only with extreme movements in the determining variables—very high unemployment relative to the NAIRU or expectations relative to, say, the target—that we can observe big enough effects on inflation to offset these difficulties. However it should be noted that it may be possible to use economic information to reliably signal the direction of change in inflation, as found in Robinson, Stone, and van Zyl (2003) for Australia and in Fisher, Liu, and Zhou (2002) for the United States, and in many contexts this might well be sufficient.

Finally, we might just make ϕ a function of some economic variables and so change the exponential weights. To get some idea of whether this would work we need a series on ϕ_t . I fitted an AR(1) to the change in inflation using a ten-year rolling horizon to get an estimate of ρ_1 that changes over time. From that I got values for q_t and ϕ_t . Figure 3.18 shows




Figure 3.18
Time-Varying Phi and the Hodrick-Prescott-Filtered Output Gap

the series on ϕ_t obtained with this method. Note that there are some missing observations in the 1982:Q3–1983:Q2 period since the estimate of ρ_1 during that time was positive, which is not admissible if the UC model is correct. Consequently, I have set this particular period's observations to zero to match the movements that were evident before and after that period. The graph is best at identifying changes in ϕ_t rather than the precise values of it, but it shows that the weighting factor on past data to be used when forecasting has varied significantly over history. In the 1982 recession the close-to-zero weights pointed to ignoring the past history of inflation and indicated that some other information needed to be used. The output gap (found with the Hodrick-Prescott filter) in figure 3.18 has a positive correlation with ϕ_t , and so measures of excess demand are likely to be useful for forecasting. This finding reinforces Stock and Watson's conclusion that the Phillips curve was useful in forecasting inflation in the first half of the 1980s.

References

Fisher, Jonas D.M., Chin Te Liu, and Ruilin Zhou. 2002. "When Can We Forecast Inflation?" *Federal Reserve Bank of Chicago Economic Perspectives* Q1: 30–42. Available at http://www.chicagofed.org/publications/economic_perspectives/2002/1qepart4.pdf.

Pesaran, M. Hashem, and Allan Timmermann. 2007. "Selection of Estimation Window in the Presence of Breaks." *Journal of Econometrics* 137(1): 134–161.

Pesaran, M. Hashem, and Andreas Pick. 2008. "Forecasting Random Walks Under Drift Instability." Cambridge Working Papers in Economics No. 0814. Faculty of Economics, University of Cambridge. Available at <http://www.econ.cam.ac.uk/dae/repec/cam/pdf/cwpe0814.pdf>.

Robinson, Tim, Andrew Stone, and Marileze van Zyl. 2003. "The Real-Time Forecasting Performance of Phillips Curves." Economic Research Department Discussion Paper 2003–12. Sydney: Reserve Bank of Australia. Available at <http://www.rba.gov.au/rdp/RDP2003-12.pdf>.

Comments on “Phillips Curve Inflation Forecasts” by James H. Stock and Mark W. Watson

Lucrezia Reichlin

1. Is the Phillips Curve Dead?

Stock and Watson present convincing evidence on the Phillips curve’s lack of predictive power over the last fifteen years—meaning the inability of models based on the Phillips relationship between inflation and unemployment to predict beyond a naïve benchmark such as the random walk.

The analysis is very convincing: it is conducted systematically on the basis of different specifications and definitions of real indicators of economic activity and is based on a simulated out-of-sample exercise. This is the right methodology to evaluate the robustness of the Phillips relation since, unlike in regression analysis, the evaluation results are not only valid under the assumption of the correct specification of the model.

The result is perhaps disturbing to the macroeconomic profession which has spent so much time formulating and discussing micro-foundations for the Phillips curve. Has the profession wasted its time trying to explain a relationship that had in fact disappeared? More constructively, what can we learn about our macroeconomic models from this result?

In this discussion I will bring some complementary evidence to the authors’ results, focusing not only on the predictive ability of the Phillips curve relationship but also on the predictability of inflation and real activity in general. The evidence I will present suggests a decline of relative predictability not only for inflation, but also for real economic activity and for statistical as well as institutional models such as the Greenbook forecasts prepared by the staff of the Federal Reserve Board.

I will then ask whether the decline in predictability can be explained by a change in the structure of the covariance between different variables in the economy, or if this decline is the result of the decline of the variability of exogenous shocks. To analyze this question I will propose a quantitative exercise.

Clearly, to the extent that evidence points to a change in the covariances, the decline in predictability should be attributed to changes in policy or structural features of the economy, occurring around the mid-1980s. In that case the way forward must be to build on these results and try to understand the breakdown of the Phillips curve as the endogenous result of changes in structural and policy parameters in structural models.

2. Decreasing Relative Predictability

In what follows I will focus on relative predictability as defined by the expression:

$$RP_{it} = \frac{\hat{E}[X_{it+bt}^{model} - X_{it+b}^2]^2}{\hat{E}[X_{it+bt}^{naive} - X_{it+b}^2]^2}.$$

Here, relative predictability is defined as the predictive ability of a given model relative to the prediction based on a random walk model and measured in terms of mean squared errors.

Recent literature has pointed to a decrease in the relative predictability of inflation. The evidence is accurately surveyed by Stock and Watson's present paper and I have little to add.

Perhaps the best way to understand the evidence is to inspect figure 3.19, which plots the GDP-deflator inflation rate (solid-diamond line) against the naïve (random walk) forecast (dashed line) and the Greenbook forecast (solid line) since 1970; the shaded areas indicate recession episodes as defined by the National Bureau of Economic Research (NBER).

The picture illustrates two points. First, since the early 1990s, when Stock and Watson's sample starts, the naïve predictor becomes more accurate due to the decline in inflation volatility. Second, for the same sample, the Greenbook and the naïve forecasts are very similar—no clear




Figure 3.19

Greenbook One-Year-Ahead Forecasts of GDP Deflator Annual Inflation

Source: Author's calculations.

advantage seems to be obtained by a sophisticated forecast such as that produced by the staff of the Federal Reserve Board.

The analysis by Stock and Watson shows that the same is true for predictive equations based on the Phillips curve. The same authors have pointed out in previous work (Stock and Watson 2007) that the result applies no matter what variables are considered.

In work focusing on complex models based on a large number of predictors, nominal and real, De Mol, Giannone, and Reichlin (2008) show that principal component regression (PC), ridge regression (Ridge) and variable selection algorithms (Lasso) all produce forecasts with no relative advantage with respect to the random walk for that sample. These are relatively complex statistical models which perform very well out-of-sample until the mid-1980s. Figure 3.20 shows the forecast for the annual rate of change of Consumer Price Index (CPI) inflation since 1970 (again, the shaded areas indicate NBER-dated recessions). Clearly, in the last 20 years, the sophisticated models have not outperformed the




Figure 3.20

One-Year-Ahead December Forecasts of Yearly Changes of Consumer Price Index Annual Inflation

Source: De Mol, Giannone, and Reichlin (2008).

naïve model. Notice that the Lasso, PC, and Ridge forecasts are highly correlated throughout the sample and hardly distinguishable from one another.

But these results are not only true for inflation. These results remain valid for real economic activity as well, as pointed out by D'Agostino, Giannone, and Surico (2006) and by De Mol, Giannone, and Reichlin (2008). In particular, D'Agostino, Giannone, and Surico (2006) emphasize that the interest rates term structure, which was a good predictor until the mid-1980s, has failed ever since. This is an interesting result since the term structure is typically thought of being a forward-looking variable, capturing expectations of future economic activity.

For real activity, figures 3.21 and 3.22 show similar features to those described for inflation. Figure 3.21 plots the annual growth rate of GDP (solid-diamond line) since 1970, the forecast based on the Greenbook




Figure 3.21

Greenbook One-Year-Ahead Forecasts for GDP Annual Growth

Source: Author's calculations.

(solid line) and the random walk (naïve) model (dashed line). Here the random walk is a model in which growth in the next period is equal to the average growth in the previous ten years.

Figure 3.22, from De Mol, Giannone, and Reichlin (2008), reports the annual change of industrial production, the naïve forecast (Naïve) again defined as a model in which growth in the next period is equal to the average growth in the previous ten years, the principal component forecast (PC), the forecast based on ridge regression (Ridge) and variable selection (Lasso) algorithms, as also shown in figure 3.20.

What I conclude from this evidence is that the Phillips curve is not the only predictive relationship that has broken down in the last 20 years. In general, we have experienced a failure of models to predict beyond a naïve benchmark, for both inflation and output.

How can we interpret this evidence? One conjecture is that the last 20 years have been a lucky period with very moderate volatility of exogenous shocks to the economy. Low volatility in exogenous shocks has




Figure 3.22

One-Year-Ahead December Forecasts for Annual Growth of Industrial Production

Source: De Mol, Giannone, and Reichlin (2008).

implied low volatility of observable variables, and in these circumstances it is not surprising that the random walk model has performed relatively well.

Alternatively, low volatility of output and inflation can be attributed to a change in the structure of the economy or to changes in policy, as extensively discussed in the literature about the Great Moderation (for a review of that debate and some new results, see Giannone, Lenza, and Reichlin 2008).

Actually, as Giannone, Lenza, and Reichlin (2008) have pointed out, decreases in inflation and output volatility combined with a decrease in the relative predictability of statistical bivariate and multivariate models and institutional models can only be explained by a change in the covariance structure of the data. Moreover, given the predictive failure of mod-

els based on many variables, including financial variables, the change must not only be in the covariance between inflation and real activity but, more generally, in the covariance between financial variables, inflation, and real activity. Finally, as we have shown in Giannone, Lenza, and Reichlin (2008), these changes can best be identified using models containing all relevant macroeconomic variables and might be lost in small models due to problems arising from omitted variables.

3. Shocks or Propagation? A Counterfactual Exercise

In this section I will ask what is the fraction of the decline of relative predictability in output and inflation that can be attributed to a decline in the variability of exogenous shocks, and what is that fraction that can be attributed to a change in the parameters of a model that include several macroeconomic indicators—among these prices, real activity, labor market, and monetary and financial indicators.

I will consider a vector autoregression (VAR) model, including 19 quarterly variables, all of which are typically used in macroeconomic models: GDP, the GDP deflator, the federal funds rate, commodity prices, consumer prices, consumption, investment, change in inventories, the producer price index, interest rates at one-, five-, and ten-year horizons, hours worked, hourly compensation, capacity utilization, stock prices, M2, total reserves, and the unemployment rate. This is the same model estimated in Giannone, Lenza, and Reichlin (2008). I refer to that paper for details on estimation and exact definitions and sources.

The VAR is estimated for two subsamples: 1959:Q1 to 1983:Q4, and 1984:Q1 to 2007:Q1.¹ Since the model is quite large and I face an issue of over-fitting, I follow Banbura, Giannone, and Reichlin (2008) and use Bayesian shrinkage. In practice I use a Litterman (random walk) prior whose tightness is set so that the in-sample fit of the interest rate equation in the large VAR models is fixed at the level achieved by a simple four-variable monetary VAR. This choice is grounded on the evidence that U.S. short-term interest rates are well described by linear functions of inflation and real activity—Taylor rules (on this point see Giannone, Lenza, and Reichlin 2008).

The models are:

$$\Delta X_t = A_{pre84}(L)\Delta X_{t-1} + e_{pre84,t} \quad e_{pre84,t} \sim WN(0, \Sigma^{pre84}),$$

$$\Delta X_t = A_{post84}(L)\Delta X_{t-1} + e_{post84,t} \quad e_{post84,t} \sim WN(0, \Sigma^{post84}).$$

The counterfactual exercise consists in simulating the shocks, assuming that their covariance matrix has remained unchanged at the level of the pre-1984 sample estimates ($\hat{\Sigma}_{pre84}$) and feeding them through the propagation mechanism estimated for the post-1984 sample ($\hat{A}_{post84}(L)$). Specifically, we consider the following counterfactual processes:

$$\Delta X_t^* = \hat{A}_{post84}(L)\Delta X_{t-1}^* + e_{pre84,t}^*, \quad e_{pre84,t}^* \sim WN(0, \hat{\Sigma}_{pre84}).$$

If the counterfactual relative predictability of GDP (or inflation) is the same as the actual standard deviation observed in the post-1984 sample, then this should indicate that the change of propagation mechanisms fully explains its decline. The change in shocks plays a role if, instead, the counterfactual decline is smaller than observed.

Reported in table 3.6 below, the results are unambiguous: for both inflation and GDP, the change in propagation explains all the decline in variance and all the decline in predictability.

Table 3.6
Counterfactual Volatility and Relative Predictability

Coefficients	Shocks	Std. Deviation		Predictability	
		GDP growth	Inflation	GDP growth	Inflation
Observed					
Pre-1984	Pre-1984	2.68	2.66	0.18	0.12
Post-1984	Post-1984	1.28	0.75	0.36	0.31
Counterfactual					
Post-1984	Pre-1984	1.30	0.69	0.47	0.33

4. Conclusion

In the last 20 years, as Stock and Watson convincingly show, the relative predictability of the Phillips curve has broken down. My discussion points out that in the same period, we have experienced a decline in the relative predictability of inflation and real activity in general. The empirical analysis I proposed suggests that these changes are attributable to changes in the multivariate covariance structure of the data.

This conclusion tells us that one direction for future research should be to study predictability as a function of the deep parameters of structural models, characterizing either structural features of the model or policy behavior. In particular, the result first presented by D'Agostino, Giannone, and Surico (2006) showing that, since the mid-1980s, the spread between short- and long-term interest rates has lost its predictive power for real activity, suggests that an important factor for declining predictability might have been changes in monetary policy which, by anchoring expectations, have broken down the predictive relation between forward-looking variables and real activity. Clearly more work is needed to explore these mechanisms.

Note

1. The models are estimated with data in log-levels except for interest rates, capacity utilization, unemployment rates, and changes in inventories, for which we do not take logarithms.

References

- Bañbura, Marta, Domenico Giannone, and Lucrezia Reichlin. 2008. "Large Bayesian VARs." Working Paper No. 966. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp966.pdf>.
- D'Agostino, Antonello, Domenico Giannone, and Paolo Surico. 2006. "(Un) Predictability and Macroeconomic Stability." Working Paper Series 605. Frankfurt: European Central Bank. Available at <http://www.ecb.int/pub/pdf/scpwps/ecbwp605.pdf>.

De Mol, Christine, Domenico Giannone, and Lucrezia Reichlin. 2008. "Forecasting Using a Large Number of Predictors: Is Bayesian Shrinkage a Valid Alternative to Principal Components?" *Journal of Econometrics* 146(2): 318–328.

Giannone, Domenico, Michele Lenza, and Lucrezia Reichlin. 2008. "Explaining the Great Moderation: It is Not the Shocks." *Journal of the European Economic Association* 6(2–3): 621–633.

Stock, James H., and Mark W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39(1): 3–34.

4

The Labor Market and the Phillips Curve

A New Method for Estimating Time Variation in the NAIRU

William T. Dickens

The non-accelerating inflation rate of unemployment (NAIRU) is frequently employed in fiscal and monetary policy deliberations. The U.S. Congressional Budget Office uses estimates of the NAIRU to compute potential GDP, that in turn is used to make budget projections that affect decisions about federal spending and taxation. Central banks consider estimates of the NAIRU to determine the likely course of inflation and what actions they should take to preserve price stability. A problem with the use of the NAIRU in policy formation is that it is thought to change over time (Ball and Mankiw 2002; Cohen, Dickens, and Posen 2001; Stock 2001; Gordon 1997, 1998). But estimates of the NAIRU and its time variation are remarkably imprecise and are far from robust (Staiger, Stock, and Watson 1997, 2001; Stock 2001).

NAIRU estimates are obtained from estimates of the Phillips curve—the relationship between the inflation rate, on the one hand, and the unemployment rate, measures of inflationary expectations, and variables representing supply shocks on the other. Typically, inflationary expectations are proxied with several lags of inflation and the unemployment rate is entered with lags as well. The NAIRU is recovered as the constant in the regression divided by the coefficient on unemployment (or the sum of the coefficient on unemployment and its lags).

The notion that the NAIRU might vary over time goes back at least to Perry (1970), who suggested that changes in the demographic composition of the labor force would change the NAIRU. He adjusted the unemployment rate to account for this. By 1990 several authors, including Gordon (1990) and Abraham (1987), had suggested that the NAIRU was probably lower in the 1960s than in the 1970s and 1980s. This

adjustment was initially accommodated by adding dummy variables or splines for certain periods to the Phillips curve regression. However, when it began to appear that the U.S. NAIRU was coming down in the 1990s, new methods were developed to track its changes. Staiger, Stock, and Watson (1997), Gordon (1997, 1998) and Stock and Watson (1999) applied time-varying coefficient models and structural break models to NAIRU estimation, and typically found evidence that it rose in the late 1960s or early 1970s and declined in the 1990s.¹ However, the timing and the magnitudes of the estimated changes differed markedly depending on the specification used. Furthermore, confidence bounds on the estimated NAIRUs were so large that the estimates had little value for policy.²

This paper presents a new approach to estimating time variation in the NAIRU. A major problem with Phillips curve-based estimates is that the complicated relationship between inflation, its own lags, supply shocks, and unemployment and its lags makes it possible to explain any particular incidence of high or low inflation a number of different ways. This problem is the root cause of both the lack of robust results and the large confidence intervals around NAIRU estimates derived from Phillips curve estimates. This paper explores an alternative source of information about time variation in the NAIRU. To the extent that such changes are due to changes in the efficiency of the labor market, these changes are reflected not just in the relationship between inflation and unemployment, but also in the relationship between unemployment and job vacancies. That relationship is much simpler and consequently much easier to model in a robust fashion. Combined estimates of the Phillips curve and Beveridge curve—the relationship between unemployment and vacancies—yield remarkably consistent estimates of the timing of changes in the NAIRU.

The next section provides a brief introduction to the literature on the Beveridge curve and on how it has shifted over time. It argues that because the Beveridge curve is much simpler and potentially better fitting than the Phillips curve, it provides a better basis for discerning shifts in the efficiency of the labor market's functioning. These shifts appear to be quite large. The second section develops a theory linking shifts in the Beveridge curve to shifts in the NAIRU. The third section presents estimates of the Beveridge curve model developed in the second section. These esti-

mates turn out to be very robust and motivate the model developed in the fourth section.

The fourth section presents estimates of a linearized version of the model using a Kalman filter. The filtered series is essentially a weighted average of the residuals of the Beveridge and Phillips curves that has been scaled to satisfy an identifying constraint—this constraint is that the coefficient on the filtered variable must be the same as minus the coefficient on the unemployment rate in the Phillips curve. As might be expected, given how precisely the Beveridge curve is estimated, the filter puts nearly all the weight on the Beveridge curve residuals. Estimates of a restricted version of the model suggest that the information from the Beveridge curve adds significantly to the explanatory power of the Phillips curve. The Beveridge curve and Phillips curve NAIRUs look fairly similar, a result which supports the theory behind both curves. Confidence intervals that account for both forecast and parametric uncertainty are about 40 percent larger for Phillips curve NAIRU series than for series derived from the combined Beveridge curve-Phillips curve model. While estimates of the magnitude of the fluctuations in the NAIRU based on the joint Beveridge curve-Phillips curve model are still fairly uncertain, there is little uncertainty about the timing of the fluctuations.

1. The Beveridge Curve

The Beveridge curve describes a convex relationship between job vacancies and unemployment. It is named after Lord William Beveridge, reflecting his work defining full employment in terms of the ratio of unemployment to job vacancies (Beveridge 1945, pp. 18–20). Hansen (1970) was the first to propose a formal model to explain the nature and shape of the relationship based on disequilibrium in two labor markets. Blanchard and Diamond (1989) offer an alternative model based on a matching function.

Until recently there was no vacancy data for the United States, but starting with Abraham (1983) the Conference Board's help-wanted index has been used to construct a proxy for the number of vacancies in several studies (Abraham 1983, 1987; Blanchard and Diamond 1989; Bleakley and Fuhrer 1997; Medoff 1983; Valetta 2006). Abraham (1983) argued for several adjustments to the help-wanted index to take account of

changes in the structure of the newspaper industry and changes in the use of advertising by business in response to equal opportunity laws and regulations. However, Zargorsky (1998) provides convincing evidence that, except for an adjustment for scale (due to Konstant and Wingard 1968), up to at least 1994 the help-wanted index tracks vacancies well without any adjustment.

Sometime after 1994 this relationship between job vacancies and unemployment falls apart. By the time the Bureau of Labor Statistics began conducting the Job Openings and Labor Turnover Survey (JOLTS) series in December 2000, measures of the vacancy rate constructed from the help-wanted index were running well below the numbers coming out of the JOLTS series. This divergence could have been anticipated given the explosive growth of the Internet as a way for people to find and apply for jobs. Monster.com, one of the first job-matching services on the Internet, started operating in 1994. At that time it listed only 200 jobs (Hernandez 2008). In late 1998 it was still listing only 50,000 jobs. But by May 1999 Monster.com was listing 204,000 jobs, held more than 1.3 million active resumes, and was recording 7.6 million hits per month (Answers.com, 2008). By July 2002 Monster was receiving over 14 million unique hits per month (Hernandez 2008). Today, sites like Careerbuilder.com and Monster.com are only some of many ways that workers connect to job openings through the Internet. Many companies' web sites advertise employment opportunities, and employment agencies use the Internet to troll for jobs and workers to fill them. There is no doubt that a smaller fraction of jobs are listed in newspaper help-wanted advertisements today than was the case 15 years ago.

In the work that follows, the help-wanted index is used with only a scale adjustment, as suggested by Zargosky (1998). From the above it seems likely that this index remains a reliable measure of job vacancies at least up to the end of 1997, but probably not much beyond that point. Thus the years 1998–2000 are dropped in the work presented here. After 2000 the JOLTS data are used to measure the number of vacancies.

Figure 4.1 presents a plot of the vacancy rate (vacant jobs over labor force) versus the unemployment rate from 1954 to 1997 and from 2001 to 2007. Two things are apparent in the graph. First, the Beveridge curve is by no means a stable trade-off. The same vacancy rate was associated with a much higher rate of unemployment in the 1980s than in the




Figure 4.1
 Beveridge Curve 1954–1996 and 2001–2007
 Source: Author’s computations.

1950s and 1960s. The trade-off in the 1990s and 2000s seems to have improved notably. Starting with Abraham (1987) several authors have offered theories to explain these movements (Cohen, Dickens, and Posen 2001; Valletta 2006).

While the trade-off moves around quite a bit, there do appear to be long periods in which the relationship is relatively stable. From 1958 through 1970 the vacancy and unemployment rates move back and forth in a relatively tight band. There is a similar period from about 1975 through 1986, then another from 1989 to 1997, and then again starting in 2001. The relationship between vacancies and unemployment over these different periods looks remarkably similar. As a result, detecting the magnitude and timing of shifts in the position of the relationship is relatively easy.

If these changes do reflect changes in the efficiency of the functioning of the labor market then these should correspond to large changes in the NAIRU. This possibility provides the point of departure for this paper. What is needed is a theory to guide the measurement of the movements in the vacancy-unemployment relationship and to translate it into movements in the NAIRU.

2. The Model

The model is an extension of Blanchard and Diamond's (1989) continuous-time labor market model modified to yield a NAIRU. Each firm in the economy hires one worker and faces a nominal price for its product at time t , the natural log of which for firm i is denoted

$$(1) \quad p_i(t) = p(t) + z_i(t),$$

where $p(t)$ is the natural log of the aggregate price level³ at time t and $z_i(t)$ is the natural log of the real price entrepreneur i faces. While $p(t)$ changes continuously, the values $z_i(t)$ change in jumps that take place at a rate s . When these change, a new $z_i(t)$ is drawn from a uniform distribution with support on the interval $[a, b]$.

Firms know that the natural log of their real costs for production (including the expected amortized cost of capital) will be w , where $(a < w < 0 < b)$, but they do not know the current price level. Thus they do not know the real profits they will be able to make should they choose to produce. Before this information is revealed they must make an irreversible purchase of capital (though they do not have to pay for the capital until it is delivered, and delivery can be delayed till a worker is hired if this is a new firm). Thus, both currently active firms and new firms will decide to produce when faced with a new price if

$$(2) \quad p_i(t) - [p(t) + e(t)] = z_i(t) - e(t) > w,$$

where $e(t)$ is the error in their perception of the log of the current price level common to all entrepreneurs and, thus, the term in brackets is their perception of the log of the price level at time t . Active firms (those who currently employ a worker) can continue employing the same worker once the new capital investment is made. New firms must post a vacancy and wait to find a worker before they can begin producing.

If we now assume that $b - a = 1$, then a fraction

$$(3) \quad F = \min(1, \max([(1 - w) - e(t)], 0))$$

of active firms facing new prices will choose to continue to operate and a fraction $1 - F$ will cease to operate. A fraction F of new firms will choose to post a vacancy while $1 - F$ choose not to post a vacancy and dissolve, as do operating firms that receive a new price and decide not to continue operating.

It is assumed that the capital cost is sufficiently large relative to the largest possible error in perception that true real prices will never be less than variable costs. Thus, firms that decide to post a vacancy or operate will continue to do so at least until they receive a new price, even if they have underestimated the real price as they are still covering a fraction of the cost of capital.

It is next assumed that unemployed workers are matched with vacant jobs at a rate $M(U, V)$, where U and V are the number of unemployed workers and vacancies respectively. M is assumed to be homogenous degree 1 with $dM/dU > 0$ and $dM/dV > 0$. The labor force contains L workers so that the equations of motion for the vacancy rate and the unemployment rate are given by⁴

$$(4) \quad dV = cgJ^*F - cV(1 - F) - M(U, V)$$

and

$$(5) \quad dU = c(L - U)(1 - F) - M(U, V).$$

New potential firms are created at a rate cgJ^* , where c is the constant rate at which old firms receive new prices, and g and J^* are constants to be defined later. New vacancies are thus created at the rate cgJ^*F (the first term in equation 4). Vacancies disappear when workers are matched to those vacancies (the last term in equation 4) or when a firm with a posted vacancy receives a new price perceived as being too low to be profitable (the second term in equation 4). Workers become unemployed when their firm receives a new price that is perceived to be unprofitable (the first term in equation 5) and leave unemployment when matched with a job (the second term in equation 5).

A permanent increase in F will cause a permanent increase in the number of vacancies and a decline in the number of unemployed, while a decline will have the opposite effects. Following Blanchard and Diamond (1989), this equilibrium locus is defined as the Beveridge curve. The equation that defines it implicitly can be found by setting dV and dU to zero and substituting F out of (4) and (5). Doing this and dividing by the number of workers in the labor force squared, L^2 , yields

$$(6) \quad (1 - u) = [1 + j/(g j^*)] m(u, v) / c,$$

where lowercase letters denote the value of their uppercase counterpart divided by L and $j = v + 1 - u$, or the ratio of jobs to workers.

The long-run equilibrium of the model is defined as the values of v and u that are obtained when $e(t) = 0$. Since the right-hand sides of both (4) and (5) must equal zero in equilibrium, if we set them equal to each other we see that equilibrium also implies $j(1 - F) = g j^* F$. So if we normalize j^* to equal the equilibrium value of j we get that $g = (1 - F)/F$, which in equilibrium equals $w/(1 - w)$. Thus, the term in brackets in equation (6) becomes $1/w$. Note also that in equilibrium if $v + 1 - u = j^*$ then

$$(7) \quad v = (j^* - 1) + u.$$

Together (6) and (7) determine the long-run equilibrium values of v and u —the latter being the NAIRU.

Figure 4.2 plots examples of equation (6) and equation (7) showing how the NAIRU is derived. Equation (6) has the familiar convex shape associated with the Beveridge curve.⁵ Equation (7) is a 45-degree line, the intercept of which is equal to the excess of the vacancy rate over the unemployment rate in equilibrium (i.e., $j^* - 1$).

3. Estimating the Vacancy-Unemployment Relationship

To obtain a Beveridge curve equation that can be estimated, equation (6) must be linearized in logs. Approximating $m(v, u) = A(t) v^b u^{1-b}$, (6) can be rewritten as

$$(6') \quad \ln\left(\frac{1-u}{u}\right) = \ln(A(t)/c) + b \ln(v/u) + \ln\left[1 + \frac{j}{g j^*}\right].$$

Treating the last term as an error term yields the Beveridge equation to be estimated

$$(6'') \quad \ln\left(\frac{1-u_t}{u_t}\right) = A'_t + b \ln(v_t/u_t) + \mu_t,$$

where A'_t is an appropriately scaled time-varying parameter that reflects changes in the efficiency of the matching process. The final term, the log of one plus the ratio of jobs to the number of jobs in equilibrium divided by g , should vary only very slightly compared to the log of the ratio of the vacancy to the unemployment rate.




Figure 4.2
 Determination of the NAIRU from the Beveridge Curve
Source: Author's computations.

Although equation (6") specifies a single variable linear relationship between $\ln((1 - u)/u)$ and $\ln(v/u)$, it cannot be estimated directly with ordinary least squares. The $A'(t)$ term is time-varying, and from inspection of figure 4.1 there is good reason to believe that the variation in that term would be correlated with the v/u ratio. It would bias the estimates of the coefficient of $\ln(v/u)$ if the time-varying component of $A'(t)$ was treated as a component of the regression error term. Further, the ratio of jobs to the number of jobs in equilibrium will be positively correlated with $\ln(v/u)$, which will tend to bias the estimate of the coefficient of $\ln(v/u)$ downward (though probably only slightly).

From figure 4.1 it appears that the variation in A' is at a much lower frequency than the movement along the Beveridge curve that is reflected in the co-movement of $\ln((1 - u)/u)$ and $\ln(v/u)$. Three different approaches are taken to removing this low frequency variation. First, equation (6") is estimated using only subperiods where the $v - u$ relationship seems stable based on inspection of figure 4.1. Second, the low frequency variation is filtered out of the data and the model estimated only on the filtered data. Finally, both the left- and right-hand sides of (6") are first differenced.

With the low frequency variation in A' removed, the relationship in (6'') should fit well if the approximations used to construct it are good.

In fact, the relationship between the high frequency variation in the left- and right-hand sides of equation (6'') are remarkably well described by a simple linear relationship as can be seen in figure 4.3. In the bottom panel, the differenced data are plotted against each other. In the top panel data that have been passed through a 25-quarter centered moving average filter are plotted against each other. In this case an unemployment rate that has been age-adjusted, as in Shimer (1999), is used rather than the total unemployment rate. The R^2 s for both regressions are .90 or higher as the observations are tightly packed around a line with a slope that is only slightly larger than .50—the value one would expect if unemployed workers and job vacancies had the same impact on the matching rate.


Nor are these two relationships atypical. Table 4.1 presents 32 different estimates of the coefficient of $\ln(v/u)$ using two different measures of unemployment (age-adjusted unemployment on the top half and total unemployment on the bottom half) and a number of different methods to remove the low-frequency variation. The instrumental variables (IV) estimates are constructed using four lagged values of the log of the vacancy rate.⁶ All of the estimated values of b fall in the interval from .45 to .56 and all are precisely estimated. It is also worth noting that the IV estimates do not vary much from the ordinary least square estimates. There is simply too little error in the relationship for endogeneity of the right-hand-side variable to matter.

4. Joint Estimation of the Phillips and Beveridge Curves

The estimation done in this section proceeds under the assumption that a single unobserved factor moves both the equilibrium ratio of jobs to workers (j^*) and the constant $A'(t)$ in the Beveridge curve, and that the relationship is deterministic. If both are arbitrary functions of that unobservable variable, equation (7) is substituted into equation (6), and j is set equal to j^* (as it is by definition in equilibrium). Equation (6) implicitly defines the NAIRU as a function of the unobservable driver. Inverting that function, linearizing it, and substituting it for the unobservable in

Filtered Model 25-Quarter Moving Average Using Age-Adjusted Unemployment

$$R^2 = 0.97, b = 0.53$$



First Difference Model, no constant, Using Total Employment

$$R^2 = 0.902, b = 0.542$$




Figure 4.3

Two Examples of Model Fit

Source: Author's computations.

Note: The sample periods used are 1954–1996 and 2001–2007.

Table 4.1
Alternative Estimates of the Vacancy-Unemployment Relationship

	First Difference		Filtered (two-sided MA)				Stable Time Periods			
	With Constant	No Constant	9 qtr. window	17 qtr. window	25 qtr. window	1958:Q3– 1968:Q4	1975:Q2– 1985:Q4	2001:Q1– 2007:Q2	2001:Q1– 2007:Q2	
Age-adjusted Unemployment	OLS Coef	0.53	0.54	0.53	0.53	0.52	0.52	0.52	0.47	
	std. error	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	
IV	Coef	0.56	0.54	0.52	0.52	0.53	0.50	0.45	0.45	
	std. error	0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.01	
Total Unemployment	OLS Coef	0.54	0.54	0.53	0.53	0.50	0.50	0.49	0.49	
	std. error	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	
IV	Coef	0.56	0.55	0.52	0.52	0.50	0.50	0.49	0.49	
	std. error	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	

Note: The regressions using first differenced data use both the JOLTS and the help-wanted index. The regressions using filtered data use only the help-wanted index data from 1955:Q1–1996:Q4; stable time periods 1958:Q3–1968:Q4 and 2001:Q1–2007:Q2; help-wanted index, 2001:Q1–2007:Q2, JOLTS.

The instrumental variables used are four lags of the vacancy rate for the stable time period regressions, four lags of the change in the vacancy rate are used for the first-difference regressions and four lags of the filtered vacancy rate for the filtered regressions.
Source: Author's computations.

the determination of $A'(t)$ yields the Beveridge curve equation to be estimated

$$(8) \quad \ln\left(\frac{1-u_t}{u_t}\right) = A'' - au_t^* + b'\ln(v_t / u_t) + \mu_t,$$

where u^* is the NAIRU, which is unobservable.

A standard price-price Phillips curve is also estimated of the form

$$(9) \quad \pi_t = \sum_{i=S_u}^{L_u} c_i(u_t^* - u_{t-i}) + \sum_{i=1}^{L_{inf}} d_i \pi_{t-i} + x_t' h + \varepsilon_t,$$

where π_t is inflation in quarter t , L_u and L_{inf} are the number of lags of unemployment and inflation included respectively, S_u is either 0 or 1 depending on whether contemporaneous unemployment is included in the equation, x_t is a vector of dummy variables capturing supply shocks, and h is a conforming vector of coefficients. The coefficients d_i are constrained to sum to 1 so that u^* can be interpreted as the NAIRU in the absence of any observed supply shocks.

It is assumed that the NAIRU u^* evolves as a random walk with an innovation that is independent of the innovations in equations (8) and (9). To identify the model it is further assumed that $\text{cov}(\mu_t, \varepsilon_t) = 0$ so that the only source of correlation between the unobservables in equations (9) and (10) is the common NAIRU.

The model is estimated using a Kalman filter. The constraint that the NAIRU must have the same coefficient as the unemployment rate in the Phillips curve, and the restriction on the covariances of the error terms, are adequate to completely identify all the model parameters. The approach used here is similar to that taken by Basistha and Startz (2008) to estimating the NAIRU with multiple indicators.

Table 4.2 presents six different estimates of the Beveridge curve-Phillips curve model. The first column presents a specification using data from 1955:Q1 to 1997:Q4 and from 2001:Q1 to 2008:Q3 using the Consumer Price Index (CPI) as the inflation measure, and the civilian unemployment rate as the unemployment measure. Included in the Phillips curve equation are contemporaneous unemployment, twelve lags of the inflation rate, and three lags of the unemployment rate. All the parameters of the model are estimated with a fair degree of precision and

Table 4.2
Kalman Filter Estimates of Beveridge Curve-Phillips Curve System and Ordinary Least Squares Phillips Curve with Kalman NAIRU

Parameter	Specification					
	Beveridge curve error variance unconstrained			Beveridge curve error variance constrained to zero		
	55-97 01-08 12 lags CPI unemp. + 3 lags	60-95 01-08 4 lags GDP consumption deflator age adj. unemp.	55-96 8 lags GDP deflator 2 lags unemp.	55-97 01-08 12 lags CPI unemp. + 3 lags	60-95 8 lags core CPI age-adj. unemploy.+ 2 lags	65-96 1 lag GDP deflator 4 lags unemp.
Phillips curve						
Sum coefficient unemployment	-.73 (.21)	-.30 (.09)	-.40 (.11)	-.73 (.21)	-.51 (.20)	-.16 (.13)
Average** Kalman gain	1.5E-7	3E-8	3E-7	0*	0*	0*
Beveridge Curve						
Constant	4.57 (.38)	4.14 (.21)	4.59 (.42)	4.56 (.37)	4.22 (.27)	3.96 (.43)
NAIRU	-19.9 (6.2)	-13.4 (3.5)	-20.4 (6.9)	-19.8 (6.1)	-14.8 (4.5)	-10.7 (6.6)
Ln(v/u)	.52 (.01)	.50 (.01)	.52 (.01)	.51 (.01)	.50 (.01)	.50 (.01)
Post-2000 dummy	-.20 (.09)	-.09 (.08)		-.19 (.08)		
Average** Kalman gain	-.059	-.080	-.057	-.057	-.080	-.051
s.d. NAIRU innovation	.001 (.0003)	.001 (.0003)	.001 (.0004)	.001 (.0003)	.001 (.0003)	.001 (.0009)
Log Likelihood	1063.3	1029.5	951.4	1060.7	783.6	732.4
NAIRU in Phillips curve						
Constant	.004 (.011)	.008 (.007)	.000 (.008)	.002 (.012)	.014 (.011)	.004 (.007)
NAIRU	.86 (.37)	.66 (.42)	.90 (.43)	.85 (.37)	.55 (.47)	.57 (.99)
Chi-squared test Constant = 0 NAIRU=1	.40 [.82]	1.15 [.56]	.26 [.88]	.41 [.81]	1.26 [.53]	.60 [.74]

Standard errors in parenthesis, significance levels in square brackets.

*Constrained to zero by assumption that Beveridge curve error variance equals zero.

** Average across all time periods.

Source: Author's computations.

the coefficient on $\ln(v/u)$ is very precisely estimated and falls in the same narrow range as the single equation estimates. The unobserved NAIRU variable enters the Beveridge curve equation with a very precisely estimated coefficient.

The help-wanted index is used to estimate the vacancy series during the earlier period and the JOLTS survey during the latter period. Can the same model fit both periods? Originally the model was estimated with a dummy variable for the latter period in the Beveridge curve and interactions between that dummy and the NAIRU and between the dummy and $\ln(v/u)$. Neither of the interactions was statistically significant individually or jointly in any specification tried, so these were dropped from the model with virtually no impact on any other model parameters. Only the dummy variable supplementing the intercept was retained.

The most interesting result in column 1 is the relative magnitude of the Kalman gain for the Phillips curve and Beveridge curve residuals. The Kalman filter model constructs the NAIRU as a weighted average of the residuals of the two equations plus the previous period's estimate of the NAIRU. The Kalman gain is the weight put on each of the two residuals. Given the estimated parameters, the residuals of the Phillips curve play virtually no role in constructing the NAIRU, while those of the Beveridge curve play a major role. The estimated variance of the Beveridge curve innovation (μ_t) is so close to zero⁷ that the model identifies the NAIRU as being nearly exactly proportional to the difference between the left-hand side of equation (8) and the constant plus $b \ln(v/u)$. This result is not unique to the model in the first column. In every specification presented—in fact, in every one of the several dozen specifications tried—the model chose to equate the NAIRU with the Beveridge curve residual nearly exactly. Thus the right three columns of the table present three specifications with the constraint that the Beveridge curve residual is exactly proportional to the NAIRU. The first of these replicates the specification in the first column. Comparing the results in the first and third columns shows what little effect the constraint has on the estimated coefficients.

That the Phillips curve residuals are seen as uninformative with respect to the magnitude of the model NAIRU might indicate that what is being measured is not a NAIRU, but simply the time variation in the inter-

cept of the Beveridge curve. On the other hand, it could be that the time variation in the intercept of the Beveridge curve very precisely measures movements in the NAIRU so that information from Phillips curve residuals is superfluous. We saw in table 4.1 and figure 4.3 just how well the Beveridge curve fits. There is no straightforward test of whether the model's NAIRU matters for explaining inflation, since it must be assumed that the coefficient on the unemployment rate is the same as the coefficient on the NAIRU for identification. However, once the variance of the Beveridge curve innovation is restricted to zero, so that the Beveridge curve residual is assumed to be proportional to the NAIRU, a constant and a separate coefficient for the NAIRU can be added to the Phillips curve. The bottom three lines of table 4.2 show the results of doing this for the six specifications presented there. If the Kalman filter model is correct then the constant term in the Phillips curve should equal zero, and the coefficient on the NAIRU should equal 1. In none of the six specifications can either hypothesis be rejected individually or jointly. Furthermore, in most of the specifications presented the hypothesis that the coefficient on the NAIRU is equal to zero can be rejected at the .10 level or better in a one-tailed test.

When models with every possible combination of inflation measure (CPI, CPI-core, GDP deflator, GDP consumption deflator), inflation lag structure, and unemployment lag structure were estimated, the results were remarkably consistent. There were 159 specifications that could not be rejected as being overly constrained compared to the specification in column one of table 4.2. Of those, there was not one in which the hypotheses that the constant was equal to zero or the coefficient on the NAIRU is equal to one could be rejected individually or jointly at the .10 level (two-tailed test). Yet in 92 of the 159 specifications, the hypothesis that the coefficient on the NAIRU is zero was rejected at the .05 level and at the .10 level in 150 specifications (one-tailed test).

The results aren't quite as good when the model is estimated with age-adjusted unemployment. There are 169 specifications that cannot be rejected when compared to one with 12 lags of inflation, contemporaneous unemployment, and 4 lags of unemployment. Once again, there is not a single specification where the hypothesis that the constant in the Phillips curve equals zero or the hypothesis that the coefficient on the

NAIRU equals 1 can be rejected at the .10 level individually or jointly. However, there are only 43 specifications where the hypothesis that the coefficient on the NAIRU is zero can be rejected at the .10 level and only 9 where it can be rejected at the .05 level. Perhaps once the variation in the NAIRU due to the age structure of the population is taken into account by adjusting the unemployment rate there is little additional information in the Beveridge curve residuals. On the other hand, the Beveridge curve model might be considered misspecified when age-adjusted unemployment is used, since the vacancy rate hasn't been similarly adjusted.

It appears that the low frequency movements in the Beveridge curve probably belong in the Phillips curve as an indicator of variation in the NAIRU—at least if unadjusted unemployment is used. The next question is whether there is any more variation in the NAIRU than the variation due to labor market efficiency reflected in the Beveridge curve residuals. With the constraint that there is no innovation in the Beveridge curve, it is possible to estimate a model with two unobservables so that the NAIRU is the sum of the Beveridge curve residual and a filtered version of the Phillips curve residual. When this model is estimated in any of a wide range of specifications the likelihood is maximized when the variance of the innovation in the Phillips curve NAIRU is zero. Thus the hypothesis that there is nothing more to variation in the NAIRU than that captured by the Beveridge curve residual cannot be rejected.

How does the standard time-varying NAIRU estimated using only the Phillips curve compare to the Beveridge curve-based NAIRU? Figure 4.4 presents examples of both. The gray line in figure 4.4 depicts the NAIRU derived from the Beveridge curve model from column 1 of table 4.2. The heavy black line in figure 4.4 depicts a NAIRU estimated using only the residuals of a Phillips curve with the same specification as that used to estimate the combined Beveridge curve-Phillips curve model. The two are similar in many respects. Except for a short period in the late 1960s to early 1970s, and the lack of a bulge in the Phillips curve NAIRU in the mid-1980s, the two track each other fairly closely. When the unemployment rate was high relative to the job vacancy rate it was also high relative to the inflation rate.

If the Beveridge curve and the Phillips curve NAIRU look similar, what is the advantage of the latter? Confidence intervals for both series




Figure 4.4
Phillips Curve and Beveridge Curve NAIURU
Source: Author's computations.

were constructed for the two models in figure 4.4 that take account of both forecast and parametric uncertainty by computing 10,000 Monte Carlo trials. Despite the Beveridge curve model having several additional parameters, the 90-percent confidence intervals for the Phillips curve NAIURU were about 40 percent larger on average.⁸ Other specifications for the two models yielded similar results—the Beveridge curve-based NAIURU had narrower confidence intervals in every specification tried.

The very precise and similar estimates of the Beveridge curve series across many different specifications, along with the narrow confidence intervals on the computed NAIURU series, suggests that there should be considerably more certainty about the position of the Beveridge curve NAIURU than there is about NAIURUs estimated from the Phillips curve alone. This is somewhat true. Figure 4.5 presents the average value of the Beveridge curve NAIURU estimated across the 159 specifications using total unemployment. The specifications varied the lags of unemployment and inflation and the inflation measure as described above. Also plotted in figure 4.5 are the minimum and maximum values in each quar-




Figure 4.5
 NAIURU with Upper and Lower Bounds
Source: Author's computations.

ter of the 90-percent confidence intervals for the NAIURU computed for each of the 159 specifications. The bounds were estimated by simulating parametric and forecast uncertainty with 10,000 Monte Carlo trials each.

Allowing for forecast, specification, and parametric uncertainty, as the bounds in figure 4.5 do, considerable uncertainty about the position of the NAIURU at any given time remains. This is particularly true at the moment because the switch from the help-wanted series to the JOLTS series adds substantially to uncertainty. Still, the results reported here improve on estimates based on only the Phillips curve in at least one dimension—there is little uncertainty about the timing of major changes in the NAIURU. All estimates show a substantial rise in the NAIURU during the 1970s and a decline in the late 1980s and early 1990s. While NAIURU values much above 6 percent can be ruled out during the 1960s and the mid-to-late 1990s, values less than that can be ruled out for the decade starting in 1978. This provides more guidance to policymakers than past estimates.

5. Conclusion

This paper has presented a new method for estimating time variation in the NAIRU using the vacancy-unemployment relationship. A simple theory of this relationship based on a matching model suggests equations that do an uncannily good job of fitting transformed vacancy and unemployment data. When the Beveridge curve model is estimated simultaneously with a Phillips curve, the parameter estimates for both equations are reasonable and the parameters of the Beveridge curve are estimated with particular accuracy. The estimates suggest that the NAIRU is nearly exactly proportional to the residual in the Beveridge curve. When this constraint is imposed it is possible to test whether the Beveridge curve residuals help explain inflation. In the 328 specifications tried, there were none in which the hypothesis that the scaled Beveridge curve residual was the NAIRU could be rejected. The hypothesis that the Beveridge curve residual did not help explain inflation could be rejected at least at the .10 level in nearly all specifications using the total unemployment rate and many where age-adjusted unemployment was used. A model augmented with a time-varying NAIRU estimated as the sum of the scaled Beveridge curve residual plus a filtered version of the Phillips curve error was estimated. The hypothesis that the filtered Phillips curve error did not help forecast inflation could not be rejected and the resulting NAIRU series differed little from the Beveridge curve residual alone. A standard Phillips curve NAIRU resembles the Beveridge curve NAIRU in the timing of its movements, which validates the theory on which both the Beveridge curve and the Phillips curve NAIRU are based. However, estimates of the Beveridge curve NAIRU are more precise.

Despite the very precise estimates of the parameters of the Beveridge curve, when forecast, specification, and parametric uncertainty are taken into account, the data are consistent with a fairly wide range of values for the NAIRU at each point in time. This is particularly true since 2001, when the JOLTS vacancy rate series replaces the help-wanted series. Still, there appears to be considerable information in the Beveridge curve model about movements in the NAIRU. Estimated NAIRU series differ in the magnitude of the fluctuations, but hardly at all in their average value or the timing of the fluctuations. Further, as we get more experience

with the JOLTS vacancy series, uncertainty about the current NAIRU will decline since the increased uncertainty post-2000 is due entirely to uncertainty about the magnitude of the post-2000 dummy variable in the Beveridge curve.

■ *I would like to thank Vania Stavrakeva for her excellent research assistance and Robert Valletta for the use of his age-adjusted unemployment series. I would also like to thank the Brookings Institution and the Russell Sage Foundation for research support. This paper further develops work I did with Jessica Cohen and Adam Posen on Beveridge curve shifts (Cohen, Dickens, and Posen 2001). A conversation with Jim Stock was helpful in thinking through identification issues, while Richard Startz pointed out a problem with the model and made several helpful comments about presentation. I would also like to thank Olivier Blanchard for comments on an earlier draft of the paper and for pointing me toward two conference papers he wrote in the late 1980s that took a very similar approach to the one taken here. This paper differs from those in the way the Beveridge curve is estimated, the joint estimation of the Phillips curve and Beveridge curve, and its focus on estimating a time-varying NAIRU.*

Notes

1. In these papers only the constant term, or the NAIRU, was allowed to vary. When Brainard and Perry (2000) estimated Phillips curves allowing all parameters to vary they found that the constant and the coefficient on unemployment were relatively stable (and thus so was the NAIRU). Instead they explained the different behavior of inflation over the decades by variation in the sum of the coefficients on lagged inflation.
2. See, for example, the results in Staiger, Stock, and Watson (1997) or compare figures 3 and 4 in Gordon (1997).
3. Where the aggregate price level is the geometric average of all prices in the economy.
4. In a departure from Blanchard and Diamond (1989), quits and layoffs are both assumed to arise from the breakdown of a match, which is signified by the arrival of a new price for the entrepreneur.
5. The figure assumes a log-linear matching function with constant returns to scale and equal weight on vacancies and unemployment.

6. Using either the level, the filtered value, or first differences corresponding to the treatment of the variables in the specification being estimated.

7. In fact, in all of the dozens of specifications tried the variance of the Beveridge curve residual was estimated to be slightly negative. This is possible because the variance is scaled and then added to the forecast error variance due to the innovation in the NAIRU so that the total forecast error variance remains positive. The slight negative value was often statistically significantly different from zero, a result that suggests that the assumptions about the innovations are not exactly correct. In fact, the small positive auto-correlation (about .25 in specifications where it was inspected) in the estimated changes in the NAIRU could explain this.

8. In general the likelihood function for the Kalman filter model for the Philips curve NAIRU pushed the variance of the innovation of the NAIRU to zero so that value was fixed in order to compute the NAIRU shown in figure 4.4. The variance of the innovation was chosen so that the Beveridge curve and Phillips curve NAIRUs had the same variance.

References

- Abraham, Katharine G. 1983. "Structural/Frictional vs. Deficient Demand Unemployment: Some New Evidence." *American Economic Review* 73(4): 708–724.
- Abraham, Katharine G. 1987. "Help-Wanted Advertising, Job Vacancies, and Unemployment." *Brookings Papers on Economic Activity* 1: 207–248.
- Answers.com. 2008. "Hoover's Profile: Monster Worldwide, Inc." Available at <http://www.answers.com/topic/monster-worldwide-inc?cat=biz-fin>. Accessed June 12.
- Ball, Lawrence, and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives* 16(4): 115–136.
- Basistha, Arabinda, and Richard Startz. 2008. "Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach." *Review of Economics and Statistics* 90(4): 805–811.
- Beveridge, William H. 1945. *Full Employment in a Free Society*. New York: W.W. Norton & Company.
- Blanchard, Olivier Jean, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1: 1–60.
- Bleakley, Hoyt, and Jeffrey C. Fuhrer. 1997. "Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics." *New England Economic Review* (September/October): 3–19. Available at <http://www.bos.frb.org/economic/neer/neer1997/neer597a.pdf>.
- Brainard, William C., and George L. Perry. 2000. "Making Policy in a Changing World" Perry, George and James Tobin (eds), In *Economic Events, Ideas, and Policies: The 1960s and After*, ed. George L. Perry and James Tobin, 43–68. Washington, DC: Brookings Institution Press.

Cohen, Jessica, William T. Dickens, and Adam Posen. 2001. "Have the New Human Resource Management Practices Lowered the Sustainable Unemployment Rate?" In *The Roaring Nineties: Can Full Employment be Sustained?*, ed. Alan B. Krueger and Robert M. Solow, 219–259. New York: The Russell Sage Foundation and the Century Foundation Press.

Gordon, Robert J. 1990. "U.S. Inflation, Labor's Share, and the Natural Rate of Unemployment." In *Economics of Wage Determination*, ed. Heinz König, 1–34. New York: Springer Verlag.

Gordon, Robert J. 1997. "The Time-Varying NAIRU and its Implications for Economic Policy." *Journal of Economic Perspectives* 11(1): 11–32.

Gordon, Robert J. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 2: 297–346.

Hansen, Bent. 1970. "Excess Demand, Unemployment, Vacancies, and Wages." *Quarterly Journal of Economics* 84(1): 1–23.

Hernandez, Judy. 2008. "The Monster.com Story." Available at <http://www.careertrainer.com/Request.jsp?lView=ViewArticle&Article=OID%3A55989>. Accessed June 12.

Konstant, Raymond A., and Irvin F. Wingard. 1968. "Analysis and Use of Job Vacancy Statistics: Part I." *Monthly Labor Review* 91(August): 22–31.

Medoff, James L. 1983. "U.S. Labor Markets: Imbalance, Wage Growth and Productivity in the 1970s." *Brookings Papers on Economic Activity* 1: 87–128.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity* 1: 411–441.

Shimer, Robert. 1999. "Why is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual 1998*, ed. Ben S. Bernanke and Julio J. Rotemberg, 11–61. Cambridge, MA: The MIT Press.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "How Precise are Estimates of the Natural Rate of Unemployment?" In *Reducing Inflation: Motivation and Strategy*, ed. Christina D. Romer and David H. Romer, 195–246. Chicago: University of Chicago Press.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 2001. "Prices, Wages, and the U.S. NAIRU in the 1990s." In *The Roaring Nineties: Can Full Employment be Sustained?*, ed. Alan B. Krueger and Robert M. Solow, 3–60. New York: The Russell Sage Foundation and the Century Foundation Press.

Stock, James H. 2001. "Monetary Policy in a Changing Economy: Indicators, Rules, and the Shift Towards Intangible Output." Paper presented at Bank of Japan Conference on Monetary Policy.

Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44(2): 293–335.

Valletta, Robert G. 2006. "Why Has the U.S. Beveridge Curve Shifted Back? New Evidence Using Regional Data." Federal Reserve Bank of San Francisco Working Paper Series No. 2005-25. Available at <http://www.frbsf.org/publications/economics/papers/2005/wp05-25bk.pdf>.

Zagorsky, Jay L. 1998. "Job Vacancies in the United States: 1923 to 1994." *Review of Economics and Statistics* 80(2): 338-345.

Comments on “A New Method for Estimating Time Variation in the NAIRU” by William T. Dickens

Olivier Blanchard

Dickens’s paper offers a promising strategy to identify shifts in the natural rate of unemployment by looking jointly at the Beveridge curve and the Phillips curve. In my comments, I want to do two things. First, sketch the conceptual framework that allows one to extract information about the natural rate of unemployment from the Beveridge curve and the Phillips curve, and the factors behind its movement. I believe this is the framework that underpins Dickens’s analysis; all I want to do is to make it more transparent. Second, bring the framework to the U.S. data. The answers one gets from this exercise are surprisingly clear-cut and, I think, interesting, as they do not conform to my (and I suspect your) priors. This is, of course, just a first look at the data, but it shows how useful this approach can potentially be for thinking about changes in the natural rate of unemployment.

The theoretical framework can be characterized simply by two equations. The first equation is a relation between the flow into unemployment and the flow out of unemployment and back to employment:

$$(1) \quad s(1-u) = mf(u, v).$$

Equilibrium in the labor market is characterized by the equality of separations and hires. The left side of the equation captures separations, where s is the separation rate and u is unemployment normalized by the labor force. The right side of equation (1) captures hires, and is given by a matching function, whose output is a flow of new hires. The function matches unemployed workers with vacancies, v . The function is increasing in u and v , while m is a scale parameter denoting the efficiency of the matching process. Equation (1) describes a negative relation between

unemployment and job vacancies—the Beveridge curve. The position along the Beveridge curve is related to the state of the business cycle. Recessions are periods when many unemployed workers are pursuing few vacancies. Conversely, in a tight labor market, which is generally associated with high labor demand, more vacancies are searching among fewer unemployed workers.

For our purposes, we are interested not so much in movements along the Beveridge curve, but in the factors that lead to *shifts* in the curve. Given the way equation (1) is written, shifts in the curve arise from two sources. The first is a change in s . A decrease in the separation rate means lower flows through unemployment. This could result from less labor market reallocation and/or more relative flows directly from employment (or entrance into the labor force) to employment. The latter case could be the result, for example, of temporary employment agencies now making it possible for workers to transition from one job to another without becoming unemployed in the process. A lower s , for a given amount of vacancies, implies less unemployment, and thus a shift of the Beveridge curve inward. The other source of shifts in the curve is a change in m , the efficiency of matching. It is reasonable to think that in recent years the technology for matching workers and jobs has become better, so that not as many workers or as many vacancies are needed to generate a certain flow of hires. The advent of Monster.com is an obvious example. Improvements in matching efficiency shift the Beveridge curve inward.

The second equation in the framework determines wages. The relationship says that the wage firms can offer must be equal to the wage implied by bargaining:

$$(2) \quad \bar{w} = w \left(\frac{v}{u}, z \right).$$

The left side of the equation denotes the wage that is consistent with normal profits. The right side of the equation is a wage function that can be derived in the context of a wage bargaining model where the surplus from matching a firm and a worker is shared in some proportion (see, for example, Pissarides 1985). According to this function, the labor market variable that matters in determining the outcome of wage bargaining is the ratio of vacancies to unemployment, v/u . This ratio determines the

bargaining power that each party possesses. A high v/u indicates that workers' bargaining strength is high relative to firms', and this yields a higher wage rate. The wage-determination function is thus increasing in v/u . For given \bar{w} and z , equation (2) describes a positive linear relationship between v and u . The parameter z summarizes the factors that affect bargaining, which may arise from the presence of unions or other institutional features pertaining to the labor market that influence the bargaining power of workers relative to firms.

Equation (2) abstracts from the presence of nominal rigidities. Wage dynamics, however, can be introduced in (2) to yield a Phillips curve specification that relates wage inflation, price inflation, v/u , and z . Movements in z will shift the Phillips curve relationship. Increasing globalization, to the extent that it reduces the bargaining power of workers, will shift the Phillips curve inward in the (u, w) space.

Using equations (1) and (2), it is possible to derive the steady-state equilibrium levels of u and v . In particular, the natural rate of unemployment can be written as:

$$(3) \quad u^* = u^*(s, m, z).$$

The natural rate of unemployment is a function of the separation rate, the efficiency of the matching process, and the factors that affect the bargaining power of workers. From equation (3), it is evident that we cannot obtain an estimate of the natural rate of unemployment by estimating either the Beveridge curve or the Phillips curve alone. Both the Beveridge curve and the Phillips curve are needed to back out s , m , and z , the three factors that affect u^* . This simple framework thus illustrates why the strategy pursued in Dickens's paper to estimate the natural rate of unemployment is potentially a good one: estimation of both the Beveridge curve and the Phillips curve can capture the three factors influencing the natural rate of unemployment.

Having laid the framework for thinking about movements in the natural rate of unemployment in the context of the Beveridge curve and the Phillips curve, I will now bring this framework to the data. The first observation that I would like to make, which I think is not controversial, is that low frequency movements in the unemployment rate strongly suggest a decline in the natural rate of unemployment since the late 1980s,

by a magnitude of at least 1 percentage point. The question is then what accounts for the decline in the natural rate of unemployment. Is it movements in s , m , or z ? As documented in Dickens's paper, the Beveridge curve has shifted noticeably inward after a transition period in the late 1980s. This inward shift in the Beveridge curve points to s and m , and not to z —the Phillips curve shifter—as potentially important factors in lowering the natural rate of unemployment. Is it possible to assess the independent contribution of s and m to the decline in u^* ? The answer to this question is relatively straightforward. We can observe the separation rate s directly in the data. An estimate of m can be obtained from estimating a matching function, as in Blanchard and Diamond (1989, 1991). Performing this exercise, it becomes apparent that a lot of the action is coming from the separation rate, which has declined noticeably since the late 1980s. This is shown in figure 4.6, which plots a time series for the separation rate from employment and for the unemployment rate, using Bureau of Labor Statistics data. The separation rate includes separations from employment to unemployment, from employment to moving out of the labor force, and from employment to employment.




Figure 4.6

The Unemployment Rate and the Separation Rate from Employment, 1970–2004

Source: Author's computations from Bureau of Labor Statistics data.

Note: Shaded areas show NBER-dated recessions.

The sharp decline in the separation rate suggests that s played a prominent role in the decline of the natural rate of unemployment. Note that the picture shows that s has declined from a monthly 3 percent rate to about 2.5 percent—a 15 percent decline. This is roughly the same decline that seems to have occurred for the natural rate of unemployment, which is usually estimated to have declined from 6 percent to about 5 percent, approximately a 15 percent decline. In sum, to a close approximation, it appears that the decline in the natural rate of unemployment can be explained entirely by a decline in the separation rate. Changes in matching efficiency and in the bargaining power of workers—with this last factor being my prior as the most likely explanation for a decline in u^* —do not appear to account for an important part of the story.

The next step, which goes beyond the scope of Dickens's paper but is nonetheless an important issue, is to understand what lies behind the decline in the separation rate. Lower worker flows could be the result of demographics. An older labor force is less inclined to move among jobs. This would result in lower worker flows even in an environment in which job flows are approximately given. But it is possible to look directly at job flows, and here we see that the decline in worker flows is not just due to demographics, as figure 4.7 shows. There has been a decline in job creation and in job destruction; that is, the amount of job reallocation in the U.S. economy has fallen. This decline does not explain entirely the decline in worker flows, but roughly two-thirds of it. Still, the reasons why churning in the labor market has decreased remain to be investigated. More wage flexibility, better inventory control, and more integrated chains of production could have contributed to the decline in job creation and job destruction.

In sum, Dickens's paper outlines a promising approach, and I hope that my comments have highlighted the usefulness of this approach to both estimating the natural rate of unemployment and discriminating among different reasons for movements in the natural rate of unemployment. The implementation of this approach to U.S. data delivers clear leads, in that it downplays decreases in the bargaining power of workers and more efficient firm-worker matching as explanations for a lower natural rate of unemployment. Somewhat surprisingly, the data point to lower worker flows, and in turn to lower job flows as the most important factor affect-




Figure 4.7

Quarterly Job Flows in the Private Sector, 1990–2005

Source: Faberman (2006); tabulated from BLS Business Employment

Note: Dynamics (BED) micro data. Shaded areas show NBER-dated recessions.

ing the decline in u^* . What lies behind the decrease in job flows remains an open issue that we need to address. But at least we know what to look for, and how to interpret it.

References

- Blanchard, Olivier Jean, and Peter A. Diamond. 1989. “The Beveridge Curve.” *Brookings Papers on Economic Activity* 1: 1–60.
- Blanchard, Olivier Jean, and Peter A. Diamond. 1991. “The Aggregate Matching Function.” Working Paper No. 3175. Cambridge, MA: National Bureau of Economic Research.
- Faberman, R. Jason. 2006. “Jobs Flows and the Recent Business Cycle: Not All ‘Recoveries’ Are Created Equal.” BLS Working Paper No. 391. Washington, DC: U.S. Bureau of Labor Statistics. Available at <http://www.bls.gov/ore/pdf/ec060030.pdf>.
- Pissarides, Christopher A. 1985. “Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages.” *American Economic Review* 75(4): 676–690.

Comments on “A New Method for Estimating Time Variation in the NAIRU” by William T. Dickens

Christopher A. Pissarides

It is fitting that there should be a paper on the Beveridge curve at this conference. Like the Phillips article, the original paper on the Beveridge curve was also published in 1958, by Dow and Dicks-Mireaux in *Oxford Economic Papers*. The article was mainly about measurement and made the case for a good correlation between vacancies and unemployment, using British historical data. It offered no theory. It launched a literature, known as “UV-analysis,” on the measurement of vacancies and unemployment and on their relation to excess demand in the labor market. It was concerned with the problem of finding how much unemployment can be reduced with Keynesian demand management policy, given the frictions in the labor market, and in this sense it was a precursor of the later critiques of the Phillips curve. Lipsey (1965) brought the Phillips curve and UV-analysis together, in a paper that addressed many of the issues examined in Bill Dickens’s paper.

Commenting on this paper has become, in Dickens’s words, an attempt to hit a “moving target.” In order to avoid writing a comment that may turn out to be irrelevant I have therefore decided to comment less directly on what Bill says, and focus instead on the problem that he has posed and discuss some thoughts on how to go about modeling it.

Dickens suggests using information derived from the Beveridge curve to calculate changes in the NAIRU. I totally agree with this objective—ever since its inception, the Beveridge curve has been used to classify reasons for changes in unemployment. These exercises were a precursor to his task. Dickens’s question can be rephrased to the question, did unemployment between t and $t + 1$ change because of a change in the NAIRU,

or because of nominal shocks? Or, more generally, did unemployment change because of a real shock or because of a nominal shock?

Distinguishing between changes in the NAIRU and other changes in unemployment requires two equations. One is the Beveridge curve, which is an equilibrium equation that summarizes the speed of structural change and the frictions in the labor market. The other equation is essentially an equation for the demand for labor. In my view, the best way to think about the Beveridge curve is in terms of the flows in and out of unemployment. By definition, the change in unemployment between period t and $t + 1$ is

$$(1) \quad u_{t+1} - u_t = \text{inflows in } t - \text{outflows in } t,$$

with the stocks measured at the beginning of the period. For the flow terms we can write

$$(2) \quad \text{inflows} = \text{new entry} + \text{job separations, and}$$

$$(3) \quad \text{outflows} = \text{exits} + \text{job acceptances.}$$

The Beveridge curve is defined as the combination of unemployment and vacancies that equates the inflows with the outflows. Writing a theory of the Beveridge curve amounts to modeling each one of the four terms in (2) and (3), and tracing the combinations of vacancies and unemployment that maintain the equality between the inflows and the outflows in the absence of shocks.

Perhaps surprisingly at first, but on reflection not so surprisingly, we get a good approximation to the dynamics of unemployment if we treat unemployment as if it were always on the Beveridge curve (Pissarides 1986, Shimer 2007). It might be surprising at first because with the change in unemployment given by the difference between inflows and outflows, and the Beveridge curve defined as the locus of equality between inflows and outflows, how does unemployment change if we are always on the Beveridge curve? The best way to think about this conundrum is in terms of speeds of adjustment and the length of the period. Treat unemployment as the only unknown in the inflows = outflows condition and assume the period is a quarter. If one of the four terms in (2) and (3) changes because of a shock, unemployment changes fast to restore equality between the new inflows and outflows. In other

words, although the labor market is characterized by frictions, given the size of the shocks that we normally observe, the frictions are sufficiently small that unemployment jumps between one flow equilibrium and the next within a quarter.¹

Consider now the shocks that might make unemployment change in the context of the Beveridge curve. The search and matching theory makes the job acceptance flow the key to the entire framework (Pissarides 2000, chapter 1). In its simplest form it assumes constant job separation rates $s(1 - u_t)$, either zero or constant entry and exit rates, and that the rate of job acceptance is given by the aggregate matching function, $m(u_t, v_t)$. The matching function gives the number of new jobs formed as a function of the workers available to take new jobs, and the number of vacant jobs, v_t . Let f_t denote the average rate of job finding, defined by $f_t = m(1, v_t/u_t)$, and assume that entry and exit are zero. The Beveridge curve is

$$(4) \quad u_t = \frac{s}{s + f(v_t / u_t)}.$$

If nominal shocks have any influence on unemployment in this framework, the channel through which they have it is the vacancy rate, v_t . The vacancy rate is given by the second equation of the system, the demand for labor. If, for example, a positive nominal shock that raises inflation increases the demand for labor because of nominal stickiness somewhere in the system, the vacancy rate increases above trend and unemployment falls. The implied negative relation between unemployment and inflation is the essence of the Phillips curve, and the channel that links the change in the demand for labor with unemployment is the vacancy rate and the matching function.

In terms of the Beveridge diagram derived from (4), the fall in unemployment induced by the nominal shock is represented by a movement along the Beveridge curve. If one were to accept the simple framework underlying equation (4) as a complete characterization of the dynamics of unemployment, the vacancy rate is the only channel through which nominal shocks can be transmitted to unemployment. Any other changes in unemployment, for given vacancies, are changes in the NAIRU. These changes are associated with changes in the rate of labor turnover, s , changes in the matching efficiency of the labor market, represented by

shifts in $f(\cdot)$ for given v_t/u_t , and with changes in the rate of entry into and exit from the labor force. For example, demographics shift the NAIRU, potentially by changing all terms in (4), the rate of labor turnover, the matching efficiency of the labor market, and the rate of entry and exit from the labor force. Unemployment insurance shifts the NAIRU by changing the intensity of search, the efficiency of matching, and so on.

In my view, the best way to uncover changes in the NAIRU associated with shifts in the matching efficiency of the labor market is not to estimate the entire Beveridge relation, as Bill has attempted to do, but to estimate the matching function directly (or the job-finding rate). When I did this for Britain in 1986 I found that most of the changes in unemployment were associated with changes in the NAIRU, although changes in the vacancy rate also played a role. This was to be expected, given that when unemployment was trending up between the late 1960s and the early 1980s the vacancy rate was fluctuating around a flat trend. Several estimates of matching functions by other authors can be used to decompose changes in unemployment between changes due to the vacancy rate and changes due to other factors.² The U.S. experience since 2001, when reliable vacancy data became available through JOLTS, is probably unique in that it attributes virtually all changes in unemployment, save for a small error term, to changes in the vacancy rate, a property that has been emphasized in some of Shimer's recent influential work (for instance, see Shimer 2005 and Elbrahimi and Shimer 2008).

Dickens finds something similar in his estimated Beveridge curves. However, this finding does not necessarily imply a constant NAIRU, even in the simple framework of equation (4). There might be causes of changes in the vacancy rate, which keep the Beveridge curve fixed, and which are real and associated with changes in the NAIRU. For example, consider material shocks. If the price of raw materials goes up and real wages are subject to inertia, vacancies might fall dramatically. Unemployment rises through a movement down the Beveridge curve. The Beveridge curve has no obvious reason to shift in this case.

This is why we need to estimate a second equation, preferably simultaneously with the matching function, before we can confidently calculate the NAIRU. The second equation is a demand for labor equation and is derived from a conventional model of the firm with costs of adjustment due to frictions. The difference between investment-type quadratic adjust-

ment costs and matching frictions is that the costs of adjustment with frictions depend on the tightness of the labor market. At high vacancy-to-unemployment ratios these costs and frictions are higher, because there is more competition between firms for the pool of unemployed workers. The implication of this property is that we can write the dynamic demand for labor equation as a vacancy supply equation and estimate it in terms of all the conventional labor demand regressors, including price misperceptions (Pissarides 1986; Yashiv 2000).

Dickens has a second equation in his model but it is not a labor demand equation. His equation is similar to the one that featured in the very first models of the Beveridge curve (Dow and Dicks-Mireaux 1958). It is essentially the 45-degree line through the origin, which defines the locus of equality points between u and v as the equilibrium points. Modern approaches to the Beveridge curve derive the second equation from optimizing models of the firm and show that the slope of the second curve is a function of the model's parameters.

A more important point about the second equation, however, is this: are we justified in focusing on the vacancy rate as the only variable that can transmit nominal shocks to unemployment? In the context of Phillips curve analysis we are asking whether all shocks to the unemployment rate other than those acting through the vacancy rate are shocks to the NAIRU. In the context of Beveridge curve analysis the question is whether the simple framework in (4) is sufficient.

There has been a lot of work on this issue recently, with reference mainly to business cycle fluctuations in unemployment. These high frequency fluctuations are also the ones that Bill studies in his paper. The upshot of the discussion is that business cycle fluctuations in unemployment are driven both by fluctuations in the inflow rate and the outflow rate (see Shimer 2007, Fujita and Ramey 2007, and Petrongolo and Pissarides 2008). Moreover, for cyclical fluctuations one can ignore the movement in and out of the labor force and focus on movements between employment and unemployment. In that context, the consensus is that about two-thirds of fluctuations are due to the outflow rate, for which the matching function approach serves us well, and another third to the inflow rate. The inflow rate in (4) is the parameter s . The recent empirical literature on the ins and outs of unemployment says that s should not be a parameter but a cyclical variable. A complete model of the NAIRU

derived from the Beveridge curve should account for the endogeneity of job separations.

Although good theoretical models of the endogeneity of job separations exist, it is much more difficult to find good empirical or quantitative models of separations.³ I think this is likely to be the main sticking point in the task that Dickens set himself. Because job separations vary and are negatively correlated with job accessions, it is plausible to assume that these are driven by the optimizing decisions of firms and workers in response to shocks. Some of those shocks are nominal, and if there are nominal rigidities of the kind analyzed in Phillips curve models, some changes in the parameter s in (4) are changes associated with nominal shocks, namely, not changes in the NAIRU. But changes in s shift the Beveridge curve. It follows that in a general model of the NAIRU there are changes in unemployment that are not caused by changes in vacancies, and which are not changes in the NAIRU. Therefore, identifying all changes in unemployment that take place for a given vacancy rate as changes in the NAIRU would be a mistake.

A challenge that is facing both search and matching theory and modern Phillips curve analysis is how to explain the fact that on average about one-third of fluctuations in unemployment are due to shocks to job separations (or, at least, to the unemployment inflow rate) and yet for long stretches of time the vacancy-unemployment scatter of points is tightly distributed around a fixed Beveridge curve. As far as I know there is no paper in the literature yet that does that, and so there is no model that can convincingly be used to provide a framework for the estimation of the NAIRU from Beveridge curve analysis. But following the approach that I outlined in this comment, under the assumption that all nonrandom shifts in the Beveridge curve are changes in the NAIRU, is a good first approximation to the data.⁴

Notes

1. In my examination of British and other European data, the only time that the assumption of flow equality in quarterly data was not a good working assumption was the two-year period of the large “Thatcher shock,” 1979–1981. See Pissarides (1986) and Petrongolo and Pissarides (2008). Shimer (2007) does not report any period when this assumption was badly violated for the United States.

2. See Blanchard and Diamond (1989) for early U.S. estimates and Petrongolo and Pissarides (2001) for a survey of several estimates.
3. For the theory see Mortensen and Pissarides (1994) and Caballero and Hammour (1996). For more discussion of the empirics of job separations see Davis and Haltiwanger (1999).
4. See Ebrahimi and Shimer (2008) for a promising attempt at explaining simultaneously the tightly distributed points in Beveridge space and the variance in the separation rate. They focus on the post-2001 data, when there are no shifts in the Beveridge curve. The problem of reconciling periodic shifts with long periods of tightly distributed $u - v$ points remains.

References

- Blanchard, Olivier Jean, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1: 1–60.
- Caballero, Ricardo J., and Mohamad L. Hammour. 1996. "On the Timing and Efficiency of Creative Destruction." *Quarterly Journal of Economics* 111(3): 805–852.
- Davis, Steven J., and John C. Haltiwanger. 1999. "Gross Job Flows." In *Handbook of Labor Economics*, Vol. 3B, ed. Orley Ashenfelter and David E. Card, 2711–2805. Amsterdam: Elsevier.
- Dow, J. C. R. and L. A. Dicks-Mireaux. 1958. "The Excess Demand for Labour: A Study of Conditions in Great Britain, 1946–56." *Oxford Economic Papers*, New Series 10(1): 1–33.
- Ebrahimi, Eshan, and R. Shimer. 2008. "Stock-Flow Matching." Working paper. Available at <http://robert.shimer.googlepages.com/stock-flow-print.pdf>.
- Fujita, Shigeru, and Garey Ramey. 2008. "The Cyclicalities of Separation and Job Finding Rates." Forthcoming in *The International Economic Review*. Available at http://weber.ucsd.edu/~gramey/Pubs/FR_IER_1-10-08.pdf.
- Lipsey, Richard G. 1965. "Structural and Deficient-Demand Unemployment Reconsidered." In *Employment Policy and the Labor Market*, ed. Arthur M. Ross, 210–255. Berkeley: University of California Press.
- Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61(3): 397–415.
- Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39(2): 390–431.
- Petrongolo, Barbara, and Christopher A. Pissarides. 2008. "The Ins and Outs of European Unemployment." *American Economic Review* 98(2): 256–262.
- Pissarides, Christopher A. 1986. "Unemployment and Vacancies in Britain." *Economic Policy* 1(3): 500–541.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*, second edition. Cambridge, MA: The MIT Press.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95(1): 25–49.

Shimer, Robert. 2007. "Reassessing the Ins and Outs of Unemployment." Working Paper No. 13421. Cambridge, MA: National Bureau of Economic Research.

Yashiv, Eran. 2000. "The Determinants of Equilibrium Unemployment." *American Economic Review* 90(5): 1297–1322.

5

Inflation Expectations

Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy

Christopher A. Sims

1. A View of the History of the Phillips Curve

The original observation by Phillips (1958) simply noted an empirical regularity: unemployment and wage inflation tended to be inversely related. This observation came at a time when Keynesian macroeconomic theory had a very simple and incomplete model of inflation. Keynesian theory treated wages as, if not fixed, then on an exogenously given time path. It was a theory of how nominal aggregate spending determined the level of output and employment, so long as supply-side limits on output and employment were not encountered. It was recognized that when aggregate demand exceeded supply-side limits, the result would be inflation, but the standard Keynesian theory had a discontinuity at the point where output hit “capacity,” and it had no quantitative predictions about the determination of the level of inflation once capacity limits were hit.

As macroeconomists began to think about quantitative modeling of the aggregate economy, the Phillips curve offered a way to make Keynesian inflation theory continuous and quantitative. The level of unemployment could be used to measure how far the economy was from capacity, and thereby to make quantitative predictions about how inflation would be affected by the level of aggregate demand. Policy, whether monetary or fiscal, was conceived as affecting inflation via a causal chain, from aggregate demand, to the level of output and employment (and thereby unemployment), to the rate of inflation. Through the 1960s and 1970s probably most economists thought about inflation-determination this way, and many still do. I am not arguing here that many economists think

such a two-equation recursive model of the economy is the full story of inflation determination, but simple one- and two-equation models are part of the mental furniture of most macroeconomists, and this particular simple model remains influential.

Primiceri (2006) models inflation-determination in the 1960s, 1970s, and 1980s as reflecting policymakers' use of a two-equation recursive model like this and learning over time about the value of its coefficients. One may be skeptical of Primiceri's results because of his assumption that the model about which the policymakers are learning is correct, with only the coefficient values uncertain. Nonetheless, the fact that Primiceri's interpretation of history works as well as it does may explain why this relatively simple way of thinking still has a hold on policymakers' thinking.¹

This is interesting, because we know that Lucas, in a series of papers in the late 1960s and early 1970s, some written with Rapping (Lucas and Rapping 1969a, 1969b; Lucas 1973), developed a model with some plausibility in which Phillips's empirical regularity could be misleading if used, as the Keynesian models were doing, to analyze the effects of policy. This new simple model arrived on the scene just as the United States entered a period in the 1970s of simultaneous high unemployment and high inflation, making the data in unemployment-inflation plots jump off the historical Phillips curve. The simple rational expectations version of this theory, in which the causal direction is reversed, with inflation surprises causing changes in unemployment, did not fit the data any better than the deteriorating standard Phillips curve model, but it provided a qualitative story about why a Phillips curve might first appear in the data, then disappear in the presence of Keynesian policymaking.

While a few of the early advocates of rational expectations modeling (such as Sargent 1981) held out the hope that it would generate cross-equation restrictions that would lead to improved quantitative policy models, the new theory was more commonly interpreted as implying the entire enterprise of large-scale policy modeling was quixotic. Simultaneous equation econometrics began to disappear from economics Ph.D. training in the United States, while every new Ph.D. could explain how the "Lucas critique" implied that Keynesian macroeconomic models

would lead to policy errors. With the simple “Lucas supply curve” (the rational expectations, reversed-direction, Phillips curve) replacing the Phillips curve, there was no further need for big policy models. The best monetary policy could do was to avoid creating surprises. Milton Friedman’s proposal of a fixed growth rate for the money stock (which he supported with a different set of arguments) fit well with the rational expectations policy analysis.

Meanwhile, those actually making monetary policy faced a continuing need to make decisions responsibly in the light of data emerging week after week. The Thatcher government’s experiment in the United Kingdom with a simple monetary growth rate policy rule showed that the historical statistical relationships among various measures of the money stock, and between the money stock and inflation and output, could deteriorate when exploited for policy purposes in the same way, and for the same reasons, that the empirical Phillips curve had decayed. With academic economic research turned almost entirely away from large-scale policy modeling, central bank economists developed their own solutions. They emerged with models that preserved many of the characteristics of the first generation of Keynesian models: equation-by-equation specification, an emphasis on flow equilibrium, and Phillips curves as the locus for the non-neutrality of monetary policy. Expectations now entered the models more pervasively, and the models, to sidestep the Lucas critique, made it at least formally possible to treat expectations as rational. The discipline of using simultaneous equations to make econometric inferences was entirely abandoned.²

For policy modeling, the simple Lucas supply curve was inadequate. Besides not fitting the data, its microeconomic underpinnings were either informal or, in formal models, highly abstract and unrealistic—for example, models of “island economies” in which people had to infer the value of the economy-wide interest rate or money stock from the price level on their own island. The policy models began by simply adding an inflation expectations term to the right-hand-side of the original Phillips curve, but there was no satisfactory theory of how such a relationship arose out of individual economic behavior. Into this gap sprang the New Keynesian Phillips curve.

2. Is the New Keynesian Phillips Curve Truly a Phillips Curve? Is It Useful?

The New Keynesian Phillips curve is not an empirical relation between unemployment and inflation. It nonetheless can play the same role as the Phillips curve in a policy model: it links a continuously varying, observable measure of “distance from capacity” to predictions about the rate of inflation. Furthermore, it provides a microeconomic story about how this relation emerges, a story in which people have rational expectations and have no money illusion. There are a number of reasons, though, to see the development of the New Keynesian Phillips curve as a Pyrrhic victory.

The theory of the New Keynesian Phillips curve is well known and documented elsewhere, as in Woodford (2003), so I will just summarize it here. A continuum of monopolistically competitive firms have control over their own prices, because of product differentiation, but have an incentive to keep their prices in line with those of other firms, because there are competitive pressures; however, these firms face some friction in price setting. The form of the friction varies, giving rise to different versions of the theory. One form holds that prices are set in contracts of fixed length, an idea first explored by Taylor (1979, 1980). Another, more convenient form suggested by Calvo (1983) is that prices are fixed for random periods, with the duration of the random period determined exogenously. There are further variations on the form of the friction, some of which will be discussed later. Because of the friction surrounding price setting, when the aggregate price level moves not all firms respond to the change at once, and this situation creates non-neutrality for monetary policy.

This New Keynesian theory sidesteps the Lucas critique because it contains expectations explicitly and assumes that expectations are rational. But the Lucas critique is only one special case of a generic problem we face in econometric modeling: we make simplifications and approximations that we realize are contingent, so that some kinds of changes in policy, or in the nature of exogenous disturbances, will force us to change the model. The New Keynesian Phillips curve is clearly unstable under some kinds of policy change—indeed under exactly the same kinds of policy

change that the Lucas critique claimed could undermine old Keynesian models. Though the agents in the New Keynesian model have rational expectations and no money illusion, the theory has simply moved the non-neutrality stemming from agent behavior itself into the constraints the agent faces, namely the price-setting frictions. The contract lengths of Taylor and Calvo theory are clearly not constants of nature; surely these durations will change systematically with the level, variability, and forecastability of inflation.

But there is a perhaps more important problem with the New Keynesian theory: it props up the simple Phillips curve way of thinking about the link from monetary policy to inflation. Though it suggests a different way of measuring real tightness—the “output gap” in place of unemployment—it still provides an equation in which real tightness appears as the crucial determinant of inflation. Of course, in principle, once inflation expectations are admitted to a Phillips curve equation, new style or old, it becomes possible for disturbances anywhere in the model to impact inflation directly, without any intermediating move in the measure of real tightness. If such influences are small, or slow-moving, it may nonetheless be helpful to think of inflation as determined, via a Phillips curve, by real tightness. But it is also possible that the opposite is true—the impact of policy and other disturbances on inflation is mainly direct, through the expectation term in the Phillips curve, so that retaining the Phillips curve as the central focus of informal thinking about inflation determination is misleading. Orphanides (2004) has explained how the U.S. inflation in the 1970s could have emerged from policymakers’ difficulties in real-time measurement of the output gap. But these difficulties played such a central role in good part because of thinking enforced by the Phillips curve paradigm—the notion that some measure based on real data, with no statistical input from inflation itself or inflation expectations, was the central determinant of inflationary or disinflationary pressure.

The New Keynesian theory gives a central role not to unemployment, but to the output gap. Recently the empirical literature (for instance, Sbordone 2003), has recognized that the output gap is actually important in the theory because it measures marginal cost, and has moved toward more direct measures of this variable, in particular to looking at the labor share of output.

It is reasonable then to ask whether we have any evidence on this issue: to what extent is some version of a Phillips curve central to the determination of inflation? For an answer we can look to monetary structural vector autoregressions—multivariate statistical models that distinguish monetary policy behavior, and disturbances to it, from other sources of variation in the economy—without imposing detailed interpretations of the estimated dynamics in terms of individual behavior. Such a model must allow for the very different monetary policy behavior during the 1979–1982 period of Volcker’s reserve targeting and for the substantial decline and shifts in relative sizes of volatility of disturbances between the pre-Volcker and post-1982 periods. The results discussed here are from models fit separately to quarterly data from 1959:Q1 through 1979:Q2 and from 1983:Q1 through 2008:Q1.³ The model was identified by assuming the pattern of zero restrictions shown in table 5.1 for the matrix of contemporaneous coefficients in a model of the form

$$(1) \quad A(L)y_t = \alpha + \varepsilon_t,$$

where the y vector consists of the federal funds rate, output, output per hour, the hourly wage rate, the price level, M1 money stock, and producer prices for crude materials (pcrm), as labeled in that order across the top of table 5.1. Output, output per hour, employee compensation, and the price level all refer to the nonfarm business sector.⁴

The first equation, labeled “monetary policy” in the table, represents monetary policy behavior. The zero restrictions in that row reflect the fact that GDP data are not available to policymakers within the quarter.⁵ The zero restrictions in the first and last column reflect an assumption that private-sector variables not set in auction markets (that is, those other than pcrm) respond only with a delay to interest rates or to rapidly fluctuating commodity prices. The triangular pattern of zeros in the lower part of the central five columns are simply normalizations. The last five equations are interpreted as a block that determines the central five variables, with the individual equations having no distinct interpretations. Though identifying assumptions like these are often characterized as controversial, behavioral models with complete interpretations often embed the same or similar timing assumptions in much more restrictive frameworks, in which these sorts of timing assumptions pass by without comment.

Table 5.1
Identifying Restrictions

	federal funds rate	output	output per hour (oph)	hourly wage rate (w)	price level (p)	M1 money stock	producer prices for crude materials (pcrm)
monetary policy	X	0	0	0	0	X	X
pcrm	X	X	X	X	X	X	X
output	0	X	X	X	X	X	0
output per hour	0	0	X	X	X	X	0
hourly wage rate	0	0	0	X	X	X	0
price level	0	0	0	0	X	X	0
M1 money stock	0	0	0	0	0	X	0

This kind of model can let us examine two questions. One, does the New Keynesian Phillips curve mechanism seem to be playing a central role in transmitting the effects of shifts in monetary policy? Two, does new information about the labor share variable and its predicted future values play an important role in changing inflation forecasts?

To answer the first question, figures 5.1 and 5.2 show the estimated impulse responses of all variables in the system to a monetary policy disturbance. All variables are measured in log units, except the interest rate, which is measured at an annual rate as a proportion (not percent). Labor share, the last entry at the lower right of each figure, is the response of the labor share variable, which is constructed from the responses to the hourly wage rate, output per hour, and the price level as $w - oph - p$. The pattern of responses is largely similar for the two subperiods, and also broadly similar to estimated responses to monetary policy disturbances in models estimated with other identifying assumptions. In particular, the federal funds rate rises, then returns to or falls below its original level; wages and prices both fall, though prices, in the earlier period especially, fall with more of a delay than for wages; output and output per hour fall. Note that the falls in prices and wages are persistent and predictable as soon as the monetary policy disturbance occurs, and the changes in




Figure 5.1
Pre-1979:Q3 Impulse Responses to Monetary Policy

output are almost as persistent. This does not fit well with stories that surprise changes in the inflation rate are what generates the real effects of monetary policy.

The most prominent difference between the two periods is that in the later period, after 1983, a monetary policy shock forecasts a humpbacked time path of further increases in the funds rate, followed by a later decline. In the earlier period, before 1979:Q3, the model estimates less of this interest-rate-smoothing behavior.

In both periods, the labor share variable moves very little in response to a monetary policy shock. In the earlier period, it moves up somewhat in response to a monetary contraction, though not by a lot relative to its overall standard deviation. This does not accord well with the notion that the New Keynesian Phillips curve can be the center of a causal chain in which restrictive monetary policy reduces inflation by reducing current




Figure 5.2
Post-1983 Impulse Responses to Monetary Policy

and expected future marginal costs measured this way. Yet this does not mean that a New Keynesian Phillips curve could not play an important role in aligning an equilibrium model with the data. Indeed, it looks as if the New Keynesian Phillips curve might help explain why the general price level is so slow to respond to monetary contraction—monetary contraction may produce a fall in productivity, rising or slowly falling marginal costs, and hence, via the New Keynesian Phillips curve mechanism, a tendency for price decreases to lag behind wage decreases.

We can also consider the second question—namely, how important is the New Keynesian Phillips curve as a way to understand determinants of inflation other than policy disturbances? For this we can look at the impulse responses of the price level to all seven sources of disturbance in the system, as in figures 5.3 and 5.4. Since labor share, in logs, is the

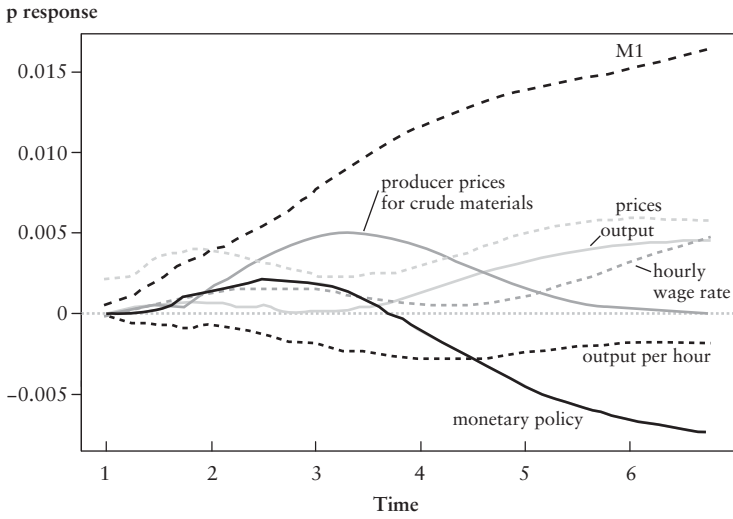


Figure 5.3
Pre-1979:Q3 Impulse Responses of the Price Level




Figure 5.4
Post-1983 Impulse Responses of the Price Level

hourly wage minus output per hour minus the price level, the effects of a surprise change in labor share are the corresponding linear combination of lines on the graph. In the pre-1979:Q3 graph, this nets out to close to zero. In the post-1983 graph, there is a strong effect of productivity surprises, with high productivity leading to increased inflation. This result is not due to any complicated behavior of productivity in response to its own shocks. Productivity shocks are the main source of variation in labor share, and these produce sustained, single-signed movements in labor share. It is not clear from this partially identified model what accounts for this pattern—but it is clear that the unidirectional New Keynesian Phillips curve causal chain is not at work here, as declining costs are associated with increasing inflation.

My conclusion is that the data show a perhaps surprisingly stable pattern of monetary policy influence on prices, wages, and output. Monetary policy is not neutral. But thinking about this pattern in terms of the New Keynesian Phillips curve does not appear to be helpful.

3. Inflation-Determination Without a Phillips Curve

If we cannot rely on a single Phillips curve-like equation to organize our thinking about inflation, what is the possible replacement? I think there are two main directions to pursue. One, which I will take up later, is to explore theories about deviations from the simple rational expectations paradigm. This may help us understand not only price stickiness and the non-neutrality of monetary policy, but also sluggishness and inertia in economic behavior more generally. The other, which can be fruitfully pursued even within the rational expectations framework, is to be more explicit and systematic in taking a full dynamic general equilibrium approach to macroeconomic modeling, and in particular to model more carefully the interaction of monetary policy with asset markets and the interaction of asset markets with the real economy.

Current and expected future fiscal and monetary policy have immediate and strong impacts on asset markets. In a fully articulated dynamic equilibrium model with rational agents, these impacts involve invoking transversality conditions. I have a colleague who interrupts every discussion of this kind of model with “Is this going to involve transversality

conditions?” His view is that few if any economists really understand transversality conditions (which is also my view) and that it is therefore unreasonable to entertain models that invoke transversality conditions to explain the behavior of actual human beings.

But transversality conditions apply even to less-than-hyper-rational agents. These conditions are really just another name for wealth effects. If monetary policy raises the rate of return on government bonds, and if agents project that this rise in the relative return of government paper will be persistent, government paper becomes more attractive, people will tend to trade other assets for government paper, and there will thus be downward pressure on the rate at which government paper trades for other goods—in other words, there will be downward pressure on the price level. But there are conditions under which a rise in interest rates on government bonds, generated by the central bank, will not lead bondholders to believe in persistently higher returns on government bonds. Higher real returns are possible, in general equilibrium, only if increased primary surpluses emerge in response to the higher interest rates. In an economy in which political economy or bureaucratic inefficiency makes increased primary surpluses impossible, the higher interest rates will only generate an increased rate of issuing government paper, with no increased rate of return—indeed with capital losses for holders of long nominal debt. It may take some time for bondholders to appreciate the nature of these fiscal dynamics, so that the inflationary effects of increased interest rates do not take hold immediately. But this only makes the real value of the outstanding debt at current prices increase more rapidly, so that when the realization that the increased debt has no real backing sinks in, the eventual effects on demand are even larger. This kind of situation is widely acknowledged to have existed in some countries and some time periods, especially where interest expense has become a large fraction of the total government debt and nominal interest rates are high.

Most macroeconomists, though, think of this type of scenario as applying perhaps to Brazil in some periods, but not ever to the United States. My view is that we should reevaluate this possibility. Our recent financial history of a stock market boom, a housing price boom, then a commodity price boom and a decline in the value of the dollar, may be best understood as reflecting the evolution of thinking by bondholders

about current and future U.S. monetary and fiscal policy. In the 1970s when the United States had its great burst of inflation, fiscal policy was by some measures much more unstable than monetary policy. On average over time, any country that can issue debt must be running primary surpluses—the conventional surplus, in addition to interest payments. The United States ran primary surpluses in all but four of the years from 1962 through 1974, for example, but ran primary deficits every year from 1975 through 1994, except for two years of small primary surpluses. Then from 1995 through 2002 it ran large primary surpluses, to the point where it seemed that the U.S. government debt might essentially vanish. And now we are again in a period of primary deficits. What ended the long period of primary deficits? What were bondholders thinking about future fiscal policy in this period? How did interest rate policy, which during the early 1980s was causing large changes in the size of the interest expense component of the budget, interact with the political economy of fiscal policy?⁶

These issues are of course only one component of a full general equilibrium approach to assessing the effects of monetary and fiscal policy on inflation. Nonetheless, it seems to me that there may be high returns to focusing more of our attention on this component, even at the expense of paying less attention to the microeconomics of price and wage dynamics.

4. Departing from Rational Expectations: New Ideas about Modeling the Effects of Uncertainty and Inertia

There is plenty of room for progress in integrating financial markets into our analysis of monetary and fiscal policy, even within the framework of rational behavior and what Sargent calls the “communist” assumption that there is a single probability measure shared by nature and by all the economic agents interacting in a given model. But increasingly economists are impatient with this assumption and interested in the implications of deviations from it. I will not try to catalog or discuss all the directions of deviation that economists have been exploring. Sargent, Williams, and Zha (2006), Evans and Honkapohja (2001), Cogley and Sargent (2005), and Sargent (1999), among many others, have explored the implications of learning, both by policymakers and private agents. Many economists,

in the area labeled “behavioral economics” that was set in motion by papers like Tversky and Kahneman (1974) and Kahneman and Tversky (1979), have tried to incorporate insights from psychology and experimental evidence about deviations from rational expectations. Mankiw and Reis (2007) have proposed a theory I think of as intermittent observation theory, in which agents process information only at certain widely separated moments.

Here, though, I want to concentrate on two lines of thinking that I find particularly interesting and promising. One is rational inattention theory, which posits why people do not use all of the information that lies in front of them “for free.” This theory invokes Shannon’s notion of a “channel” with finite “capacity” to process information, and assumes that people are such finite-capacity channels.⁷ This assumption implies that there are limits on how quickly and precisely people’s behavior can react to information about a stochastically evolving economic environment. The attractive feature of Shannon’s theory for engineers is that it allows discussion of information flows and the capacity of information channels in a way that is quantitatively precise, yet abstracts from the physical characteristics of the channel and of the information. These days we are all familiar with the notion that our Internet connections can be characterized by the bits per second figure that measures their Shannon capacity, and that this is a good measure of transmission speed whether we are transmitting photos of grandchildren, sending spreadsheets of historical GDP data, or downloading MP3 files from eMusic. The bits per second figure means the same thing for copper wire connections, fiber-optic connections, and cable connections.

This same hardware independence makes the theory attractive for modeling economic behavior, at least from the viewpoint of economists. It frees us from needing to know the details of the mental and physical limitations that prevent us from reacting at every moment to every bit of information impinging on us—we only need to know that the limitations exist, and to make the economist’s usual assumption that information processing capacity, like other resources, is used optimally. I have explored these ideas in several papers (1998, 2003, 2006). The 2003 paper shows that the theory implies modifications in the permanent

income model that bring it more closely in line with observed behavior. The 2006 paper considers a two-period savings model and shows that the theory can generate discretely distributed behavior, even in the face of continuously distributed information. By now a number of other economists have taken up these ideas, including Maćkowiak and Wiederholt (2009) and Matějka (2008), who show that some of the observed puzzling facts about microeconomic price behavior can be explained in the rational inattention framework.

The other area of recent research activity that I find interesting, though I have not contributed to it myself in published work, is exploring models in which rational agents, sharing the same information set and the same idea of the range of possible states of the world, have different probability distributions over those states. There is no reason why rational optimizing agents need share the same probability distribution. When optimizing agents with different probability distributions interact in markets, they will be attracted to betting with each other, if not explicitly, then by borrowing, lending, and making speculative investments. Furthermore, if optimistic investors, having borrowed from pessimists, discover they were mistaken in their expectations, there will be a rapid adjustment in asset prices, shifts in wealth between agents, and a high volume of transactions—all phenomena we see, and are concerned about, in actual asset markets. Papers that explore these possibilities include Scheinkman and Xiong (2003) and Brunnermeier and Julliard (2008).

5. Implications for Monetary Policy

So what are the implications of these new strands of research for the Phillips curve, monetary policy, and macroeconomics more generally? I do not have space to consider all the possibilities here, but some interrelated implications are worth drawing out.

Rational inattention implies that people will behave as if they are observing market signals with error, and that agents with a bigger stake will invest more of their Shannon capacity in obtaining a precise observation of a given signal. The rational inattention theory therefore provides one rationale for why economic agents might have different probability

distributions over the state of the economy, and for why these differences might persist despite the accumulation of freely observable evidence. Rational inattention and differences of opinion both may be related to why it is so hard, and yet so important, to model the interaction of asset markets with monetary policy and with the economy. Hard as it may be to model how a set of “communist” rational agents would have modeled the future of fiscal policy in the 1970s and 1980s, it is harder still to imagine that every agent—whether he held bonds or not, whether she was 23 years old or 75 years old, whether she was thinking of taking out a mortgage to buy a first home or had lived in the same house for 40 years and paid off her mortgage—had the same views about the future of fiscal policy and, therefore, the values of nominally denominated assets. Differences of views, learning, and rational inattention might explain why the interaction of monetary policy and fiscal policy with asset markets seems sometimes to work itself out on a long time scale. Not everyone will make the same assessment, at the same time, of the implications of transversality conditions. It may be that this divergence can lead to wide swings in asset markets, and to delayed and unpredictable effects of monetary policy shifts.

We cannot model every person’s beliefs individually, and working formally with rational inattention theory, at least at this point in time, seems hard.⁸ Nonetheless it seems important, especially in the light of the recent history of asset markets and their interaction with monetary policy, to get some working approximation of the effects of rational inattention and differences of beliefs into our models.

Where do these potential new approaches leave the Phillips curve model? Something like the Phillips curve will continue to have a place in a general equilibrium model, as part of characterizing the interaction of costs, prices, wages, and output. But the rational inattention perspective suggests that locating stickiness and inertia in that one equation may be a mistake. The same limits on information processing may be at work in the slow adaptation of prices and wages to each other that are at work in the sluggish reactions of consumption to income, or of investment to interest rates. Recognizing that sluggishness of various kinds may be related, through dependence on a common resource constraint, and that sluggish responses represent conservation of a valuable resource, may

lead to new modeling insights and to new ways of assessing the welfare implications of price stability and instability.

As with many important theories, the long-run value of Phillips curve theories may lie in the new flames that are emerging from their dying embers.

■ *This research was supported by National Science Foundation (NSF) grant SES-0719055. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the NSF.*

Notes

1. Sargent, Williams, and Zha (2006) also have a model of policymakers who learn about a relationship between inflation and unemployment, but their policymakers use regressions with unemployment on the left, depending on current and past inflation, unlike actual policymakers, who in this period thought of current and past unemployment as determining current inflation. Furthermore, their estimates implausibly imply that policymakers acted on beliefs in the 1970s that unemployment would undergo wild oscillations without deliberate counteracting movements in inflation generated by policy.
2. I discussed the state of central bank modeling in a 2002 Brookings paper.
3. The data were all drawn from the FRED website of the Federal Reserve Bank of St. Louis in 2008, during the last week of May and the first week of June.
4. The model was estimated using a dummy observation “Minnesota prior,” shrinking toward an independent random walks prior mean. The R programs used to estimate the model are available via the subversion internet protocol at <svn://sims.princeton.edu/R>.
5. Employment data are available monthly, and the data from before 1979:Q3 provide some indication that in that period a positive within-period response of the funds rate to employment changes should not have been ruled out.
6. In a 2008 paper I elaborate these points and present a model in which fiscal policy might have prevented the Fed from controlling inflation in the 1970s, even though the Fed was capable of creating recessions and corresponding temporary pauses in inflation.
7. Shannon’s theory is presented in MacKay (2003) or Cover and Thomas (1991), for example.
8. Working formally with rational expectations theory also appeared hard at one point, though, so in the future this gloomy prognosis may change for the better.

References

- Brunnermeier, Markus K., and Christian Julliard. 2008. "Money Illusion and Housing Frenzies." *Review of Financial Studies* 21(1): 135–180.
- Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–398.
- Cogley, Timothy, and Thomas J. Sargent. 2005. "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics* 8(2): 528–563.
- Cover, Thomas M., and Joy A. Thomas. 1991. *Elements of Information Theory*. New York: Wiley-Interscience.
- Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2): 263–291.
- Lucas, Robert E. Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63(3): 326–334.
- Lucas, Robert E. Jr., and Leonard A. Rapping. 1969a. "Price Expectations and the Phillips Curve." *American Economic Review* 59(3): 342–350.
- Lucas, Robert E. Jr., and Leonard A. Rapping. 1969b. "Real Wages, Employment, and Inflation." *Journal of Political Economy* 77(5): 721–754.
- MacKay, David J. C. 2003. *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99(3): 769–803.
- Mankiw, N. Gregory, and Ricardo Reis. 2007. "Sticky Information in General Equilibrium." *Journal of the European Economic Association* 5(2–3): 603–613.
- Matějka, Filip. 2008. "Rationally Inattentive Seller: Sales and Discrete Pricing." Working Paper, Program in Applied and Computational Mathematics, Princeton University. Available at http://www.pacm.princeton.edu/publications/Matejka_F_2008-wp.pdf.
- Orphanides, Athanasios. 2004. "Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches." *Journal of Money, Credit and Banking* 36(2): 151–175.
- Phillips, A. W. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957." *Economica* 25(100): 283–299.
- Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121(3): 867–901.

- Sargent, Thomas J. 1981. "Interpreting Economic Time Series." *Journal of Political Economy* 89(2): 213–248.
- Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.
- Sargent, Thomas, Noah Williams, and Tao Zha. 2006. "Shocks and Government Beliefs: The Rise and Fall of American Inflation." *American Economic Review* 96(4): 1193–1224.
- Sbordone, Argia. 2003. "U.S. Wage and Price Dynamics: A Limited-Information Approach." *International Journal of Central Banking* 2(3): 155–191. Available at <http://www.ijcb.org/journal/ijcb06q3a5.pdf>.
- Scheinkman, José, and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." *Journal of Political Economy* 111(6): 1183–1219.
- Sims, Christopher A. 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49 (December): 317–356.
- Sims, Christopher A. 2002. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity* (2): 1–62.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665–690.
- Sims, Christopher A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96(2): 158–163.
- Sims, Christopher A. 2008. "Stepping on a Rake: The Role of Fiscal Policy in the Inflation of the 1970s." Forthcoming in *The Great Inflation*, ed. Michael D. Bordo and Athanasios Orphanides. Chicago: University of Chicago Press. Available at <http://www.nber.org/chapters/c9162.pdf>.
- Taylor, John B. 1979. "Staggered Wage Setting in a Macro Model." *American Economic Review* 69(2): 108–113.
- Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88(1): 1–23.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185(4157): 1124–1131.
- Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

Comments on “Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy” by Christopher A. Sims

Michael T. Kiley

Inflation expectations play a central role in models of the Phillips curve. At long time horizons, inflation expectations may reflect the credibility of a monetary authority’s commitment to price stability. These observations highlight the importance of inflation expectations for monetary policy. These comments touch on three issues regarding inflation expectations:

- The evolving treatment of inflation expectations in empirical Phillips curve models
- Three recent models of information imperfections and inflation expectations
- Potential policy implications of different models.

The discussion will highlight two points: while historical experience suggests an important role for some deviation from the most restricted form of rational expectations in inflation dynamics, it also shows that other aspects of sluggish price adjustment—such as nominal rigidities—are important. The available indicators of inflation expectations show that imperfect information regarding central bank intentions has been one source of inertia in the formation of inflation expectations.

1. Inflation Expectations in the Phillips Curve

The Phillips curve has come a long way from its original 1958 specification. At the Federal Reserve Board’s conference on empirical work on price determination held in 1970, the dominant paradigm was adaptive expectations (Eckstein 1972). Robert Lucas’s contribution at this conference is stunningly familiar to a reader today in its approach and emphasis

on rational expectations—the idea that agents form expectations optimally given their understanding of the economy and the information available to them (Lucas 1972).

The rational expectations revolution had two quite opposite effects on subsequent empirical research regarding inflation. One branch of the literature, exemplified in the early contributions by John Taylor (1980) and by Julio Rotemberg (1982), followed Lucas's suggestion closely and specified tightly parameterized models incorporating various types of nominal rigidities, which led to various restrictions on a system of equations that would allow econometric identification. Another branch responded to the broader criticism leveled by Christopher Sims (1980) that the types of schemes traditionally used for identification in reduced-form Phillips curves and other empirical research were fatally flawed—"incredible," to use the terminology Sims employed—and looked to develop empirical techniques that imposed fewer restrictions on the data. Both lines of research bore significant fruit: modern dynamic general equilibrium models, with a large number of frictions, fit U.S. macroeconomic data quite well (see, for instance, Smets and Wouters 2007; Edge, Kiley, and Laforce 2008), and the set of stylized facts gleaned from analyses of vector autoregressions with minimal identifying restrictions has had a profound impact on the way such dynamic general equilibrium models are specified (see Christiano, Eichenbaum, and Evans 2005). However, it seems fair to say that the need for structural models to use when considering policy changes that represent significant departures from historically typical behavior implies the case may often be that tightly parameterized structural models will play an important role. This need for better predictive structural models is, at some level, troublesome. As someone actively involved in specifying relatively large and rich dynamic equilibrium models, my perspective is that the underlying assumptions used to achieve identification are clearly "incredible." But making empirical progress sometimes requires that economists make incredible assumptions; subsequent research strives to remove the need for such assumptions.

One area where such research on nominal price and wage rigidities is already being enriched, and may result in more plausible models, is in relaxing the assumption that perfect and homogeneous information sets

underlie price-setting behavior, which has been the primary assumption since Taylor (1980) and Rotemberg (1982). It is interesting to note that Robert Lucas emphasized expectations that were rational subject to an information constraint—in his model presented at the Fed's 1970 conference and in related work (Lucas 1972, 1973), agents in the economy only imperfectly perceived aggregate conditions.¹

While research has found greater empirical support for the simple full-information rational-expectations model in recent data (for instance, see Kiley 2007), my reading of the evidence from the aggregate inflation dynamics literature suggests that some type of information constraint is needed to explain fluctuations in U.S. inflation over the past forty years. These stylized facts about postwar U.S. inflation are well known:

- Inflation seems to respond sluggishly to (some) aggregate disturbances (see Christiano, Eichenbaum, and Evans 2000)
- The costs of disinflation are sizable (see Ball 1994)
- Inflation dynamics seem well characterized by a Phillips curve in which both leads and lags of inflation are important, especially for data including the 1970s and early 1980s (Fuhrer and Moore 1995, Kiley 2007).

But these stylized facts provide little guidance regarding what types of information imperfections are important to help explain fluctuations in the U.S. inflation rate.

2. Models of Imperfect Information and Inflation Expectations

Imperfections in information, and how this may influence inflation expectations and the behavior of economic agents, has been the subject of some research, which includes:

- Learning about the structure of the economy (see Orphanides and Williams 2005, 2007)
- How imperfect information impacts upon the goals (or credibility) of the central bank (see Ball 1995, Bomfim et al. 1997, Erceg and Levin 2003, Kiley 2008)
- Understanding the costs or constraints on information acquisition or processing (Caballero 1989; Kiley 2000, 2007; Carroll 2003; Reis 2006a, 2006b; Sims 1998, 2003, 2006)

Interestingly, it is primarily the third example of imperfect information, which emphasizes the costs of acquiring or processing information, that gives rise to differences in information sets across agents, as emphasized in Lucas (1972, 1973).

Each of these theories is capable of explaining costly disinflations and providing evidence on inflation dynamics. To date, research has not compared the ability of each model to fit the data relative to the other models. I will focus my attention on two of these models: imperfect information regarding the inflation objective of the monetary authority, and models of costly information acquisition.

The idea behind models emphasizing imperfect information regarding the inflation objective is simple: in an environment where the inflation objective of the monetary authority is not explicit or widely known, households and firms will need to infer where the monetary authority intends to bring inflation from its policy actions. As a result, agents make persistent mistakes regarding the inflation objective during a transition period after the objective has shifted. This idea seems a plausible explanation of inflation dynamics. For example, the Federal Reserve did not reveal an explicit objective for inflation in the period around 1980 when it began its effort to bring inflation down from undesirably high levels. And models incorporating imperfect information can explain many of the stylized facts regarding inflation dynamics, including the costs of the Volcker disinflation (see Erceg and Levin 2003) and the slow evolution of survey measures of long-run inflation expectations (Kiley 2007).

Models of costly information acquisition can also explain many of the stylized facts regarding inflation. Importantly, research in both the information processing tradition (such as Sims 1998, 2003, 2006) and in the information cost/infrequent updating tradition (for instance, Reis 2006a, 2006b) have emphasized that these models can explain the sluggishness of adjustment in much broader contexts—meaning that in addition to helping explain inflation, such models may help us better explain the observed patterns in consumption, investment, and other variables. The ability of such models to explain a range of facts is a great strength.

Christopher Sims's work on rational inattention is built upon an especially solid foundation: the communications literature has developed

axiomatic descriptions of uncertainty and analyzed how constraints on information processing capacity affect choices regarding information flow. However, the payoff of applying this research to economic modeling has yet to be realized, as this area is very complex. The complexity of economic problems reflects their dynamic nature, the endogeneity of aggregate variables that can provide information, and the possibility of rich information production functions.

Moreover, the information imperfections emphasized by Sims (2003, 2006) lie in how information is processed by economic agents. An added area of complexity in economic problems involves the computation of optimal actions by firms and workers under uncertainty and with highly nonlinear objective functions and constraints. While the computational tools to solve such problems are well understood by economists for many simple parametric examples, it is not obvious that the costs of such computations are trivial to the economic agents making these decisions. Modeling approaches that emphasize such costs, and their impact on the form of decisions made by workers and firms, may prove just as valuable in furthering our understanding of wage and price behavior as the information-processing approach that has been the subject of research by Sims and others.²

Finally, I should also note that some recent research has suggested that a model in which inflation and other expectations reflect some type of information imperfection may be capable of explaining the data on prices without reference to nominal price rigidities. I think this is unlikely for at least three reasons. First, my own empirical work on the Phillips curve that compared some sticky price and sticky information models showed clear evidence supporting the sticky price specification (Kiley 2007). Second, my dissertation considered a model with endogenous price and information rigidities and showed that sticky price and imperfect information models implied different effects on the form of the Phillips curve from the trend inflation rate and the higher moments of inflation. Cross-country evidence clearly shows evidence of the link between mean inflation and the form of the Phillips curve, suggesting an important role for sticky prices (Kiley 2000). Finally, casual observation and microeconomic evidence supports a role for infrequent adjustment of nominal prices and wages in macroeconomic models (see Nakamura and

Steinsson 2008). These considerations suggest that a model combining a microeconomic foundation for nominal price rigidities with the costs of acquiring or processing information will best account for the microeconomic and macroeconomic aspects of price adjustment that have been documented. Woodford (2008) presents a step in this direction.

3. Inflation Expectations and Monetary Policy

Given the model emphasizing imperfect information, I will turn to some of the implications this poses for policymakers' inflation objective.

In the United States, we have a few data sources regarding the long-horizon of inflation expectations; these sources are surveys of household and professional forecasters and measures of inflation compensation implied by yields on nominal and inflation-indexed Treasury securities. Models emphasizing imperfect information regarding the policymakers' inflation objective imply that these data should show a link between monetary policy actions or the policy regime and these long-horizon inflation expectations—as inflation expectations at long horizons should reflect, to a significant extent, the expectations of households and firms regarding the inflation objective. The data show such links in several areas.

The top panel of figure 5.5 presents data on long-horizon inflation expectations from the Reuters/Michigan Survey of Households and the Survey of Professional Forecasters.³ Long-horizon expectations from the Michigan Survey or from professional forecasters are sporadically available prior to the 1990s; both series are available continuously since the early 1990s. It is clear that long-horizon inflation expectations have fallen over the past 15 to 25 years; the recent data arguably provide evidence of some degree of anchoring.

Kiley (2008) shows that the survey measures of long-horizon inflation expectations in the United States behave very much as implied by models emphasizing uncertainty regarding the central bank's inflation objective: long-horizon inflation expectations respond to policy actions that are deemed either restrictive or expansionary, as judged by deviations from the Taylor rule, and the quantitative magnitude of such responses is consistent with plausible specifications and estimates of the costs of disinflation. This relationship can be expressed graphically. I define the

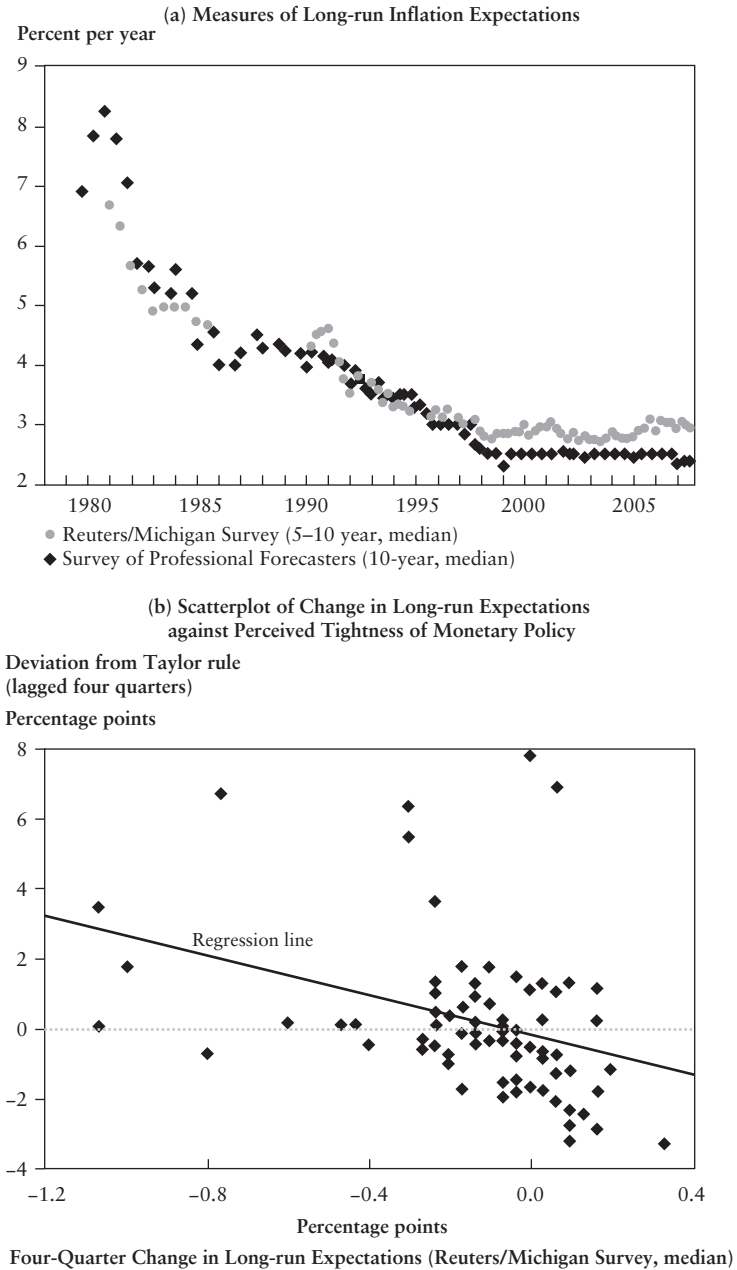


Figure 5.5
 U.S. Inflation Expectations Over a Long-Term Horizon
Source: Author's calculations.

perceived tightness of monetary policy as the gap between the nominal federal funds rate ($r(t)$) and the level predicted by the following Taylor rule involving consumer price inflation ($p(t)$) (as measured by the CPI), a perceived inflation target measured by long-run inflation expectations from the Reuters/Michigan Survey ($p^*(t)$), and the output gap⁴ ($y(t)$)

$$r(t) = 2 + p(t) + 0.5 * (p(t) - p^*(t)) + 0.5 * y(t).$$

The bottom panel of figure 5.5 plots the four-quarter change in the level of long-run expected inflation against this measure of perceived tightness of monetary policy lagged four quarters; there is a strong negative relationship.

Gürkaynak, Levin, and Swanson (2006) and Beechey, Johannsen, and Levin (2008) examine differences in the properties of measures of inflation compensation and inflation expectations between the United States and other countries that have explicit inflation objectives. Their analyses demonstrate that both survey measures of inflation expectations and inflation compensation implied by government-issued securities appear more stable in countries with an explicit inflation objective. Interestingly, there is pretty clear evidence that the dispersion in long-run expected inflation is much smaller in some inflation targeting countries. Figure 5.6 compares the dispersion in long-run expected inflation among professional forecasters for the euro area and the United States, following Beechey, Johannsen, and Levin (2008); dispersion is much lower in the euro area, which has an explicit long-run objective for inflation. This seems clearly consistent with a model emphasizing uncertainty regarding the long-run inflation objective as an important aspect of the link between monetary policy, expectations, and inflation or economic activity.

Finally, recent research has emphasized that inflation uncertainty is a significant factor determining term premia on nominal bonds (and hence the slope of the term structure; see Wright 2008). The link between the monetary policy regime and uncertainty about long-horizon inflation objectives suggests that the nature of such regimes can be expected to have significant effects on the term structure of interest rates. Figure 5.7 provides such an example. The Bank of England was given operational independence on




Figure 5.6

Cross-Sectional Dispersion in Long-Run Inflation Expectations

Source: European Central Bank Survey of Professional Forecasters and Survey of Professional Forecasters (from the Federal Reserve Bank of Philadelphia; see note 4) .

May 6, 1997, and the slope of the nominal yield curve flattened considerably that day, while the slope of the real yield curve changed very little.

Taken together, I interpret the set of results illustrated in these three figures as suggesting that the setting of monetary policy and the nature of the policy regime are important determinants of inflation expectations and, potentially, macroeconomic performance more generally.

These concrete results provide an example of a type of analysis that cannot yet be supported by research in, for example, the rational-inattention vein. It may be, however, that these stylized facts are very consistent with a model of rational inattention. Given the potentially wide-ranging implications of such models for macroeconomic dynamics that have been tentatively suggested in previous work, I view research in this direction as very promising. In the meantime, I also think that research emphasizing particular stories that may be applicable to policy considerations in the short-to-medium run, like that motivated by the model of imperfect information regarding the inflation objective, is likely to have a direct impact on policy discussions.

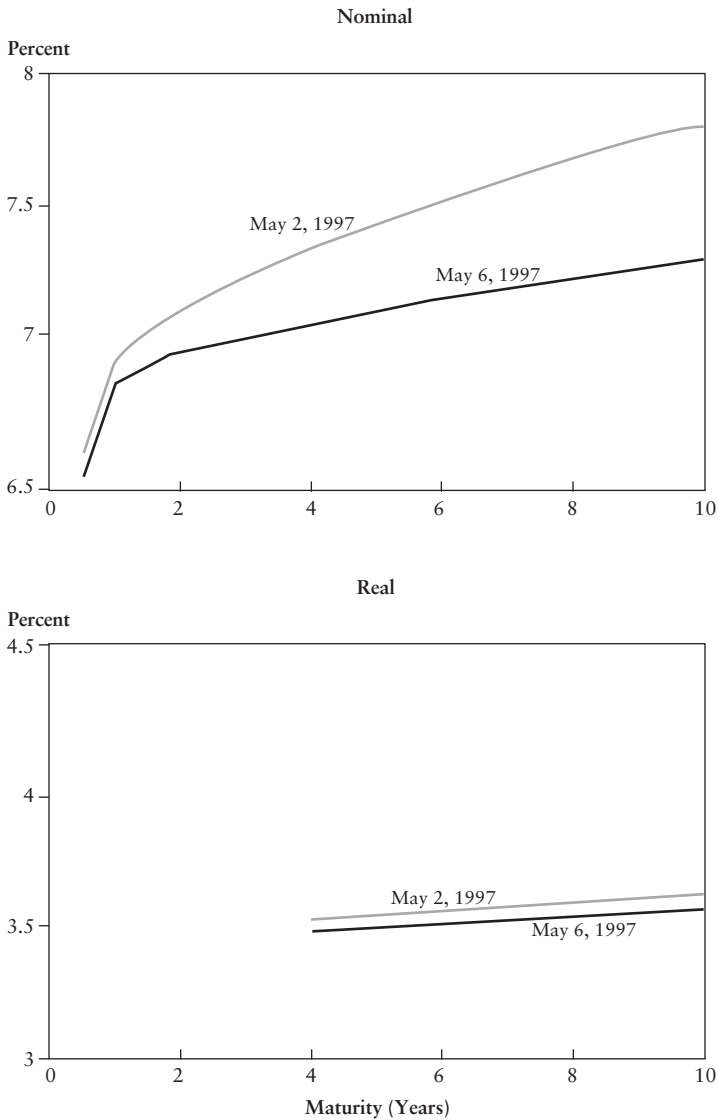


Figure 5.7
Nominal and Real Forward Curves in the United Kingdom around Granting of Independence to the Bank of England
Source: Wright 2008.

■ *I would like to thank Andrew Levin for helping me think about the issues discussed in this paper, and Ben Johanssen and Jonathan Wright for providing assistance with data and figures. The views expressed here are the author's, and do not reflect those of the Federal Reserve Board or its staff.*

Notes

1. A glance back at John Muth's (1961) article on rational expectations reveals that his idea was more in line with the information-constrained version of rational expectations: "The hypothesis can be rephrased a little more precisely as follows: that expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the 'objective' probability distributions of outcomes)" (Muth 1961, 316). In particular, Muth suggests that expectations are distributed around the mathematical expectation, implying some difference in signals available to different agents or errors in expectations.
2. Gabaix and Laibson (2000) present an example of this type of research. Their analysis is sufficiently distant from the form of a dynamic price adjustment problem that development of this type of reasoning to the problem of inflation dynamics is a substantial challenge, with unclear payoff.
3. This survey was formerly conducted by the American Statistical Association (ASA) and the National Bureau of Economic Research (NBER) and was known as the ASA/NBER survey. The Federal Reserve Bank of Philadelphia, in cooperation with the NBER, assumed responsibility for the survey in June 1990. With the exception of two observations, namely 1990:Q1 and 1991:Q2, the data from 1979 through 1991 are from the Blue Chip Survey. The two exceptions are from The Livingston Survey.
4. From the Federal Reserve Board's FRB/US model.

References

- Ball, Laurence M. 1994. "What Determines the Sacrifice Ratio?" In *Monetary Policy*, ed. N. Gregory Mankiw, 155–182. Chicago: University of Chicago Press.
- Ball, Laurence. 1995. "Disinflation with Imperfect Credibility." *Journal of Monetary Economics* 35(1): 5–23.

- Beechey, Meredith J., Benjamin K. Johansson, and Andrew Levin. 2008. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area than in the United States?" Federal Reserve Board Finance and Economics Discussion Series 2008–23. Available at <http://www.federalreserve.gov/Pubs/feds/2008/200823/200823pap.pdf>.
- Bomfim, Ant3lio, Robert Tetlow, Peter von zur Muehlen, and John Williams. 1997. "Expectations, Learning and the Costs of Disinflation: Experiments using the FRB/US Model." Federal Reserve Board Finance and Economics Discussion Series 1997–42. Available at <http://www.federalreserve.gov/Pubs/feds/1997/199742/199742pap.pdf>.
- Caballero, Ricardo J. 1989. "Time-Dependent Rules, Aggregate Stickiness, and Information Externalities." Columbia University Department of Economics Discussion Paper Series No. 428. Available at <http://econ-www.mit.edu/files/134>.
- Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118(1): 269–298.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2000. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford, 65–148. New York: Elsevier.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.
- Eckstein, Otto, ed. 1972. *The Econometrics of Price Determination*. Washington, DC: Board of Governors of the Federal Reserve System.
- Edge, Rochelle, Michael T. Kiley, and Jean-Philippe Laforte. 2008. "A Comparison of Forecast Performance between the FRB Staff Forecasts, Simple Reduced-Form Models, and a DSGE Model." Working Paper. Available at http://www.indexmeasures.com/dc2008/papers/ekl_forecast_2008.pdf.
- Erceg, Christopher J., and Andrew T. Levin. 2003. "Imperfect Credibility and Inflation Persistence." *Journal of Monetary Economics* 50(4): 915–944.
- Fuhrer, Jeff, and George Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110(1): 127–159.
- Gabaix, Xavier, and David Laibson. 2000. "A Boundedly Rational Decision Algorithm." *American Economic Review* 90(2): 433–438.
- G3rkaynak, Refet S., Andrew T. Levin, and Eric T. Swanson. 2006. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-Term Bond Yields in the U.S., U.K., and Sweden." Federal Reserve Bank of San Francisco Working Paper Series 2006-09. San Francisco: Federal Reserve Bank of San Francisco. Available at <http://www.frbsf.org/publications/economics/papers/2006/wp06-09bk.pdf>.

- Kiley, Michael T. 2000. "Endogenous Price Stickiness and Business Cycle Persistence." *Journal of Money, Credit and Banking* 32(1): 28–53.
- Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit and Banking* 39(1, Supplement): 101–125.
- Kiley, Michael T. 2008. "Monetary Policy Actions and Long-Run Inflation Expectations." Federal Reserve Board Finance and Economics Discussion Paper Series 2008-03. Available at <http://www.federalreserve.gov/Pubs/feds/2008/200803/200803pap.pdf>.
- Lucas, Robert E. Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4(2): 103–124.
- Lucas, Robert E. Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63: 326–334.
- Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29(3): 315–335.
- Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415–1464.
- Orphanides, Athanasios, and John C. Williams. 2005. "Inflation Scares and Forecast-Based Monetary Policy." *Review of Economic Dynamics* 8(2): 498–527.
- Orphanides, Athanasios, and John C. Williams. 2007. "Robust Monetary Policy with Imperfect Knowledge." *Journal of Monetary Economics* 54(5): 1406–1435.
- Reis, Ricardo. 2006a. "Inattentive Producers." *Review of Economic Studies* 73(3): 793–821.
- Reis, Ricardo. 2006b. "Inattentive Consumers." *Journal of Monetary Economics* 53(8): 1761–1800.
- Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90(6): 1187–1211.
- Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica* 48(1): 1–48.
- Sims, Christopher A. 1998. "Stickiness." Carnegie-Rochester Series on Public Policy 49(1): 317–356.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665–690.
- Sims, Christopher A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96(2): 158–163.
- Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.

Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88(1): 1–23.

Wright, Jonathan H. 2008. "Term Premiums and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset." Federal Reserve Board Finance and Economics Discussion Paper Series 2008-25. Available at <http://www.federalreserve.gov/pubs/feds/2008/200825/200825pap.pdf>.

Woodford, Michael. 2008. "Information-Constrained State-Dependent Pricing." Working Paper 14620. Cambridge, MA: National Bureau of Economic Research.

Comments on “Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy” by Christopher A. Sims

Athanasios Orphanides

The paper gives us a view of the history of the Phillips curve stressing the role of inflation expectations. Christopher Sims talks about the usefulness of the Phillips curve concept, with a focus on the “new” Keynesian Phillips curve, for analyzing price determination. He pretty much concludes that it is not particularly useful, preferring instead to talk about inflation determination without a Phillips curve. Sims then examines departures from rational expectations and the implications of new ideas about expectations formation for monetary policy, building on the critique of what Tom Sargent has termed the “communism” of rational expectations.¹

Why the emphasis on inflation expectations and the Phillips curve? Why do we have a whole session on it at this conference? The reason is because the Phillips curve has been at the core of thinking about macroeconomic stabilization for as long as macroeconomics has existed—indeed much before Phillips’s 1958 paper. As early as the 1920s, for example, academics and policy researchers at the Federal Reserve were talking in terms of concepts that today we identify with the Phillips curve.²

There are two central elements for understanding inflation in the Phillips curve framework. One is the concept of economic slack (the output gap) and the other is inflation expectations. Both elements are unobservable and potentially problematic, subject to misunderstanding and misuse. Regarding the output gap in particular, we are so ignorant about its proper definition and measurement in real time that it leaves little, if any, room for it to be useful for policy.³ So when we think about the Phillips curve, inflation expectations must be at the center of policy considerations.

Why are inflation expectations so important? To begin with, inflation expectations are a crucial determinant of actual price and wage setting and, therefore, actual inflation over time. In the Phillips curve context, this is the case regardless of what one may think the proper measure of the output gap may be. Well-anchored inflation expectations are essential not only for securing price stability, but also for facilitating overall economic stability over time. As we know, when private inflation expectations become unmoored from the central bank's objective—episodes characterized by Marvin Goodfriend (1993) as “inflation scares”—macroeconomic stabilization can suffer. With unanchored inflation expectations, the Phillips curve becomes too unpredictable to be a useful concept for policy guidance.

Well-anchored inflation expectations are arguably most important during times that may be most challenging for monetary policy—such as what we are going through at present. In such times, two things are useful to keep in mind. One is that monetary policy can have considerably greater leeway in responding to adverse supply shocks when inflation expectations are well anchored. This can be easily understood in the context of Phillips curve analysis. The other is that the central bank can also have greater flexibility for swift responses to financial disturbances if it can be confident that inflation expectations will remain well anchored. I think we've seen this in practice since August 2007, again and again.

So what about the importance of inflation expectations under the present circumstances? We need only look at very recent statements about monetary policy by the European Central Bank (ECB) and the Federal Reserve to appreciate the central role policymakers attach to them. Consider the introductory statement read by ECB President Jean-Claude Trichet at the press conference following the Governing Council's monetary policy decision meeting on June 5, 2008: “Against this background, it is imperative to secure that medium to longer-term inflation expectations remain firmly anchored in line with price stability” (Trichet 2008, p. 2). And compare this with what I thought was a key sentence from Federal Reserve Chairman Ben Bernanke's address at this conference: “The Federal Open Market Committee will strongly resist an erosion of longer-term inflation expectations, as an unanchoring of those expectations would be destabilizing for growth as well as for inflation.”⁴ When

comparing the two statements, it is clear that well-anchored inflation expectations are the underlying thinking in both central banks.

Despite the central role of inflation expectations in the Phillips curve and the emphasis policy practitioners place on them, most of the Phillips curve models we have seen in the last 50 years rely on rather simplistic and unrealistic models of inflation expectations. Two examples serve as useful illustrations of this simplicity and lack of realism. In the early days of Phillips curve modeling, we had the linear fixed-distributed-lag models, and later the rational expectations models under perfect knowledge. At present the traditional modeling still, by and large, imposes rational expectations in a world with fixed and perfectly known structures, including known and stable policy preferences. As Jim Stock reminded us earlier at this conference, this practice is really equivalent in some sense to the old-fashioned modeling of expectations, which was incorrectly called “adaptive expectations” back in the 1950s and 1960s. I say “incorrectly” because nothing adapted in these models—parameters were kept fixed, by assumption. Under these assumptions, using either the old-fashioned distributed-lag models or the new, more modern, if you wish, linear rational expectations models, the monetary policy problem is rather trivial and anchoring inflation expectations is a simple matter of policy adopting and adhering to a stable policy rule.

By downplaying the information limitations that either policymakers or economic agents likely face in reality and oversimplifying the expectations formation mechanism, both “old” and “new” Phillips curve models have some common issues. They may suggest, for example, that a Phillips curve model should be able to forecast inflation better than is likely achievable in practice. They may also suggest the existence of an exploitable short-run trade-off between price stability and economic stability. This may raise hopes among academics, and perhaps even among some policymakers, about what monetary policy might be able to achieve. Indeed, one should be concerned that, if misused, these models could lead both forecasters and policymakers astray. This is an issue I hope Sims will elaborate further in the context of his model.

I fully agree with Sims that it is useful to deviate from the assumption of rational expectations in Phillips curve models and to think a little bit more about learning and alternative models of expectations formation.

Recent work has explored various avenues for improving the expectations formation mechanisms embedded in Phillips curve models, thus making them more useful for policy analysis. A common element in these models that I wish to stress is the acknowledgment of the presence of imperfections in the formation of expectations, relative to the simplistic rational expectations model we continue to use as the benchmark model. As a result, these models can better capture the inherent limitations in gathering and processing information. One way of proceeding has been to posit that private agents may act as econometricians. I have in mind here an exercise along the lines of Stock and Watson (1999) who consider respecifying and reestimating forecasting models with the objective of obtaining in an adaptive manner the best forecast they can with available data. Why do they posit that the forecaster should respecify and reestimate the model in order to achieve this? This is because they are concerned about structural change and uncertainty, imperfections that influence expectations formation in practice.

To illustrate the implications of alternative treatments of expectations, I will present to you a simple example, reproduced from earlier work with John C. Williams (Orphanides and Williams 2005). Consider two economies with an identical Phillips curve and the same monetary policy rule. One treatment assumes that expectations are always well anchored, the rational expectations outcome based on perfect knowledge. In the alternative treatment, agents learn from recent economic outcomes in forming expectations. One way to interpret the experiment is to compare these two models and analyze the impact of an adverse supply shock to the economy. In other words, assess whether second-round effects from an adverse supply shock can be avoided, as we would say on the other side of the Atlantic. If so, then the impact of the adverse supply shock can be avoided and everything appears to be perfect. Otherwise, if these effects cannot be avoided, the economy may experience conditions that bring back memories of stagflation. In figure 5.8, the gray lines show what a supply shock would look like with well-anchored inflation expectations. We have a temporary hump in inflation and a mild slowdown in the economy. Yet in the case of unanchored inflation expectations, where expectations are formed with a learning mechanism, as the black lines




Figure 5.8

Evolution of Economy Following an Adverse Supply Shock

Source: Based on Orphanides and Williams (2005).

Notes: Gray line—with well-anchored inflation expectations, Black line—with unanchored inflation expectations.

show, second-round effects on inflation follow the original supply disturbance, which, if unchecked, lead to a protracted stagflationary episode.

This deviation from the simple linear rational expectations paradigm has a number of implications. Learning behavior in the formation of expectations introduces nonlinear dynamics in otherwise linear economies. Specifically, it induces time-variation in the formation of expectations, and hence in the structure of the economy, even in the absence of fundamental regime changes. This complicates empirical modeling, including estimation and forecasting of what might otherwise be a fixed-coefficient linear model. I note that all these issues arise in this environment simply because of the presence of an imperfection in expectations formation.

Even more interesting are the implications for monetary policy. Learning behavior in the formation of expectations may impart additional persistence to inflation for a given monetary policy, thereby diminishing the policymaker's ability to stabilize business cycle fluctuations in addition to maintaining price stability. This provides an explanation as to why the appearance of an exploitable policy trade-off in an estimated linear rational expectations Phillips curve model is unlikely to be useful in practice. Furthermore, perpetual learning with imperfect knowledge induces the endogenous inflation scares that can be particularly damaging to the economy without a forceful policy response. This provides an explanation as to why policymakers monitor inflation expectations so closely and place a premium on maintaining well-anchored inflation expectations.

There are also implications for policy communication. Recognition of the role of learning in the formation of expectations introduces a role for central bank communication that is absent in traditional models. To the extent that central bank communication can facilitate the formation of more accurate inflation expectations, it can prove useful for improving policy outcomes. In this light, clarity regarding the central bank's price stability objective may improve macroeconomic performance.

In summary, properly accounting for the formation of inflation expectations is essential for understanding the potential usefulness and inherent limitations of Phillips curve models. One may love or one may hate these models, but I think much of the difference in perspectives comes

from improper versus better accounting of how expectations are assumed to be formed in these models. Many of the puzzles which seem to be associated with Phillips curves could potentially be solved once we recognize the richness introduced by deviating from the old-fashioned fixed-coefficient-distributed-lag models or the new, but still old-fashioned, linear rational expectations models of the Phillips curve. New approaches that incorporate learning may better capture the formation of inflation expectations, and I think we can make a lot of progress by moving in this direction without completely throwing away the intellectual framework of the Phillips curve. Nonetheless, some of the pertinent lessons may be quite simple. First, clarity regarding the central bank's price stability objective may facilitate efforts to maintain well-behaved inflation expectations, even in the presence of a series of adverse shocks. And second, maintaining well-anchored inflation expectations over time enhances the central bank's ability to flexibly respond to financial disturbances as well as adverse supply shocks.

Notes

1. See Evans and Honkapohja (2005).
2. See Orphanides (2003) for references to early 1920s versions of the Phillips curve and their relation to the modern policy debate.
3. See Orphanides and van Norden (2005) for an empirical examination of the problems associated with real-time inflation forecasts based on output gaps.
4. This quotation is from a brief topical portion of Chairman Bernanke's remarks not included in this volume. Speaking about the immediate conditions faced by the U.S. economy in June 2008, he commented that:

Inflation has remained high, largely reflecting sharp increases in the prices of globally traded commodities. Thus far, the pass-through of high raw materials costs to the prices of most other products and to domestic labor costs has been limited, in part because of softening domestic demand. However, the continuation of this pattern is not guaranteed and future developments in this regard will bear close attention. Moreover, the latest round of increases in energy prices has added to the upside risks to inflation and inflation expectations. The Federal Open Market Committee will strongly resist an erosion of longer-term expectations, as an unanchoring of those expectations would be destabilizing for growth as well as for inflation.

References

- Evans, George, and Seppo Honkapohja. 2005. "An Interview with Thomas J. Sargent." *Macroeconomic Dynamics* 9(4): 561–583.
- Goodfriend, Marvin. 1993. "Interest Rate Policy and the Inflation Scare Problem: 1979–1992." *Federal Reserve Bank of Richmond Economic Quarterly* 79(1): 1–23.
- Orphanides, Athanasios. 2003. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50(5): 605–631.
- Orphanides, Athanasios, and Simon van Norden. 2005. "The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time." *Journal of Money, Credit and Banking* 37(3): 583–601.
- Orphanides, Athanasios, and John C. Williams. 2005. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation Targeting Debate*, ed. Ben S. Bernanke and Michael Woodford, 201–246. Chicago: University of Chicago Press.
- Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44(2): 293–335.
- Trichet, Jean-Claude. 2008. "Introductory Statement." Monetary Policy Decision Meeting, Frankfurt am Main, June 5. Available at <http://www.ecb.int/press/pressconf/2008/html/is080605.en.html>.

6

Evidence on Price Determination

Implications of Microeconomic Price Data for Macroeconomic Models

Bartosz Maćkowiak and Frank Smets

1. Introduction

In his *Handbook of Macroeconomics* chapter written a decade ago, John B. Taylor took stock of the microeconomic evidence on price and wage setting available at the time. He summarized the micro evidence in four points (Taylor 1999, pp. 1020–1021). First, Taylor concluded that studies of micro data did not support the casual observation that prices are changed more frequently than wages. Instead, he contended that studies of micro data suggest that price changes and wage changes have about the same average frequency—about one year. Second, Taylor noted that there is a great deal of heterogeneity in price setting and wage setting. Third, he concluded that neither price-setting nor wage-setting behaviors are synchronized. Fourth, Taylor observed that the frequency of price changes and wage changes depends positively on the average rate of inflation.

Most of the microeconomic evidence on price setting that Taylor (1999) reviewed was based on a relatively narrow set of products such as magazines (Cecchetti 1986) and goods sold from retail catalogs (Kashyap 1995). In the decade since Taylor wrote his assessment, economists have gained access to new, more detailed micro data on prices. These include data collected in order to compute consumer and producer price indexes in a number of countries as well as scanner data from supermarket chains. In this paper, we review the literature that studies the new micro data, and in the process revisit some of Taylor's conclusions.¹ Furthermore, we discuss implications of the new micro data for macroeconomic models.

The new micro data confirm that there is a great deal of heterogeneity in price setting, in terms of the frequency and the size of price changes,

but also in terms of how often sales occur and what form sales take. Despite all the heterogeneity, micro price data display a number of regularities. First, prices remain constant for extended periods of time in many sectors of the economy. Second, prices appear to change less frequently in the euro area than in the United States. Depending on whether price changes related to sales and item substitutions are included or excluded, the median consumer price lasts about four to nine months in the U.S. economy. The median consumer price lasts about 11 months in the euro area. Third, when prices change, on average these change by large amounts relative to inflation. This finding suggests that idiosyncratic shocks are a much more important cause of variation in prices than aggregate shocks. At the same time, many price changes are small. Fourth, new cross-country evidence confirms that the frequency of price changes depends positively on the average rate of inflation. Fifth, there is indeed little evidence of synchronization of price changes.

One reason why macroeconomic modelers are interested in micro price data is that in familiar and tractable macroeconomic models, the frequency of price changes maps easily into impulse responses of prices and quantities to aggregate shocks. The New Keynesian model with Calvo pricing is a case in point. We provide two examples to show that there may be no simple mapping between the frequency of price changes in microeconomic data and the speed of impulse responses of prices and quantities to shocks. The first example involves time series models estimated using sectoral price data, as in Boivin, Giannoni, and Mihov (2007) and Maćkowiak, Moench, and Wiederholt (2008). Both papers find that sectoral price indexes respond quickly to sector-specific shocks. In this sense, prices are not sticky. At the same time, sectoral price indexes respond slowly to macroeconomic shocks. In this sense, prices are sticky. If the frequency of price changes were decisive for impulse responses, one would expect prices to respond with roughly equal speed to both kinds of shocks. It turns out that, in the data, the degree of price stickiness—defined as the speed at which prices respond to shocks—appears to be conditional on the source of the shock. This does not imply that the frequency of price changes is irrelevant. The frequency of price changes helps explain the speed of impulse responses of prices to macroeconomic shocks in a cross-section of sectors.

The second example involves the dynamic stochastic general equilibrium (DSGE) model of Smets and Wouters (2003). When this model is estimated using macroeconomic data from, alternatively, the United States and the euro area, the estimated value of the Calvo parameter is somewhat larger for the euro area than for the United States. The difference is not statistically significant. Equivalently, the slope of the New Keynesian Phillips curve is estimated to be somewhat lower in the euro area than in the United States. Therefore, the Smets-Wouters model reflects the notion that prices are changed less frequently in the euro area than in the United States, as is apparent from microeconomic data, but the model's evidence is weak, not robust. Furthermore, the slope of the New Keynesian Phillips curve is estimated to be very low both in the United States and in the euro area. Whether prices change on average every 4 months or every 11 months, nominal rigidity by itself does not suffice to explain the low slope of the New Keynesian Phillips curve in the Smets-Wouters model.² One needs to combine nominal stickiness with a sufficient degree of real rigidity.³ The impulse responses, and therefore the predictions concerning the effects of macroeconomic shocks, depend on the entire DSGE model and all its parameters, not only on the price-setting mechanism and the parameter governing the frequency of price changes.

In an ideal world, macroeconomists would set up, solve, and understand a DSGE model with many heterogeneous firms and households, in which sellers and buyers in different sectors interact differently, and which matches in detail both microeconomic data and macroeconomic data. Realistically, all we can hope for, at least for some time, is a model that matches macroeconomic data well, while telling a reasonable story that is broadly in line with microeconomic data. Recent DSGE models have made progress in matching macroeconomic data. We survey the recent literature on models of price setting. We search for a story or a set of stories that would simultaneously (1) imply persistent impulse responses to macroeconomic shocks, similar to the impulse responses in the Smets-Wouters model, and (2) be broadly in line with micro price data. We do not think that macroeconomists have developed such a story yet. The most promising lines of research, from our point of view, involve work with models in which prices can respond quickly and by large amounts to idiosyncratic shocks and, at the same time, prices respond slowly and by

small amounts to macroeconomic shocks. One line of research with this feature involves models in which a high degree of real rigidity arises conditionally on a macroeconomic shock. Another line of research with this feature involves imperfect information about macroeconomic shocks. In both lines of research, some degree of flexibility at the microeconomic level can potentially coexist with a sizable amount of stickiness at the macroeconomic level. Both lines of research, or some combination of these, may eventually be capable of producing a DSGE model that fits macroeconomic data as well as, or better than, the Smets-Wouters model while providing a reasonable story that is broadly in line with micro price data.

Section 2 of this paper reviews the recent literature that studies new, more detailed micro data on prices. In section 3, we discuss the recent literature that compares standard models of price setting used in macroeconomics to the new micro price data. In section 4, we argue that there is no simple mapping from the frequency of price changes in micro data to impulse responses of prices and quantities to shocks. In section 5, we discuss promising lines of research. Concluding remarks are in section 6.

2. What Do the New Micro Price Data Say?

In this section, we review the recent literature that studies new, detailed microeconomic data on prices from the United States, the euro area, and other countries. We begin by discussing the literature that studies U.S. data. Next, we discuss the literature that studies euro-area data, and we summarize new cross-country evidence. At the end of this section, we discuss the evidence on causes of price rigidity from surveys of firms in the United States and in the euro area.

Micro Price Data in the United States

The data underlying the consumer price index in the United States are so rich that questions of the form “what do the micro data say?” have no simple answers. The question “how frequently do prices change?” has no simple answer, because in microeconomic data there is a distribution of the frequency of price changes. Klenow and Kryvtsov (2008) and

Nakamura and Steinsson (2008a) study the Consumer Price Index (CPI) Research Database provided by the Bureau of Labor Statistics (BLS). The CPI Research Database contains the nonshelter component of the data collected by the BLS in order to compute the CPI. The CPI Research Database covers all goods and services other than shelter, or about 70 percent of the CPI. The BLS divides goods and services into about 300 Entry Level Items, known as ELIs. On a monthly basis, the BLS collects prices for all products in the three largest metropolitan areas (New York, Los Angeles, and Chicago). Each month, the BLS collects prices in all areas for food and fuel products. The BLS collects prices for other products and other areas only bimonthly. The CPI Research Database begins in January 1988. Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008a) focus on prices collected monthly. Klenow and Kryvtsov use the data in the CPI Research Database through January 2005. Nakamura and Steinsson focus on the data from 1998.⁴ Figure 6.1 shows the distribution of the frequency of price changes for ELIs from Nakamura and Steinsson (2008a). The dispersion of the distribution is striking. The ELI category with the lowest frequency of price changes, 1.6 percent, is Legal Services. Only 1.6 percent of prices in the category Legal Services change on average from month to month. The ELI with the highest frequency of price changes, 100 percent, is Used Cars. All prices in this category change from month to month.⁵

Table 6.1 illustrates in another way the heterogeneity in the frequency of price changes. Table 6.1 reports the distribution of the frequency of price changes for so-called Major Groups, fairly broad subbaskets of the CPI basket. The median frequency of price changes for Major Groups ranges from 6.6 percent for Services Excluding Travel to 87.6 percent for Vehicle Fuel.⁶

Klenow and Kryvtsov (2008) estimate the median frequency of price changes between 1988 and 2004 to be 27.3 percent. The implied median price duration is the inverse of this number, 3.7 months. This means that half of all prices in the U.S. economy last less than 3.7 months.⁷ Klenow and Kryvtsov estimate the mean frequency of price changes between 1988 and 2004 to be 36.2 percent. One can understand why the mean is higher than the median by noting in figure 6.1 that there are a few ELIs




Figure 6.1
Frequency of Price Changes by Entry Level Item for 1998–2005,
from Nakamura and Steinsson (2008a)

Table 6.1
Frequency of Price Changes by Major Group for 1998–2005, from Nakamura
and Steinsson (2008a)

Major Group	Median Frequency of Price Changes (percent)
Services excluding Travel	6.6
Recreation Goods	11.9
Other Goods	15.5
Household Furnishings	19.4
Processed Food	25.9
Apparel	31.0
Transportation Goods	31.3
Unprocessed Food	37.3
Utilities	38.1
Travel	42.8
Vehicle Fuel	87.6

for which price changes are very frequent. The implied mean price duration is 6.8 months.

What does the heterogeneity in the frequency of price changes imply for macroeconomic models? In these models, we almost always use the convenient assumptions of a representative firm or many homogeneous firms. The new micro price data show that, in the real world, there are many heterogeneous firms and there is no single “representative firm.” In general, we cannot expect a macroeconomic model with many heterogeneous firms to behave like a macroeconomic model with one representative firm. One possible way forward for macroeconomic modelers is to construct a model with a representative firm and compare selected predictions of this model to a version of the same model with many heterogeneous firms. One can calibrate the representative-firm model so that its selected predictions come close to the version of the model with many heterogeneous firms. One can then have more trust in other predictions of the representative-firm model. This may be wishful thinking, because no representative-firm model may come close to the model with many heterogeneous firms. Furthermore, conclusions from this kind of analysis will be model-specific. Note also that the heterogeneity in the frequency of price changes is only one form of the heterogeneity present in micro price data. Even a model that takes into account fully the heterogeneity in the frequency of price changes will, in general, neglect other forms of heterogeneity in price-setting behavior. Carvalho (2006) and Nakamura and Steinsson (2008b) are examples of recent papers that introduce heterogeneity in price-setting behavior into macroeconomic models. Carvalho studies heterogeneity in price-setting behavior in time-dependent models, whereas Nakamura and Steinsson use a menu cost model. Both papers find that heterogeneity in price-setting behavior magnifies real effects of nominal shocks.

Answering the question “how frequently do prices change?” is complicated further by the presence of “sales” and “forced item substitutions” in micro price data. BLS employees who visit retail outlets in order to collect data on prices record certain prices as “sale” prices.⁸ Klenow and Kryvtsov (2008) report that about 11 percent of price quotes in their dataset are sale prices. “Forced item substitutions” occur when an item in

the sample has been discontinued from an outlet, and the BLS employee records the price of a similar replacement item in the outlet. This often takes the form of a product upgrade or model changeover, and about 80 percent of the time this involves a price change. The monthly rate of forced item substitutions is about 3 percent. So far, we have reported the frequency of price changes based on all posted prices. It turns out that the answer to the question “how frequently do prices change?” depends to a considerable degree on whether one excludes price changes related to sales and forced item substitutions, and how exactly one goes about doing this. Klenow and Kryvtsov (2008) find that looking only at “regular” prices—i.e., removing all sale-related price changes from the data—raises the estimated median price duration from 3.7 months to 7.2 months.⁹ Next, removing all forced item substitutions from the data increases the estimated median price duration to 8.7 months. Then, looking only at adjacent prices raises the estimated median price duration further to 9.3 months, and to 10.6 months if one restricts attention to the period from 1998 through 2004. Note that this is an estimate one obtains when looking only at consecutive monthly regular prices between forced item substitutions.¹⁰

Nakamura and Steinsson (2008a) observe that sale-related price changes are different from regular price changes. For example, sale-related price changes are larger and more transient than regular price changes. Klenow and Kryvtsov (2008) point out that a sale price is more likely to differ from the previous sale price than a regular price is to differ from the previous regular price. Sales are stochastic; these are not a fixed discount from the regular price. Matters are complicated yet further by the fact that sales and forced item substitutions are important in some sectors but not in other sectors, and firms operating in different sectors have different kinds of sales. Nakamura and Steinsson (2008a) emphasize the heterogeneity in the prevalence of sales and forced item substitutions across Major Groups. For example, 87 percent of price changes in Apparel, 67 percent of price changes in Household Furnishings, and 58 percent of price changes in Processed Food are sale-related price changes. The monthly rate of forced item substitutions is about 10 percent in Apparel and Transportation Goods, and about 6 percent in Recreation Groups, Utilities, Vehicle Fuel, Travel, and Services Excluding

Travel are sectors of the economy in which the fraction of sale-related price changes is close to zero. And there are sectors of the economy in which the monthly rate of forced item substitutions is close to zero—for example, Vehicle Fuel and Utilities. Klenow and Kryvtsov (2008) report that about 60 percent of sales are associated with a V-shape in a price quote, as the price goes down and after some time returns to the same presale regular price. The other common type of a sale is a clearance sale, often associated with repeated markdowns. Nakamura and Steinsson (2008a) report that clearance sales occur frequently in the Major Group Apparel, sometimes in the Major Groups Household Furnishings and Recreation Groups, and almost never in other Major Groups.

What does the presence of sales and forced item substitutions in micro price data imply for macroeconomic models? An optimizing model of sales will, in general, predict that the magnitude, frequency, and duration of sales respond to macroeconomic shocks. When aggregate productivity is growing quickly, products are likely to sell at bigger discounts and to be marked down more frequently and for longer time spans. Holding to the prior assumption of “exclude all sale-related price changes from macroeconomic models” may therefore be unjustified. Klenow and Willis (2007) estimate that sale-related price changes in the BLS data are at least as sensitive to inflation as are regular price changes. Furthermore, even if sales are caused by shocks orthogonal to the macroeconomy, the presence of sales may matter for the response of prices to macroeconomic shocks. In the rational inattention model of Maćkowiak and Wiederholt (2009), the idiosyncratic uncertainty faced by firms matters for the response of prices to macroeconomic shocks. As far as forced item substitutions are concerned, it seems reasonable to assume that product turnover is not caused by a desire for a price change. However, we find it intuitive that macroeconomic shocks play a role in a firm’s decision whether to use product turnover as a repricing opportunity or not, whatever the reasons behind product turnover. Macroeconomic shocks may also influence the timing and the frequency of product turnover. Isn’t it a good idea to replace old products with new ones when the aggregate economy is strong? Klenow and Willis (2007) estimate that price changes associated with forced item substitutions in the BLS data are sensitive to inflation.¹¹ Kehoe and Midrigan (2008) is a recent paper incorporating V-shaped sales

into a menu cost model. Kehoe and Midrigan find that in the menu cost model with sales, nominal shocks have considerably larger real effects compared to the same model without sales calibrated to all posted prices. The reason is that sale-related price changes are transient, and therefore do not offer firms much of an opportunity to respond to persistent nominal shocks.¹²

Bils and Klenow (2004) analyze determinants of the frequency of price changes. They find that products sold in competitive markets, as measured by concentration ratios or wholesale markups, display more frequent price changes. However, this result disappears once Bils and Klenow control for a good being energy-related, or a fresh food. We later discuss the determinants of the frequency of price changes using evidence from the euro area.

We have argued that the question “how frequently do prices change?” has no simple answer. There is no simple answer either to the related question “by how much do prices change?” In micro price data, there is a distribution of the size of price changes, and it turns out that this distribution has fat tails. A typical price change is large, and many price changes are small. Klenow and Kryvtsov (2008) find that, conditional on a price change, the median absolute size of the price change is 11.5 percent. When sale-related price changes are excluded, the median absolute size of the price change falls somewhat to 9.7 percent.¹³ Note that the average monthly inflation rate in the sample period of Klenow and Kryvtsov was 0.2 percent. This means that, even excluding sale-related price changes, price changes are on average large relative to inflation. At the same time, small price changes are common. About 44 percent of price changes excluding sales are smaller than 5 percent in absolute value, 25 percent are smaller than 2.5 percent, and 12 percent are smaller than 1 percent.

Nakamura and Steinsson (2008a) document the heterogeneity with respect to the size of price changes across Major Groups. The median absolute size of price changes ranges from about 6 percent in Utilities and Vehicle Fuel to about 30 percent in Apparel, Unprocessed Food, and Processed Food. Note that price changes in all Major Groups are on average large relative to inflation. After excluding sale-related price changes,

the median absolute size of price changes in Apparel, Unprocessed Food, and Processed Food falls, but remains large (12–14 percent) relative to inflation.

Nakamura and Steinsson (2008a) study the data underlying the producer price index in the United States. Producer goods can be divided into three groups based on stages of processing: finished goods, intermediate goods, and crude materials. Nakamura and Steinsson estimate the median price duration of finished goods between 1998 and 2005 to be 8.7 months. Nakamura and Steinsson's estimate of the median price duration of intermediate goods is 7 months. Crude materials have almost perfectly flexible prices, the estimated median price duration being 0.2 months. These estimates are based on data excluding forced item substitutions. The frequency of forced item substitutions varies across Producer Price Index (PPI) Major Groups from 0 percent in Farm Products to 16.6 percent in Transportation Goods. Sales are very rare in the PPI data.¹⁴ As in the case of consumer prices, there is a large amount of heterogeneity across PPI Major Groups with respect to the frequency of price changes. Nakamura and Steinsson estimate the median absolute size of price changes for finished producer goods to be 7.7 percent.

Nakamura and Steinsson (2008a) note that interpreting PPI data is difficult. The BLS collects this PPI data by means of a survey of firms.¹⁵ This method gives rise to the concern that firms report list prices rather than transaction prices. The BLS attempts to address this concern by requesting the prices of actual shipments of goods. Furthermore, many prices collected in order to compile the PPI are likely to be part of explicit or implicit contracts. This raises the possibility that Barro's (1977) criticism applies to some degree, and many observed producer prices do not map easily into allocations. An observed price in the data differs from the actual price faced by the buyer, and this actual price is unobserved by researchers. A related point made by Nakamura and Steinsson is that in wholesale markets suppliers may vary quality margins, such as delivery lags, instead of changing the price. We suspect that a version of Barro's criticism applies also to consumer prices in some sectors. Repeated interactions arise also in some sectors of the retail economy. For example, long-term relationships could play a role in explaining why the ELI with

the lowest frequency of price changes turns out to be Legal Services. Furthermore, suppliers may also vary quality margins in retail markets. For example, consumers sometimes must wait in order to purchase a good at a published price. The unobserved cost of waiting affects the shadow price of the good to the consumer.¹⁶ For these reasons, it may be wise to think of the available estimates of the frequency of consumer price changes as lower bounds.

A number of recent papers analyze scanner data from supermarket chains. Scanner datasets include data on quantities in addition to data on prices, and sometimes scanner datasets also include data on costs. Furthermore, scanner data typically are collected weekly. This means that fewer price changes are missed compared with monthly data from national statistical authorities. On the other hand, scanner data are not as representative as data from national statistical authorities. Eichenbaum, Jaimovich, and Rebelo (2008) study a new weekly scanner dataset from a major U.S. retailer that contains information on prices, quantities, and costs for over 1,000 stores. They find that prices and costs fluctuate around reference values, which tend to remain constant for extended periods of time. Prices have an average duration of three weeks. Reference prices have an average duration of about one year, where the reference price of a given item is defined as the most common price of that item during a given time interval. It is possible that variation in reference prices captures most of the variation in prices that matters for macroeconomics; that is, most of the variation in prices reflecting the response to macroeconomic shocks. Deviations from reference prices tend to be transient, whereas macroeconomic shocks tend to be persistent.

Micro Price Data in the Euro Area

The analysis of microeconomic price data in the euro area has been carried out in a project called the Inflation Persistence Network (IPN). The IPN has been a joint undertaking of the European Central Bank and national central banks of the euro area member countries.

Dhyne et al. (2005) summarize the findings of the IPN concerning microeconomic data on consumer prices in the euro area.¹⁷ Dhyne et al. analyze a sample of 50 goods and services common across the euro area member countries. The data are monthly and run from January 1996 to

January 2001. The sample of 50 products is representative in the sense that computing the mean frequency of price changes for the same 50 products in the dataset of Bils and Klenow (2004) yields a number close to the mean frequency of price changes in the entire Bils and Klenow dataset.

Dhyne et al. (2005) find considerable heterogeneity in the frequency of price changes in the euro area, just like in the United States. The frequency of price changes in the euro area ranges from 5.6 percent in Services to 28.3 percent in Unprocessed Food and 78 percent in Oil Products. In most sectors, the frequency of price changes in the euro area (and its individual member countries) is lower than in the same sector in the United States. These comparisons are shown in table 6.2.¹⁸

Dhyne et al. (2005) estimate the median price duration of consumer prices in the euro area to be 10.6 months.¹⁹ Forced item substitutions are included in the data analyzed. Unfortunately, different national statistical institutes within the euro area treat sale-related price changes differently. In some countries the reported price excludes the discount even when a sale is known to be in place.²⁰ Therefore, it is possible that the concept underlying the statement “the median price duration equals 10.6 months in the euro area” is closer to regular prices than to posted prices. Having said that, it does appear that consumer prices are changed less frequently in the euro area than in the United States. An interesting research question is what explains this apparent difference in the degree of price rigidity between the two economies.

Dhyne et al. (2005) estimate that the mean size of a price increase in the euro area is 8 percent and the mean size of a price decrease in the euro area is 10 percent. The euro-area estimates of the size of price changes are somewhat smaller than the U.S. estimates.²¹ At the same time, price changes in the euro area are on average large relative to inflation, just like in the United States.

Vermeulen et al. (2007) provide a comparative analysis of the data underlying the producer price indexes in Belgium, France, Germany, Italy, Portugal, and Spain. The findings are similar to the findings of Nakamura and Steinsson (2008a) concerning producer prices in the United States. In both the U.S. and euro-area economies, there is a great deal of heterogeneity in the frequency of price changes at the wholesale level. Furthermore, producer price changes are large relative to inflation. Vermeulen

Table 6.2
Average Frequency of Price Changes by Product Type, Euro Area versus the United States, from Dhyne et al. (2005)

	Unprocessed Food	Processed Food	Energy (Oil Products)	Non-Energy Industrial Goods	Services
Austria	37.5	15.5	72.3	8.4	7.1
Belgium	31.5	9.1	81.6	5.9	3.0
Germany	25.2	8.9	91.4	5.4	4.3
Spain	50.9	17.7	n.a.	6.1	4.6
Finland	52.7	12.8	89.3	18.1	11.6
France	24.7	20.3	76.9	18.0	7.4
Italy	19.3	9.4	61.6	5.8	4.6
Luxembourg	54.6	10.5	73.9	14.5	4.8
Netherlands	30.8	17.3	72.6	14.2	7.9
Portugal	55.3	24.5	15.9	14.3	13.6
Euro Area	28.3	13.7	78.0	9.2	5.6
U.S.	47.7	27.1	74.1	22.4	15.0

Note: The estimates for the United States are based on Bills and Klenow (2004).

Table 6.3
 Factors Affecting the Frequency of Producer Price Changes, from Vermeulen et al. (2007)

	Belgium	France	Germany	Italy	Portugal	Spain
Share of labor in costs	Yes	Yes	Yes	Yes	n.a.	Yes
Share of intermediate inputs in costs						
energy	Yes	n.a.	Yes	Yes	Yes	Yes
nonenergy	Yes	n.a.	Yes	Yes	Yes	Yes
Inflation	n.a.	Yes	Yes	n.a.	n.a.	Yes
Competition	Yes	Yes	n.a.	n.a.	n.a.	Yes
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes
Attractive prices	n.a.	n.a.	Yes	Yes	n.a.	Yes
Regulated prices	n.a.	n.a.	n.a.	n.a.	n.a.	Yes
Changes in VAT rates	n.a.	Yes	Yes	n.a.	n.a.	Yes

Note: “Yes” denotes that the factor has an impact on the frequency of price changes; “n.a.” indicates that the impact of the factor on the frequency of price changes has not been analyzed.

et al. investigate what explains the differences in the frequency of price changes across sectors. Table 6.3 summarizes these findings concerning what factors matter for the frequency of price changes.²² We would like to highlight a few factors. The cost structure matters. Vermeulen et al. (2007) find that firms with a higher labor share in total costs tend to change prices less frequently. Firms with a higher share of intermediate inputs, both energy and nonenergy, tend to change prices more frequently. Furthermore, the degree of competition matters. Vermeulen et al. (2007) find that firms operating in more competitive sectors tend to change prices more frequently. Vermeulen et al.'s findings concerning the impact of energy inputs and the degree of competition accord with the evidence presented by Bils and Klenow (2004) for the United States. Álvarez and Hernando (2007) investigate what explains the differences in the frequency of price changes across sectors using the IPN survey evidence on price setting behavior. Álvarez and Hernando find that prices tend to be more rigid in countries in which product markets are more regulated, as proxied by an index of product market regulation. More intense product market regulation may be one reason why prices change less frequently in the euro area than in the United States.

The IPN provides evidence concerning the degree of synchronization of price changes. Dhyne et al. compute an index of synchronization due to Fisher and Konieczny (2000). Dhyne et al. (2005) find that the degree of synchronization of price changes is low except for energy prices. The degree of synchronization typically increases as more narrow product categories are considered.

Cross-Country Evidence

Bils and Klenow (2004) and the IPN have inspired research on the frequency of price changes in many countries. Álvarez (2008) provides a survey of the recent cross-country evidence based on micro price data. It is useful to think of each study of micro price data as providing a pair of data points, where one data point is the monthly frequency of price changes in a given country in a given period, and the other data point is the average monthly rate of inflation in that country in that period. Using the studies listed in Álvarez (2008) and adding a few studies, we obtain 33 observations. Regressing the monthly frequency of price changes on




Figure 6.2
Frequency of Price Changes versus Inflation in a Cross-Section of Countries

the average monthly rate of inflation yields a statistically significant positive relationship (see figure 6.2). Thus the recent cross-country evidence confirms Taylor's (1999) conclusion that the frequency of price changes depends positively on the average rate of inflation.

The studies listed in Álvarez (2008) include countries with diverse experiences. According to one study, the frequency of price changes in Sierra Leone was 51 percent between 1999 and 2003, when the average rate of inflation was about 1.5 percent per month. According to another study, the frequency of price changes in Italy was 9.5 percent between 1996 and 2003, when the average rate of inflation was about 0.3 percent per month. Dotsey, King, and Wolman (1999) show that a general equilibrium menu cost model predicts a positive relationship between the frequency of price changes and steady-state inflation. Golosov and Lucas (2007) calibrate a menu cost model to match some features of the BLS data, including the frequency of price changes and the average rate of inflation. Golosov and Lucas show that the menu cost model calibrated to the U.S. data fits the frequency of price changes and the average rate

of inflation during two episodes of high inflation in Israel studied in Lach and Tsiddon (1992), an episode of low inflation in Israel studied in Baharad and Eden (2004), an episode of high inflation in Poland studied in Konieczny and Skrzypacz (2005), and a low-inflation episode and a high-inflation episode in Mexico studied in Gagnon (2007).

Survey Evidence

The IPN has surveyed euro-area firms asking about their price setting behavior. Fabiani et al. (2005) analyze the results of the survey.²³ The survey question we would like to focus on is: “If there are reasons for changing the price of your main product, which of the following factors may well prevent an immediate price adjustment?” This survey question was followed by a number of possible answers, each answer expressing in simple terms one economic theory. The respondents could indicate their degree of agreement with each economic theory. The responses indicate that firms refrain from changing prices mainly because of explicit and implicit contracts with customers. The physical costs of changing prices (also called menu costs) are among the reasons for price rigidity least favored by firms, as are the costs of obtaining information.

This recent euro-area evidence from surveys of firms matches well with the same kind of evidence collected earlier in the United States. Blinder, Canetti, Lebow, and Rudd (1998) report that when managers of U.S. manufacturing firms were asked why they do not change prices more often than they do, the most common answer was that doing so would “antagonize” customers. This response confirms that business firms view recurring customer relationships as important. Zbaracki et al. (2004) analyze in detail the pricing behavior of a large manufacturing U.S. company. They find that the most important costs of changing prices are the “customer costs,” meaning the costs of communicating and negotiating with customers. The customer costs arise in the presence of long-term relationships. Zbaracki et al. find that the customer costs are followed in terms of importance by the “managerial costs,” meaning the costs of information gathering and decisionmaking. Zbaracki et al. find that the customer costs are more than 20 times higher compared to the menu costs, and the managerial costs are more than six times higher than the menu costs.

The survey evidence, though it comes with known problems, may well indicate the main reason why prices remain constant for extended periods of time in many sectors of the economy. Firms keep prices unchanged because they worry about jeopardizing future interactions with customers. We find it remarkable that the survey evidence matches so well with the classic analysis of “customer markets” by Okun (1981), who wrote:

The firm recognizes its ability to discourage customers from shopping elsewhere by convincing them of the continuity of the firm’s policy on pricing, services, and the like. . . . Customers are attracted by continuity because it helps to minimize shopping costs. They know the terms of the previous supplier’s offer without shopping if they can count on its continuance, but they must shop to determine the offers of unfamiliar sellers. That information is available, but it can be obtained only at a cost. . . . [C]ustomer markets share the characteristics of career labor markets. Both feature search costs, information costs, and bilateral-monopoly surpluses associated with established relations (Okun 1981, pp. 141–142).

Nakamura and Steinsson (2009) is a recent paper that makes further progress modeling the idea of customer markets.

Two more comments concerning the survey evidence are in order. About 75 percent of the firms surveyed by Fabiani et al. (2005) sell their output mainly to other firms.²⁴ Similarly, Blinder, Canetti, Lebow, and Rudd (1998) and Zbaracki et al. (2004) study manufacturing firms operating in wholesale markets. Economists familiar with these studies sometimes express the view that firms operating in retail markets view recurrent interactions as unimportant compared with firms operating in wholesale markets. We think that recurrent interactions matter in some wholesale sectors and in some retail sectors. Holding the prior widely accepted modeling assumption that recurrent interactions matter only for producer prices appears unjustified. At the same time, it may be that considerations other than long-term relationships are the main cause of price rigidity in some sectors of the retail economy.

We also think that the survey evidence does not speak against all models of price setting with imperfect information. In the rational inattention model of Maćkowiak and Wiederholt (2009), an individual firm is very well informed about its environment. The firm’s private marginal value of information is low. Therefore, it is to be expected when conducting a survey that the firm will respond that it is very well informed.

Let us conclude section 2 with a summary of the main points. The new micro price data from the United States and the euro area share a number of characteristics. Both in the United States and in the euro area, there is a lot of heterogeneity in the frequency of price changes, prices remain constant for extended periods of time in many sectors of the economy, prices change on average by large amounts relative to inflation, and the survey evidence indicates that firms perceive long-term relationships as the main reason why prices remain constant for some time. Consumer prices appear to change less frequently in the euro area than in the United States. In the U.S. economy the median consumer price lasts about four months; it lasts about seven months when price changes related to sales are excluded, and lasts about nine months when price changes related to both sales and forced item substitutions are excluded. The median consumer price lasts about 11 months in the euro-area economy. The frequency of price changes depends positively on the average rate of inflation in a cross-section of countries.

3. Rejecting the Null Hypothesis

In this section, we discuss the recent literature that compares standard models of price setting used in macroeconomics to the new micro data. The main theme of this section is that the new micro data support the basic premise underlying the New Keynesian or New Neoclassical Synthesis perspective: the prices of many goods and services remain constant for extended periods of time. At the same time, standard models of price setting used in macroeconomics are so simple that each of the models is bound to be rejected. Each of the models is at odds with some aspect of the detailed micro data that we now have.²⁵

Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008a) document a number of features of the BLS micro price data, in addition to the frequency and the size of price changes. Both papers compare the main features of the micro data to standard models of price setting used in macroeconomics. Klenow and Kryvtsov (2008) emphasize the following features of the BLS price data, in addition to the frequency and the size of price changes. Price durations for a given product are variable. Hazard rates for a given product are approximately flat. The size of

price changes for a given product is unrelated to the time since the previous change. The intensive margin dominates the variance of inflation.²⁶ Nakamura and Steinsson (2008a) stress some features of the data that do not receive emphasis from Klenow and Kryvtsov. One feature of the data noted by Nakamura and Steinsson that seems important is seasonality. The frequency of price changes is highly seasonal. It is highest in the first quarter and lowest in the fourth quarter.

Klenow and Kryvtsov (2008) note that the menu cost model of Golosov and Lucas (2007) fails to generate enough small price changes. The reason is that the model has a single, fairly large menu cost. The Golosov-Lucas model needs a fairly large menu cost in order to match the large average absolute size of price changes. Furthermore, only if the elasticity of substitution between products is set equal to a number like 2 can the Golosov-Lucas model yield flat hazard rates and show approximately no relationship between the size of price changes and the time since the previous price change. Interestingly, the intensive margin drives inflation movements in the Golosov-Lucas model. The fraction of price changes is stable, because prices change mostly in response to large idiosyncratic shocks rather than small macroeconomic shocks. Since the Golosov-Lucas model also produces small real effects of nominal shocks, we now know that real effects of nominal shocks can be small even if the extensive margin is unimportant.

Klenow and Kryvtsov (2008) study a version of the Calvo model with idiosyncratic shocks. They find that the model predicts a larger absolute size of price changes for older prices, in contradiction to the micro data. Klenow and Kryvtsov also study a Taylor model with multiple sectors. The Taylor model fails to produce variable price durations and flat hazard rates. We would like to add that it is an open question whether either a menu cost model or a Calvo model can give a plausible account of the seasonality in the micro data.

A number of recent papers analyze scanner data and compare the data to standard models of price setting used in macroeconomics. Eichenbaum, Jaimovich, and Rebelo (2008) study a new weekly scanner dataset from a major U.S. retailer that contains information on prices, quantities, and costs for over 1,000 stores. They find that prices are more volatile than marginal costs. Furthermore, the probability of a price change

increases in the deviation of the markup from its mean. Prices typically adjust when the markup deviates by more than 20 percent from its mean. Eichenbaum, Jaimovich, and Rebelo argue that neither the Calvo model nor the menu cost model can match the data. The Calvo model is inconsistent with the observed state dependence of prices. The menu cost model predicts that prices should be less volatile than marginal costs. Campbell and Eden (2007) analyze a different scanner dataset, but also find that prices are state dependent. The probability of a price change is highest when a store's price differs substantially from the average of other stores' prices.

Macroeconomists recently gained access to new, more detailed micro price data. We confront the data with familiar and tractable models of price setting. These models are such crude approximations to the real world and the microeconomic data are so detailed that, not surprisingly, the models fail. The microeconomic data are not only rich, in the sense that a lot goes on for a given product category, but also reveal a lot of heterogeneity across products. There is heterogeneity in terms of the frequency and size of price changes, but also in terms of how frequently sales take place, what form sales take, and how frequently forced item substitutions occur. The wholesale economy seems to differ from the retail economy. The heterogeneity in the microeconomic data makes us skeptical that economists should aim at developing "the model" of price rigidity. Different models of price rigidity may be necessary for different sectors of the wholesale and retail economy. Furthermore, it is an open question whether models of price rigidity that fit the micro price data well will imply the kinds of impulse responses to macroeconomic shocks that we see in aggregate data.

4. Mapping Micro Price Data into Macroeconomic Models

One reason why macroeconomists are interested in microeconomic price data is that in familiar and tractable macroeconomic models, the frequency of price changes maps easily into impulse responses of prices and quantities to shocks. In the Calvo model, a lower frequency of price changes implies larger and more persistent real effects of nominal shocks. The same is typically true in a menu cost model. In this section, we argue that there is

no simple mapping from the frequency of price changes in microeconomic data to the impulse responses of prices and quantities to shocks.

This point is not new. We know there are models with physically rigid prices in which nominal shocks have no real effects, and we know there are models with perfectly flexible prices in which nominal shocks can have large and persistent real effects. In the menu cost model of Caplin and Spulber (1987), nominal shocks have no real effects, despite the fact that individual prices are adjusted infrequently. Sims (1998) provides a more recent, perhaps less familiar example. Sims considers a model with nominal wage contracts that do not represent open-ended commitments to provide labor at a given wage rate. This is in contrast to most models of nominal rigidity, which assume that workers supply as much labor as firms demand at a given wage rate, and firms produce as much output as consumers demand at a given price. In Sims's model, individual workers either are employed or unemployed. If their contract specifies a nominal wage that turns out to be low in real terms, they lose at the expense of firms. But since workers own firms, an "expansionary" nominal shock in Sims's model only "redistributes" wealth from workers to firms and back to workers. Output is unaffected.

At the opposite end of the spectrum of models, Woodford (2002), Mankiw and Reis (2002), Reis (2006), and Maćkowiak and Wiederholt (2009) develop the idea of Phelps (1969) and Lucas (1972) that the real effects of nominal shocks are due to imperfect information. In these recent models, nominal shocks can have large and persistent real effects, despite the fact that individual prices are adjusted frequently. In the rest of this section, we would like to use our own separate research to illustrate anew, in two different ways, the point that mapping micro price data into macroeconomic models is complicated.

Impulse Responses of Sectoral Price Indexes

Maćkowiak, Moench, and Wiederholt (2008) estimate impulse responses of sectoral price indexes to common shocks and sector-specific shocks. If the frequency of price changes were decisive for impulse responses of prices to shocks, one would expect sectoral price indexes to respond with roughly equal speed to both kinds of shocks. Maćkowiak, Moench, and Wiederholt estimate a Bayesian unobservable index model using monthly

sectoral consumer price indexes from the U.S. economy for the period from 1985 until 2005. The unobservable index model is motivated by the optimal pricing equation in the rational inattention model of Maćkowiak and Wiederholt (2009). Maćkowiak, Moench, and Wiederholt find that most of the variation in sectoral price indexes, about 85 percent, is caused by sector-specific shocks. Sectoral price indexes are very volatile relative to the aggregate price level. Figure 6.3 shows the cross-section of the impulse responses of sectoral price indexes to sector-specific shocks. Note that this is a posterior distribution taking into account both parameter uncertainty and variation across sectors.²⁷ It is apparent that sectoral price indexes respond quickly to sector-specific shocks. Essentially 100 percent of the long-run response occurs within one month. In this sense, sectoral price indexes are not sticky at all. Figure 6.3 is reminiscent of the one-to-one impulse response of the aggregate price level to money in the Caplin-Spulber model.

Figure 6.4 shows the cross-section of the impulse responses of sectoral price indexes to an aggregate shock. Sectoral price indexes respond slowly to aggregate shocks. About 15 percent of the long-run response occurs within one month. In this sense sectoral price indexes are sticky. These




Figure 6.3
Impulse Responses of Sectoral Prices Indexes to Sector-Specific Shocks
Source: Maćkowiak, Moench, and Wiederholt (2008)




Figure 6.4

Impulse Responses of Sectoral Prices Indexes to a Common Shock

Source: Maćkowiak, Moench, and Wiederholt (2008)

findings accord with the idea that the degree of price stickiness—defined as the speed of response of prices to disturbances—is conditional on the source of the disturbance. Prices in the same sector, which get assigned the same frequency of changes in microeconomic studies, respond quickly to sector-specific shocks, but only slowly to macroeconomic shocks.

The findings of Maćkowiak, Moench, and Wiederholt (2008) are complementary to the findings of Boivin, Giannoni, and Mihov (2007), who use different sectoral data from the United States, a different model, and a different estimation methodology to address a similar set of questions. Boivin, Giannoni, and Mihov also find that sectoral price indexes respond quickly to sector-specific shocks and slowly to aggregate shocks. Having said that, there is evidence that the frequency of price changes helps explain the speed of impulse responses of prices to macroeconomic shocks in a cross-section of sectors. The results of Boivin, Giannoni, and Mihov (2007) and Maćkowiak, Moench, and Wiederholt (2008) are consistent with the idea that sectors in which prices are adjusted more frequently tend to respond faster to macroeconomic shocks. In the data, there is a strong positive relationship between the frequency of price changes and the size of sector-specific shocks.²⁸ Therefore, it is reasonable

to conclude that prices in sectors facing greater sector-specific uncertainty tend to respond faster to macroeconomic shocks. This is consistent with the menu cost model, the imperfect information model of Reis (2006), and the rational inattention model of Maćkowiak and Wiederholt (2009).

In a recent paper, McCallum and Smets (2008) estimate the effects of monetary policy shocks in the euro area using a large dataset of area-wide, country-specific, and sector-specific time series for the period from 1987 to 2005. McCallum and Smets (2008) employ the factor-augmented vector autoregression (VAR) methodology proposed by Bernanke, Boivin, and Elias (2005). We added five sectoral consumer price indexes from the three largest euro-area countries (Germany, France, and Italy) to the dataset of McCallum and Smets (2008). The five sectors are unprocessed food, processed food, nonenergy industrial goods, energy, and services.²⁹ We computed the impulse responses of the sectoral price indexes to a monetary policy shock. The impulse responses are plotted in figure 6.5, along with the impulse response of the euro-area Harmonized Index of Consumer Prices to the same shock. The sectoral price indexes fall slowly after a contractionary monetary policy shock, as does the aggregate price level. We assigned to each sectoral price index a frequency of price changes, based on the results of the IPN. Figure 6.6 shows a statistically significant, negative relationship between the response to a monetary policy shock after eight quarters in a given sector and the frequency of price changes in that sector. Sectors in which prices change frequently tend to respond more strongly to a monetary policy shock. The findings based on the model of McCallum and Smets (2008) confirm that the frequency of price changes helps explain the speed of impulse responses of prices to macroeconomic shocks in a cross-section of sectors.

DSGE Models

The second example involves the DSGE model of Smets and Wouters (2003). This model has a marginal likelihood that is comparable to that of an unconstrained, low-order VAR. The model's price-setting behavior and its wage-setting behavior are a mixture of the Calvo mechanism and a backward-looking component (indexation). Sahuc and Smets (2008) estimate a variant of the Smets-Wouters model using quarterly aggregate data from 1985 to 2004 for, alternatively, the United States and the




Figure 6.5

Impulse Responses to a Monetary Policy Shock, based on McCallum and Smets (2008)

euro area. The slope of the New Keynesian Phillips curve is estimated to be very low and somewhat lower in the euro area than in the United States (0.008 versus 0.012). Since the slope of the New Keynesian Phillips curve is inversely related to the Calvo parameter, the Smets-Wouters model reflects the notion that prices change less frequently in the euro area compared to the United States. However, this difference is only weakly demonstrated in the model: the difference is not statistically significant and not large economically.³⁰ Furthermore, without real rigidity of some form, the estimated value of the Calvo parameter is very high (0.91 in the euro area and 0.89 in the United States). Smets and Wouters




Figure 6.6

Impulse Response to a Monetary Policy Shock versus Frequency of Price Changes, based on McCallum and Smets (2008)

(2007) replace the Dixit-Stiglitz aggregator with the Kimball aggregator in estimation using U.S. data.³¹ This increases the degree of real rigidity in the model, as the price elasticity of demand becomes an increasing function of the firm's relative price. Therefore, prices in the Smets-Wouters model respond by smaller amounts to shocks, for a given frequency of price changes. Smets and Wouters obtain a much smaller estimate of the Calvo parameter, about two-thirds. A value for the Calvo parameter of about two-thirds implies that price contracts last three quarters on average. The Kimball aggregator plus the Calvo parameter equal to about two-thirds yield roughly the same fit to U.S. macroeconomic data as the Dixit-Stiglitz aggregator plus the Calvo parameter equal to about 0.9. That the frequency of price changes in the microeconomic data does not map easily to impulse responses in the Smets-Wouters DSGE model should not come as a surprise. The Smets-Wouters model is intended to capture the impulse responses of the aggregate price level and other macroeconomic variables to macroeconomic shocks. The price-setting mechanism embedded in the Smets-Wouters model is a stand-in for "how

prices respond to macroeconomic shocks,” not for “how prices are being set at the micro level.”

Note also that in the Smets-Wouters model, because of indexation, all prices change every quarter. This frequency of price changes makes it difficult to map the Calvo parameter in the Smets-Wouters model to the micro price data. It is possible that a model with the “rule-of-thumb” behavior, as in Galí and Gertler (1999), instead of indexation would fit the macroeconomic data equally well. In the Galí-Gertler model, some firms keep prices unchanged each period. The importance of the backward-looking elements for matching the macroeconomic data accords well with the idea that some form of imperfect information about macroeconomic shocks matters for aggregate dynamics. More generally, the fact that prices change does not imply that prices reflect perfectly all available information.

The likelihood of a DSGE model, such as the Smets-Wouters model, peaks in a region of the parameter space, which implies that prices and quantities move slowly in response to most macroeconomic shocks. The Smets-Wouters DSGE model is capable of matching this pattern in the macroeconomic data via a combination of some Calvo rigidity and some form of sufficient real rigidity. How frequently prices change is not decisive. What is decisive is how price-setting behavior interacts with other features of the economy to produce sluggish impulse responses to macroeconomic shocks.³² We discuss modeling different kinds of real rigidity in section 5.

5. The Road Ahead

In an ideal world, macroeconomists would set up, solve, and understand a DSGE model with many heterogeneous firms and households, in which buyers and sellers in different sectors interact differently, and which matches in detail both microeconomic data and macroeconomic data. Realistically, all we can hope for, at least for some time, is a model that matches the macroeconomic data well while telling a reasonable story that is broadly in line with microeconomic data. The recent DSGE models have made progress in matching macroeconomic data. In this section, we survey the recent literature on models of price setting. We search

for a story or a set of stories that would simultaneously imply persistent impulse responses to macroeconomic shocks, similar to the impulse responses in the Smets-Wouters model, and be broadly in line with micro price data.

Making a State-Dependent Model Behave Like an Exogenous Timing Model

After Bilal and Klenow (2004) and others had begun analyzing the BLS data, Golosov and Lucas (2007) constructed a menu cost model and calibrated the model to match some features of the BLS data. Golosov and Lucas found that the model needed large firm-specific productivity shocks in order to match the large average absolute size of price changes in the data. Furthermore, Golosov and Lucas found that their calibrated model predicted that nominal shocks would have small real effects.

There are two reasons why the real effects of nominal shocks are small in the calibrated Golosov-Lucas model. Price adjustment in the model is triggered almost always by idiosyncratic shocks, because idiosyncratic shocks in the model are much larger than aggregate shocks. But in the menu cost model, the identity of shocks does not matter. Instead, it is the size of the shocks that matters. Once a firm decides to adjust its price, the firm makes the price respond to all shocks, independent of which shock has triggered the price adjustment. Since idiosyncratic shocks trigger price adjustments fairly frequently, prices respond to nominal shocks fairly quickly.

Furthermore, there is a selection effect in a state-dependent model. In the absence of macroeconomic shocks, some firms increase prices and some firms decrease prices, in response to idiosyncratic shocks of sufficient magnitude. Suppose that an expansionary nominal shock arrives. Some firms that were going to decrease prices by a lot in the absence of the macroeconomic shock now keep prices constant. In an exogenous timing model, there are no such firms. Also, some firms that barely decided not to increase prices in the absence of the macroeconomic shock now increase prices by a sizable amount. In an exogenous timing model, there are no such firms either. Due to the selection effect, the price level increases by more in a state-dependent model than in an exogenous tim-

ing model. The selection effects can be quantitatively important. The analysis of Caballero and Engel (2007) suggests that the price level's degree of flexibility in the state dependent model is about three times larger compared with an exogenous timing model, with the same average frequency of price changes.

Midrigan (2008) shows how one can dampen the selection effect. He notes in scanner data that prices of goods sold by a particular retailer, especially those in narrow product categories, tend to be adjusted simultaneously. Midrigan uses this observation to motivate a model in which a two-product firm faces a fixed cost of changing one of its prices, but, conditional on paying this cost, a zero marginal cost for resetting the other price. The firm's profit function is affected by nominal shocks and idiosyncratic shocks drawn from a density with fat tails. Midrigan's model can match simultaneously the observation that the average absolute size of price changes is large and the observation that many price changes are small. This result is in contrast to the implications of the Golosov-Lucas model. Also in contrast to the Golosov-Lucas model, Midrigan's model predicts real effects of nominal shocks roughly similar in size to an exogenous timing model. Compared to the Golosov-Lucas model, in Midrigan's model fewer firms that were going to decrease prices in the absence of an expansionary nominal shock do keep prices constant when the shock occurs. Similarly, fewer firms that were going to keep prices unchanged in the absence of an expansionary nominal shock do increase prices when the shock occurs. The reason is that the density of price changes is leptokurtic. There is a small mass of firms near the points where the decision whether to change prices or not is made.³³

Modeling the Kind of Real Rigidity that Works

Modelers who introduce real rigidity into a DSGE model usually do so at the level of individual firms. For example, with the Kimball aggregator the price elasticity of demand faced by an individual firm is increasing along with the firm's relative price. As another example, with firm-specific inputs an individual firm's marginal cost function is increasing in the firm's output. In the presence of a real rigidity of this kind, firms find it optimal to change prices by small amounts in response to all shocks.³⁴

This kind of real rigidity is not easy to reconcile with micro price data, because in these data price changes are large relative to inflation. Furthermore, sectoral price indexes respond quickly to sector-specific shocks. Conditional on having changed prices, firms do not seem to be worried about changing prices a lot in response to firm-specific and sectoral shocks. There appears to be little evidence for significant real rigidity at the level of individual firms and sectors. Burnstein and Hellwig (2007), Dotsey and King (2005), and Klenow and Willis (2006) investigate the effects of firm-level real rigidity in menu cost models. All these studies find that the menu cost model with significant firm-level real rigidity needs a combination of very large idiosyncratic shocks and very large menu costs in order to match the large size of price changes in the data.

The kind of real rigidity that seems promising for macroeconomic models is one arising conditionally on aggregate shocks. What we have in mind are features of the economy such that firms find it optimal to change prices by large amounts in response to firm-specific and sectoral shocks and, at the same time, firms find it optimal to change prices by small amounts in response to aggregate shocks. Nakamura and Steinsson (2008b) develop a multisector menu cost model introducing real rigidity via intermediate inputs, as in Basu (1995). The degree of monetary non-neutrality generated by the model with intermediate inputs is roughly triple that of the model without intermediate inputs. Firms that change prices soon after a nominal shock adjust less than they otherwise would, because the prices of many of their inputs have not yet responded to the shock. At the same time, introducing real rigidity via intermediate inputs does not dampen the size of price changes in response to idiosyncratic shocks. Nakamura and Steinsson (2008b) also find that their multisector menu cost model generates larger real effects of nominal shocks than a single-sector menu cost calibrated to the mean frequency of price changes of all firms. The interaction between heterogeneity and aggregate-level real rigidity seems worth exploring further in future research. Furthermore, recall from section 2 that the IPN evidence on determinants of the frequency of price changes points to the labor share as one determinant of the frequency of price changes. The slow responsiveness of wages may be another source of aggregate-level real rigidity that interacts with price-

setting behavior. Kryvtsov and Midrigan (2009) is a recent paper assessing the degree of real rigidity that exists, where real rigidity may take the form of sluggish wages.

Rational Inattention and Sticky Information

Lucas (1972) formalized the idea that the real effects of nominal shocks are due to imperfect information. Lucas assumed that agents observe the current state of monetary policy with a delay. His model has been criticized on the grounds that information concerning monetary policy is published with little delay. However, Sims (2003) points out that, if agents cannot process all available information perfectly, there is a difference between publicly available information and the information actually reflected in agents' decisions.

Maćkowiak and Wiederholt (2009) develop a model in which information concerning the current state of monetary policy is publicly available, and it is optimal for agents to pay little attention to this information. In the model, price-setting firms decide what to pay attention to in their decisionmaking. Firms' inability to process all available information is modeled as a constraint on information flow, as in Sims (2003). Firms face a trade-off between paying attention to aggregate conditions and paying attention to idiosyncratic conditions. Impulse responses of prices to shocks are sticky—dampened and delayed relative to the impulse responses under perfect information. When idiosyncratic conditions are more variable or more important than aggregate conditions, firms pay more attention to idiosyncratic conditions than to aggregate conditions. Prices then respond strongly and quickly to idiosyncratic shocks, and, at the same time, prices respond weakly and slowly to aggregate shocks. In addition, there are feedback effects, because firms track endogenous aggregate variables (the price level and real aggregate demand). When other firms pay limited attention to aggregate conditions, the price level responds less to a nominal shock than under the assumption of firms having perfect information. If prices are strategic complements to the information a firm possesses, this implies that each firm has even less incentive to pay attention to aggregate conditions. The price level responds even less to a nominal shock, and so on.

Maćkowiak and Wiederholt (2009) find, like Golosov and Lucas (2007) and other studies working with menu cost models, that in order to match the large average absolute size of price changes in the microeconomic data, idiosyncratic volatility in the rational inattention model has to be one order of magnitude larger than aggregate volatility. This result implies that firms allocate almost all their attention to idiosyncratic conditions. Therefore, prices respond strongly and quickly to idiosyncratic shocks, and prices respond only weakly and slowly to nominal shocks. Nominal shocks have strong and persistent real effects.

Mankiw and Reis (2002) develop a different model in which information disseminates slowly. They assume that in every period a fraction of firms obtains perfect information concerning all current and past disturbances, while all other firms continue to set prices based on old information. Reis (2006) shows that a model with a fixed cost of obtaining perfect information can rationalize this kind of slow information diffusion. Note that in Mankiw and Reis (2002) and Reis (2006) prices respond with equal speed to all disturbances.

Modeling Rigid Prices without a Menu Cost

Models in which prices remain constant for some time due to a menu cost are useful exploratory devices. The same can be said about models in which firms face different menu costs for different products, firms face different menu costs for changing regular prices and sale prices, or firms face menu costs that vary randomly with the variation being orthogonal to the firms' economic environment. Ultimately, one would like to know why fixed costs of changing prices arise and how such costs depend on firms' economic environment. The idea of customer markets seems promising in this regard. Nakamura and Steinsson (2008c) model the idea of customer markets by observing that if consumers form habits regarding individual goods, then firms face a time-inconsistency problem. Consumers' habits imply that demand is forward-looking. Low prices in the future help attract consumers at present. Therefore, firms want to promise low prices in the future. But when the future arrives, firms have an incentive to exploit consumers' habits and raise prices. Nakamura and Steinsson show that implicit contracts involving price rigidity can be sustained as equilibria in the infinitely repeated game played by a firm and its customers.³⁵

6. Concluding Remarks

It is possible that several DSGE models will emerge consistent with the impulse responses in macroeconomic data. Each model may have different implications for some feature of micro data. Therefore, using micro data may be helpful for discriminating between the models. Furthermore, each model may have different implications for welfare consequences of macroeconomic fluctuations. For example, in a rational inattention model, the welfare costs of day-to-day macroeconomic fluctuations tend to be modest, at least so long as the marginal value of information is small from the viewpoint of individual agents in the model. Welfare costs of day-to-day macroeconomic fluctuations are much larger in the Calvo model. The reason is that, on average, agents in the Calvo model are further away from their first-best decisions compared to agents in the rational inattention model. In sticky information models, welfare costs of day-to-day macroeconomic fluctuations tend to be large, because at any time some agents are quite poorly informed about their economic environment. On the other hand, in rational inattention models an increase in aggregate volatility can be very costly, because tracking the aggregate economy becomes increasingly difficult as the macroeconomic environment becomes less predictable.

■ *We thank Robert King, Virgiliu Midrigan, Elizabeth Murry, Emi Nakamura, and the conference participants for comments. Gianluca Cafi so provided excellent research assistance. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the European Central Bank. E-mail: bartosz.mackowiak@ecb.int and frank.smets@ecb.int.*

Notes

1. The closing sentences of Taylor (1999, p. 1044) call for more research: “Understanding these [staggered price- and wage-setting] models more thoroughly takes one well beyond macroeconomics into the heart of the price discovery and adjustment process in competitive and imperfectly competitive markets. Further research on the empirical robustness and microeconomic accuracy of staggered contracts models is thus both interesting and practically important.”

2. Similarly, whether prices are changed on average every four months or every 11 months, nominal rigidity by itself does not suffice to explain why the impulse response of real aggregate output to monetary policy shocks lasts two to three years in structural VAR models. See, for example, Christiano, Eichenbaum, and Evans (1999); Kim (1999); and Leeper, Sims, and Zha (1996).
3. The concept of real rigidity is due to Ball and Romer (1990).
4. The BLS revised the ELI structure in 1998. Prior to 1998, there were about 360 ELIs. Since 1998, there are about 270 ELIs.
5. Figure 6.1 is based on table 17 in the supplement to Nakamura and Steinsson (2008a) titled “More Facts About Prices,” available at <http://www.columbia.edu/~en2198.papers.html>.
6. Table 6.1 is based on table II in Nakamura and Steinsson (2008a).
7. Bils and Klenow (2004) report the median price duration of 4.3 months. This estimate is based on summary statistics from the BLS for the period 1995 to 1997.
8. According to the BLS, a “sale” price is (a) temporarily lower than the “regular” price, (b) available to all consumers, and (c) usually identified by a sign or statement on the price tag.
9. Removing all sale-related price changes from the data raises the estimated mean price duration from 6.8 months to 8.6 months. The estimated median frequency of price changes falls from 27.3 percent to 13.9 percent. The estimated mean frequency of price changes falls from 36.2 percent to 29.9 percent.
10. The corresponding estimates of Nakamura and Steinsson (2008a) are almost identical. While Klenow and Kryvtsov (2008) report 7.2, 8.7, 9.3, and 10.6 months, Nakamura and Steinsson (2008a) find 7.5, 8.7, 9.6, and 11 months.
11. We do not mean to suggest that the findings of Klenow and Willis (2007) close the debate on whether sales and forced item substitutions respond to macroeconomic shocks. The findings depend to some degree on their modeling assumptions.
12. Nakamura and Steinsson (2008b) model price changes due to product turnover as time-dependent. They find that product turnover introduced in this way has a trivial effect on real effects of nominal shocks in a menu cost model.
13. Price increases are 2 to 3 percentage points smaller on average than price decreases. Both price increases and price decreases are common.
14. Unlike the CPI database, the PPI database does not record certain prices as sale prices. Nakamura and Steinsson (2008a) use sales filters to assess the importance of sales in the PPI data. The sales filters identify very few sales.
15. In contrast, BLS employees who collect prices used to compile the CPI record prices of goods actually on the shelf.
16. Think of waiting to get a haircut on a Saturday, and then getting it at the same published price as on a weekday.
17. See also Altissimo, Ehrmann, and Smets (2006) for a survey of the IPN evidence concerning price setting in the euro area.

18. Table 6.2 is based on table 3 in Dhyne et al. (2005). Note that they find little heterogeneity in the frequency of price changes by country.
19. Dhyne et al. (2005) estimate the mean price duration to be 13 months.
20. Three large European countries (Germany, Italy, and Spain) are among the countries in which sale prices are not recorded. In Germany sale prices are recorded only outside an explicit “seasonal sale.” See the technical appendixes 2 and 3 in Dhyne et al. (2005) regarding forced item substitutions and sales, respectively.
21. Klenow and Kryvtsov (2008) find that the mean absolute size of price changes is 14 percent and 11.3 percent excluding sale-related price changes.
22. Table 6.3 reproduces table 8 in Vermeulen et al. (2007).
23. See also Fabiani et al. (2007).
24. About two-thirds of the firms in the IPN survey indicate that long-term customers account for the bulk of their sales.
25. See Álvarez (2008) and Gaspar et al. (2007) for surveys of how the new micro data compare to standard models of price setting used in macroeconomics.
26. Inflation is the product of the fraction of products with price changes (the extensive margin) and the average size of those price changes (the intensive margin). Klenow and Kryvtsov (2008) estimate that the intensive margin accounts for 94 percent of inflation’s variance for posted prices and 91 percent for regular prices.
27. To derive this figure and the subsequent figure, sectoral inflation rates have been normalized to have variance unity. This normalization makes impulse responses comparable across sectors. Figures 6.3 and 6.4 are based on the results of Maćkowiak, Moench, and Wiederholt (2008).
28. A positive relationship between the frequency of price changes and the size of sector-specific shocks is consistent with the menu cost model.
29. The data on sectoral consumer price indexes were kindly provided by Benoît Mojon. See Altissimo, Mojon, and Zaffaroni (2007).
30. Galí and Gertler (1999) and Galí, Gertler, and López-Salido (2001) estimate the New Keynesian Phillips curve using generalized method of moments for the United States and the euro area. The findings are consistent with the notion that prices change less frequently in the euro area compared to the United States.
31. See Kimball (1995).
32. Sahuc and Smets (2008) compare impulse responses to a monetary policy shock in a model with a higher Calvo parameter like in the euro area with those in a model with a lower Calvo parameter like in the United States, keeping the other frictions constant. They find that the differences are minimal.
33. Gertler and Leahy (2006) also show how to dampen the selection effect in a menu cost model. In their model, price adjustments can only be triggered by idiosyncratic shocks.

34. Eichenbaum and Fisher (2007) estimate different variants of the Calvo model, finding that introducing the Kimball aggregator and firm-specific capital increases the estimated frequency of price adjustment. Similarly, Altig, Christiano, Eichenbaum, and Linde (2005) show that introducing firm-specific capital into a DSGE model with the Calvo mechanism increases the estimated frequency of price adjustment. However, introducing firm-specific capital also decreases the size of price changes. Altig, Christiano, Eichenbaum, and Linde estimate an elasticity of prices to real marginal costs of 0.04.

35. In a different strand of the literature, Rotemberg (2004) develops a model in which firms wishing to avoid customers' anger keep prices rigid under some circumstances when prices would change under standard assumptions.

References

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde. 2005. "Firm-Specific Capital, Nominal Rigidities and the Business Cycle." Working Paper 11034. Cambridge, MA: National Bureau of Economic Research.

Altissimo, Filippo, Michael Ehrmann, and Frank Smets. 2006. "Inflation Persistence and Price-Setting Behaviour in the Euro Area: A Summary of the IPN Evidence." European Central Bank Occasional Paper 46. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.edu/pub/pdf/scpops/ecbocp46.pdf>.

Altissimo, Filippo, Benoît Mojon, and Paolo Zaffaroni. 2007. "Fast Micro and Slow Macro: Can Aggregation Explain the Persistence of Inflation?" European Central Bank Working Paper 729. Frankfurt: European Central Bank. Available at <http://www.ecb.int/pub/pdf/scpwps/ecbwp729.pdf>.

Álvarez, Luis J. 2008. "What Do Micro Price Data Tell Us About the Validity of the New Keynesian Phillips Curve?" *Economics: The Open-Access, Open-Assessment E-Journal* 2(19). Available at <http://www.economics-ejournal.org/economics/journalarticles/2008-19>.

Álvarez, Luis J., and Ignacio Hernando. 2007. "Competition and Price Adjustment in the Euro Area." In *Pricing Decisions in the Euro Area: How Firms Set Prices and Why*, ed. Silvia Fabiani, Claire Loupias, Fernando Martins, and Roberto Sabbatini, 200–217. Oxford: Oxford University Press.

Baharad, Eyal, and Benjamin Eden. 2004. "Price Rigidity and Price Dispersion: Evidence from Micro Data." *Review of Economic Dynamics* 7(3): 613–641.

Ball, Lawrence, and David Romer. 1990. "Real Rigidities and the Non-Neutrality of Money." *Review of Economic Studies* 57(2): 183–203.

Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3(3): 305–316.

Basu, Susanto. 1995. "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare." *American Economic Review* 85(3): 512–531.

- Bernanke, Ben S., Jean Boivin, and Piotr Elias. 2005. "Measuring Monetary Policy: A Factor Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120(1): 869–902.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.
- Blinder, Alan S., Elie D. Canetti, David E. Lebow, and Jeremy B. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Boivin, Jean, Marc P. Giannoni, and Ilian Mihov. 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated U.S. Data." *American Economic Review* 99(1): 350–384.
- Burstein, Ariel T., and Christian Hellwig. 2007. "Prices and Market Shares in a Menu Cost Model." Working Paper 13455. Cambridge, MA: National Bureau of Economic Research.
- Caballero, Ricardo J., and Eduardo M.R.A. Engel. 2007. "Price Stickiness in Ss Models: New Interpretations of Old Results." *Journal of Monetary Economics* 54(1, Supplement): 100–121.
- Campbell, Jeffrey R., and Benjamin Eden. 2007. "Rigid Prices: Evidence from U.S. Scanner Data." Working Paper 2005-08. Chicago: Federal Reserve Bank of Chicago. Available at http://www.chicagofed.org/publications/workingpapers/wp2005_08.pdf.
- Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102(4): 703–725.
- Carvalho, Carlos M. 2006. "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks." *Frontiers of Macroeconomics* 2(1): Article 1. Available at <http://www.bepress.com/bejm/frontiers/vol2/iss1/art1>.
- Cecchetti, Stephen G. 1986. "The Frequency of Price Adjustment: A Study of Newsstand Prices of Magazines." *Journal of Econometrics* 31(3): 255–274.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford, 65–148. New York: Elsevier.
- Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünemann, Fabio Rünstler, and Jouko Vilmunen. 2005. "Price Setting in the Euro Area: Some Stylised Facts from Individual Consumer Price Data." European Central Bank Working Paper 524. Frankfurt: European Central Bank. Available at <http://www.ecb.int/pub/pdf/scpwps/ecbwp524.pdf>.
- Dotsey, Michael, and Robert G. King. 2005. "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models." *Journal of Monetary Economics* 52(1): 213–242.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics* 114(2): 655–690.

Eichenbaum, Martin, and Jonas Fisher. 2007. "Estimating the Frequency of Price Reoptimization in Calvo Style Pricing Models with Dynamic Indexation." *Journal of Monetary Economics* 54(7): 2032–2047.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2008. "Reference Prices and Nominal Rigidities." Centre for Economic Policy Research Discussion Paper 6709. London: Centre for Economic Policy Research.

Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Y. Mathä, Roberto Sabbatini, Harald Stahl, and Ad C. J. Stokman. 2005. "The Pricing Behaviour of Firms in the Euro Area: New Survey Evidence." European Central Bank Working Paper 535. Frankfurt: European Central Bank. Available at <http://www.ecb.int/pub/pdf/scpwps/ecbwp535.pdf>.

Fabiani, Silvia, Claire Loupias, Fernando Martins, and Roberto Sabbatini. 2007. *Pricing Decisions in the Euro Area: How Firms Set Prices and Why*. Oxford: Oxford University Press.

Fisher, Timothy C. G., and Jerzy D. Konieczny. 2000. "Synchronization of Price Changes by Multiproduct Firms: Evidence from Canadian Newspaper Prices." *Economics Letters* 68(3): 271–277.

Gagnon, Etienne. 2007. "Price Setting During Low and High Inflation: Evidence from Mexico." International Finance Discussion Paper No. 896. Washington, DC: Board of Governors of the Federal Reserve System. Available at <http://www.federalreserve.gov/pubs/ifdp/2007/896/ifdp896.pdf>.

Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.

Galí, Jordi, Mark Gertler, and J. David López-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45(7): 1237–1270.

Gaspar, Vítor, Andrew Levin, Fernando Manuel Martins, and Frank Smets. 2007. "Evidence from Surveys of Price-Setting Managers: Policy Lessons and Directions for Ongoing Research." Centre for Economic Policy Research Discussion Paper 6227. London: Centre for Economic Policy Research.

Gertler, Mark, and John Leahy. 2006. "A Phillips Curve with an Ss Foundation." Working Paper 11971. Cambridge, MA: National Bureau of Economic Research.

Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115(2): 171–199.

Kashyap, Anil K. 1995. "Sticky Prices: New Evidence from Retail Catalogs." *Quarterly Journal of Economics* 110(1): 245–274.

Kehoe, Patrick J., and Virgiliu Midrigan. 2008. "Temporary Price Changes and the Real Effects of Monetary Policy." Working Paper 14392. Cambridge, MA: National Bureau of Economic Research.

- Kim, Soyoung. 1999. "Do Monetary Policy Shocks Matter in the G-7 Countries? Using Common Identifying Assumptions About Monetary Policy Across Countries." *Journal of International Economics* 48(2): 387–412.
- Kimball, Miles S. 1995. "The Quantitative Analysis of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking* 27(4, Part 2): 1241–1277.
- Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(3): 863–904.
- Klenow, Peter J., and Jonathan L. Willis. 2006. "Real Rigidities and Nominal Price Changes." Working Paper RWP06–03. Kansas City, MO: Federal Reserve Bank of Kansas City.
- Klenow, Peter J., and Jonathan L. Willis. 2007. "Sticky Information and Sticky Prices." *Journal of Monetary Economics* 54 (Supplement 1): 79–99.
- Konieczny, Jerzy D., and Andrzej Skrzypacz. 2005. "Inflation and Price Setting in a Natural Experiment." *Journal of Monetary Economics* 52(3): 621–632.
- Kryvtsov, Oleksiy, and Virgiliu Midrigan. 2009. "Inventories, Markups, and Real Rigidities in Menu Cost Models." Working Paper 14651. Cambridge, MA: National Bureau of Economic Research.
- Lach, Saul, and Daniel Tsiddon. 1992. "The Behavior of Prices and Inflation: An Empirical Analysis Using Disaggregated Price Data." *Journal of Political Economy* 100(2): 349–389.
- Leeper, Eric M., Christopher A. Sims, and Tao Zha. 1996. "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity*, 1996(2): 1–63.
- Lucas, Robert E. Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4(2): 103–124.
- Maćkowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt. 2008. "Sectoral Price Data and Models of Price Setting." Manuscript. Available at <http://www2.wiwi.hu-berlin.de/bartosz/MMW081205.pdf>.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99(3): 769–803.
- Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117(4): 1295–1328.
- McCallum, Andrew, and Frank Smets. 2008. "Real Wages and Monetary Policy Transmission in the Euro Area." Mimeo. European Central Bank.
- Midrigan, Virgiliu. 2008. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations." Manuscript, New York University. Available at <http://www2.wiwi.hu-berlin.de/bartosz/MMW081205.pdf>.
- Nakamura, Emi, and Jón Steinsson. 2008a. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415–1464.

Nakamura, Emi, and Jón Steinsson. 2008b. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model." Working Paper 14001. Cambridge, MA: National Bureau of Economic Research.

Nakamura, Emi, and Jón Steinsson. 2009. "Price Setting in Forward-Looking Customer Markets." Manuscript, Columbia University. Available at <http://www.columbia.edu/~js3204/papers/habit.pdf>. Accessed April 18, 2009.

Okun, Arthur M. 1981. *Prices and Quantities: A Macroeconomic Analysis*. Washington, DC: Brookings Institution Press.

Phelps, Edmund S. 1969. "The New Microeconomics in Inflation and Employment Theory." *American Economic Review* 59(2): 147–160.

Reis, Ricardo. 2006. "Inattentive Producers." *Review of Economic Studies* 73(3): 793–821.

Rotemberg, Julio J. 2004. "Fair Pricing." Working Paper 10915. Cambridge, MA: National Bureau of Economic Research.

Sahuc, Jean-Guillaume, and Frank Smets. 2008. "Differences in Interest Rate Policy at the ECB and the Fed: An Investigation with a Medium-Scale DSGE Model." *Journal of Money, Credit and Banking* 40(2–3): 505–521.

Sims, Christopher A. 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49(1): 317–356.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665–690.

Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1(5): 1123–1175.

Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.

Taylor, John B. 1999. "Staggered Price and Wage Setting in Macroeconomics." In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford, 1009–1050. New York: Elsevier.

Vermeulen, Philip, Daniel Dias, Maarten Dossche, Erwan Gautier, Ignacio Hernando, Roberto Sabbatin, and Harald Stahl. 2007. "Price Setting in the Euro Area: Some Stylised Facts from Individual Producer Price Data." European Central Bank Working Paper 727. Frankfurt: European Central Bank. Available at <http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp727.pdf>.

Woodford, Michael. 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, 25–58. Princeton, NJ: Princeton University Press.

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics* 86(2): 514–533.

Comments on “Implications of Microeconomic Price Data for Macroeconomic Models” by Bartosz Maćkowiak and Frank Smets

Robert G. King

When I was about five, my parents first brought me to Cape Cod. Reading “Implications of Micro Price Data for Macroeconomic Models” by Bartosz Maćkowiak and Frank Smets is like a trip to an old-fashioned candy store on such a childhood vacation. At first glance, it is a wonderful experience simply because there are so many excellent items on the shelves. Later on, one realizes just how judiciously selective the proprietors have been: they have chosen from a large universe of potential items, presenting the best and most interesting ones to the customer, and thus providing an opportunity for a real treat. But one still must make judicious choices oneself.

In this paper, Maćkowiak and Smets have worked hard to bring us the best candy: in this case the most salient features of modern research on pricing, with a stress on the implications of micro data for macroeconomic models. It is absolutely wonderful that Maćkowiak and Smets have had to work so hard. In the last ten years, there has been an explosion of research on pricing: the work of Bils and Klenow (2004) has stimulated a new industry producing studies of price dynamics, based principally on new access to survey data collected for the consumer and producer price indexes in many countries around the world. The Inflation Persistence Network, supported by the central banks of the euro system, has produced a wealth of studies for European countries. Recent investigations extend the coverage to a wide range of other countries, including countries with higher average rates of inflation. All in all, there are new opportunities and new challenges associated with this new, more detailed information on price dynamics.

In my discussion, like Maćkowiak and Smets, I am going to selectively draw from the available studies of micro price data and highlight what I see as several key implications. In particular, I am going to argue that the micro data indicates that we need to organize our thinking around a dynamic pricing model that is very far from the Calvo (1983) model. This is the model that we presently teach to first-year Ph.D. students and its near relatives are used in many modern quantitative dynamic stochastic general equilibrium (DSGE) models. The Calvo model, like other time-dependent pricing models, abstracts from a firm's choice of the timing of the price adjustment and focuses on the magnitude of a firm's price adjustment. Calvo's model is attractive theoretically and empirically because it leads to a simple forward-looking theory of inflation and potentially is compatible with large non-neutralities arising from sticky prices. It was precisely those aspects of the Calvo model that led me to use it in the initial quantitative DSGE studies that I undertook in the mid-1990s. But those early studies used a degree of price stickiness that is simply implausible given the micro price data that we now have. Some researchers have sought to modify the Calvo model along dynamic indexation lines to generate a backward-looking component of inflation and such modifications are now a key part of many quantitative DSGE models. I argue that these modifications are so grossly at variance with the microeconomic data that they should be scrapped as devices to improve the empirical performance of macroeconomic models.

Proceeding further, I think it is useful to ask: suppose that we are forced to choose between using a model that explains only the magnitude of price adjustment (as in the Calvo model) or only the timing of price adjustment? Drawing on recent empirical work of Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008a,b) on U.S. CPI data, I conclude that we want the endogenous timing model. That is, to understand inflation, we should focus on understanding the timing rather than the magnitude of price adjustment.

1. DSGE Models with Sticky Prices

It was not always the case that there was a wealth of microeconomic pricing data. In particular, it was not the case in the mid-1990s, when econo-

mists began building a new class of small-scale DSGE models designed to allow explicit microeconomic foundations. These models included optimizing price formation by forward-looking firms, so as to undertake analysis of monetary policy consistent with the Lucas critique.

This 1990s model-building activity focused on prices, rather than wages, for several reasons. First, beginning with controversies in the 1930s, macroeconomists had become convinced that there was no firm difference between cyclical price and wage movements, so that in turn there was no strong cyclical pattern of real wages. Second, many economists found convincing the views of Barro (1977) and Hall (1980) that nominal wage bargains between firms and workers need not be allocative. For this pair of reasons, New Keynesian economists like Mankiw (1990) recommended that price stickiness be the centerpiece of new research activity, paired with imperfect competition in product markets. However, incorporating price stickiness exacts a cost in dynamic modeling: a distribution of prices is the relevant state of the economy.

The Nature of Price Dynamics

In the 1990s, as Maćkowiak and Smets stress, there was a relatively small amount of data on micro price dynamics, largely limited to studies of newspapers and catalogs. But most macroeconomists had a sense that there was important price stickiness, based on casual observation. Continuing the discussion of confections from earlier, figure 6.7 shows the weekly price of a particular cookie at Dominick's Fine Foods during the 1989–1997 period. The figure highlights the pattern of price dynamics that makes macroeconomists interested in price stickiness as a potential source of monetary non-neutrality. The time scale is weekly, so that there are several periods of six months or more during which there are no changes in the “regular” price. In fact, the price of a package of cookies is \$1.99 for most weeks during a two-year interval in the midst of the sample. The second impression is that these “constant price spells” are not of equal duration: sometimes the periods of price fixity are lengthy and sometimes these are short. The third impression is that the intervals of stickiness are occasionally interrupted by declines in the product price: there are “sales” of varying sizes and there is some tendency for the post-sale price to return to its prior level.




Figure 6.7

Cookie Price Dynamics: Price of a 5.5 Ounce Package of Pepperidge Farm Geneva Cookies at Dominick's Fine Foods, 393 Weeks Starting 09/28/89
Source: Dashed line is price obtained from the Dominick's database at the University of Chicago Booth Graduate School of Business. Solid line is an estimate of the regular price, using a sales removal filter developed by Johnston (2007).

Figure 6.7 includes the underlying actual price data, including sales, as shown by the dashed line. Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008a), and other studies remove price changes that are temporary, such as sales. Such a filtering of price data yields an estimate of the product's "regular price," shown as the solid line in figure 6.7. There are ten regular price changes in the figure, so that the product price is constant for about 35 weeks on average and there is a regular price change in about 3 percent of the weeks in the sample. There are 13 sales intervals, so that including these episodes, which involve both a decrease and an increase in the price, leads to a greater estimate of the frequency of price changes.

Given the price dynamics such as those described in figure 6.7, as argued by Rotemberg (1987), a natural first approach is that of Calvo (1983) since it can capture periods of price fixity of apparently random

length, thus producing consistency with the first two facts. The crucial characteristic of the Calvo model is that there is an exogenous constant probability of price adjustment, unrelated to macroeconomic factors or the length of time since the last adjustment. At the level of the firm, the focus is thus on the intensive margin of price adjustment: how should the size of the steps in figure 6.7 be determined?

However, in terms of developing quantitative DSGE models, there are other reasons that the Calvo model is attractive, in that it allows for a simple aggregation of a distribution of prices with only a single free parameter in the aggregator. It is that aspect of the model that led it to be used in early quantitative DSGE models and mainly accounts for its continued popularity in dynamic macroeconomic analysis.

Early DSGE Models with Sticky Prices

For a combination of tractability and empirical relevance, an initial set of DSGE models was built around a real business cycle core, modified by the introduction of monopolistic competition, sticky prices, and with various approaches to money demand. Yun (1996) developed a coherent aggregation theory for a version of the Calvo setup, constructing a framework within which Solow growth accounting and, in particular, the extraction of productivity shocks was legitimate under sticky prices. Empirically, Yun used his framework to explore the dynamic interaction of inflation, output, productivity, and monetary variables. King and Wolman (1996) focused on the policy implications of a broadly similar DSGE model, stressing that such a “St. Louis model of the 21st century” provided strong support for inflation targeting; a smooth price path made the model operate as if prices were not sticky, so that real activity responded to productivity shocks just as in the real business cycle model—although the level of real activity was reduced due to monopoly distortions. King and Wolman (1996) studied the empirical performance of the DSGE Calvo sticky price model, contrasting its explanatory power for money, interest, prices, and the business cycle with some competitor macroeconomic models.

Models constructed along DSGE-Calvo lines are now prominent in two settings. First, these are part of the standard fare of first-year graduate macroeconomics classes at many universities. Second, descen-

dants of these models are now routinely used for certain monetary policy analyses at central banks.

There has been substantial growth in computational capacity since the mid-1990s, so that much more elaborate time-dependent pricing models can easily be constructed. Such models share the Calvo model's focus on the magnitude of price adjustment—the size of the jumps in figure 6.7—while relaxing the assumption that the probability of price change is independent of time since last adjustment. I will return to discussing aspects of these more elaborate models in section 4 below, but will concentrate on the Calvo model itself as a representative of the broader class. The Calvo model allows for neat aggregation of the influence of the past and the future, as discussed next, and transparent analytical expressions not available in richer time-dependent pricing models.

Simple Dynamics of the Price Level

Suppose that the probability of price adjustment is θ , the optimal price chosen by all adjusting firms at date t is p_t^* , and the price level is P_t . Then, as is familiar, the Calvo model with a constant elasticity of substitution aggregator implies that the price level evolves according to

$$(1) \quad P_t = [\theta(P_t^*)^{1-\varepsilon} + (1-\theta)P_{t-1}^{1-\varepsilon}]^{\frac{1}{1-\varepsilon}}$$

when there is a relative demand elasticity of ε . Approximation around a zero inflation stationary point leads to

$$(2) \quad \log P_t = \theta \log P_t^* + (1-\theta) \log(P_{t-1})$$

as a convenient expression for the evolution of the price level.

These expressions highlight two key features of the Calvo model that have led researchers to use it in the construction of analytical and quantitative models. One is that this model does not track a distribution of prices because the lagged price level is the relevant summary statistic for the distribution of prices. Another is that there is a single parameter, θ , which governs the dynamics of the approximate price level.

Forward-looking Price Setting

In the Calvo model, as discussed above, the focus is on the *magnitude* of price adjustment, not the *timing* of price adjustment. Further, the model

produces a direct link between inflation, $\pi_t = (P_t/P_{t-1}) - 1$, and the real “reset” price chosen by adjusting firms,

$$(3) \quad \pi_t = [\theta(P_t^* / P_{t-1})^{1-\varepsilon} + (1-\theta)]^{\frac{1}{1-\varepsilon}} - 1.$$

In turn, this real reset price can be modeled as an optimizing decision, with Sargent’s (1978) principle that “lags imply leads” coming strongly into play. That is, given that its nominal price is sticky, a firm has a substantial incentive to forecast the inflation rate that will prevail over the duration of stickiness. In fact, the Calvo model means that firms need to forecast inflation over many future periods. That is, with probability $(1 - \theta)^j$ the firm that sets its price at t will have a real price

$$\frac{P_t^*}{P_{t+j}} = \frac{P_t^*}{P_t[(1 + \pi_{t+1})(1 + \pi_{t+2}) \dots (1 + \pi_{t+j})]}$$

in $t + j$ so that it has strong incentive to forecast inflation when setting its price.

Optimal forward-looking pricing in the Calvo model also links the optimal reset price to current and future nominal marginal cost. To a first approximation around zero inflation, the optimal reset price takes the form $\log(P_t^*) = \frac{1}{1-\beta(1-\theta)} \sum_{j=0}^{\infty} (\beta(1-\theta))^j E_t[\log(P_{t+j}) + \log(\psi_{t+j} / \psi)]$, where ψ_t is real marginal cost at date t and ψ is the corresponding steady-state value. This can conveniently be written as

$$(4) \quad \log(P_t^*) = \frac{1}{1-\beta(1-\theta)} [\log P_t + \log(\psi_t / \psi)] + \beta E_t \log(P_{t+1}^*)$$

so that there is a simple recursive structure to both the backward (2) and forward (4) components of the price block under the Calvo model.

Circa 1987, at the time of Rotemberg’s survey, the price structure was a very attractive modeling assumption. The Calvo model made a firm’s nominal prices resemble those in figure 6.7: constant for periods of time that were uneven, as available microeconomic data suggested. It led to convenient expressions for DSGE model development. But the microdata was pretty sketchy, limited to the prices of a relatively small number of products.

Dynamics of the Price Level

When quarterly models along DSGE-Calvo lines were parameterized in the mid-1990s, a standard value for θ was 0.1: it was assumed that only 10 percent of firms had the opportunity to adjust prices each quarter. The specification (2), $\log(P_t) = \theta \log(P_t^*) + (1 - \theta) \log(P_{t-1})$, thus meant that the price level response to a step change in P_t^* would be very gradual, leading to a substantial period of non-neutrality. Put another way, the average duration of a price is the reciprocal of the adjustment fraction ($1/\theta$), so that an average price was assumed to be sticky for about ten quarters. This very gradual price level adjustment seemed promising in terms of developing a lengthy pattern of non-neutrality, so that a sticky price DSGE model might behave very differently from its underlying real business cycle core.

The Discipline of the Micro Data

The recent explosion of work on micro data contains controversies, nicely reviewed by Maćkowiak and Smets, about how to measure price changes and the consequences of alternative procedures for the extent of price stickiness (durations of price fixity). One important issue is highlighted by looking back at figure 6.7: one must decide whether to treat temporary (“sale”) price declines as price changes or not, for the purpose of studying aggregate monetary non-neutrality. Analysts differ on this topic, so that there are a range of estimates of the duration of price stickiness and the frequency of price change.

Despite these differences, the recent work on micro price data has led to a sharply different view about the degree of price stickiness relative to that which prevailed in the mid-1990s. These findings are reflected in Klenow and Kryvtsov (2008), tables I and II, which conduct a sensitivity analysis of frequencies of price change and duration of price fixity to various measurement issues. Median frequencies of price changes range from 14 percent to 27 percent per *month*, while mean frequencies of price change range from 30 to 36 percent. Implied median durations range from 3.7 to 10.6 months, while implied mean durations range from 6.6 to 13.4 months. There is a clear message: prices are less sticky than was commonly assumed in the early DSGE literature, so that there is sig-

nificant discipline imposed by the micro data on price adjustment parameters in quantitative models. Thus, in Maćkowiak and Smets, there is a substantial emphasis on finding real mechanisms that can substitute for price stickiness in delivering large and protracted responses to nominal disturbances.

2. Capturing Inflation Persistence

An additional problem with the Calvo model of price dynamics for some analysts, stressed by Fuhrer and Moore (1995), is that there is no intrinsic inflation persistence. Combining the various equations discussed earlier and loglinearizing around a zero inflation steady state in ways that are now familiar, one obtains

$$(5) \quad \pi_t \approx \beta E_t \pi_{t+1} + \lambda \log(\psi_t / \psi)$$

where γ is a function of β and θ . This specification is widely employed in applied work, i.e., in the extensive empirical literature exploring inflation dynamics following Galí and Gertler (1999). From this empirical perspective, the Calvo model is attractive because it is parsimonious: there is a single parameter indicating the duration of price stickiness that is a key determinant of the Phillips curve slope λ .

An Inflation Persistence Mechanism

A number of studies have sought to add backward-looking components to a forward-looking inflation specification like (5) by a variety of schemes. For example, Galí and Gertler (1999) discuss rule of thumb price-setters. More recent empirical studies have opted to use a scheme of “dynamic indexation,” by which a firm i may update its nominal price P_{it} by

$$(6) \quad P_{it} = (1 + \pi_{t-1})P_{i,t-1}$$

if it does not adjust to P_t^* . Such assumptions are used by many currently state-of-the-art DSGE models, such as that of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003).

While there are many variants of this dynamic indexation approach, an elegant recent presentation is that of Dennis (2006). In his framework,

a fraction of firms θ adjusts its price, but only a fraction ω of these adjust to P_t^* . The fraction $(1 - \omega)$ uses the dynamic indexation scheme (6). On net, this combination of assumptions yields an inflation equation of the form

$$\pi_t = b(\theta, \omega)\pi_{t-1} + f(\theta, \omega)E_t\pi_{t+1} + \gamma(\theta, \omega)\log(\psi_t / \psi),$$

which allows for a mixture of forward-looking and backward-looking components.

Estimating this model on quarterly U.S. data using Bayesian methods, Dennis (2006) finds that 60 percent of firms change prices each quarter ($\theta = 6$), but that 90 percent of these adopt the dynamic indexation rule ($\omega = 0.1$). That is, 54 percent of all firms have a price change that is equal to the inflation rate. These parameter estimates generate a substantial backward-looking component to the inflation, as well as a relatively low response of inflation to marginal cost. Other studies, such as that of Eichenbaum and Fisher (2007), simply impose that *all* firms adjust prices every quarter ($\theta = 1$), but estimate that only a much smaller fraction “reoptimize,” setting their price to P^* .

The Discipline of the Micro Data

The dynamic indexation model—some variant of which is now widely employed in DSGE models designed for monetary policy analysis—is highly inconsistent with the micro data on two dimensions. First, looking at figure 6.7, we see intervals of prices that are constant in nominal terms: cookies stay at \$1.99 rather than being updated by the lagged inflation rate.

Looking more broadly, figure 6.8 shows the distribution of price changes in the consumer price index (CPI). This figure is taken from the research of Klenow and Kryvtsov (2008), where it appears as figure II, and was kindly provided by Pete Klenow. There is a lot of relative price variability, with large positive and negative price changes being a feature of the data in both the United States and in other countries.

To think through the implications of the dynamic indexation model, let’s imagine that there is a small positive inflation rate. Then, in the Calvo model, a fraction $1 - \theta$ of firms will not change price at all and a fraction $\theta(1 - \theta)^j$ will make a (log) price change of $j \log(1 + \pi) \approx j\pi$. Of




Figure 6.8

Weighted Distribution of Regular Price Changes in US CPI, 1988–2005

Source: July 2007 working paper version of Klenow and Kryvtsov (2008).

course, this completely misses on negative price changes, but it captures the fact that some firms change prices and others don't.

With dynamic indexation, a fraction $1 - \theta$ of firms do not change price, a fraction all firms $\theta(1 - \omega)$ have a price change of exactly π_{t-1} , while a fraction $\theta(1 - \theta)^j \omega$ have a price change of $j \log(\pi)$. According to the Dennis estimates described above, we should see 90 percent of all the adjustments in figure 6.8 being exactly at the lagged inflation rate. While there are many small price changes in most datasets, which Maćkowiak and Smets appropriately stress as surprising, there is no tendency for these to cluster at last month's inflation rate. Thus, figure 6.8 seems particularly problematic for the dynamic indexation because there is no spike in the distribution of price changes at the inflation rate, in contrast to the first-order prediction of the dynamic indexation model.¹

Thus, there is substantial discipline present in the micro data: there is no evidence of dynamic indexation. It is my view that such discipline

from the micro data is an important source of restrictions on the pricing equations in macroeconomic models, which must be imposed if we are to avoid a return to the vacuity of distributed lag econometrics latent in an earlier generation of macroeconomic policy models.

3. Micro Data and State-Dependent Pricing

Many aspects of micro price data indicate that there is considerable heterogeneity in the timing of price adjustment. Starting with Bils and Klenow (2004), many researchers have documented that the average frequency of adjustment differs across product categories, across months of the year, and so on. It is hard, at least for me, to look at this considerable heterogeneity through the lens of the Calvo model or variants of it that feature heterogeneity in exogenous adjustment frequencies (the θ parameter earlier). But some studies reviewed by Maćkowiak and Smets do follow this strategy and these indicate that heterogeneity in adjustment frequency is itself important for macroeconomic adjustment dynamics. Fortunately, since the mid-1990s, we have the computational capability to build much larger macroeconomic models, so that it is feasible to think about heterogeneity and macroeconomics, in pricing and in other areas.

But I don't think that this sort of exogenous adjustment frequency heterogeneity is enough: we need to understand how firms choose the *timing* of their price adjustments. One particular look at the micro data suggests that it is not a sideshow, but that it is quite likely critical in terms of understanding inflation dynamics.

A Stark Choice

To put the issue sharply, let's ask a very specific question. Suppose that, despite all of the increases in computational capacity, we were forced to choose between two simple and extreme structures of pricing. One option is the familiar Calvo model which, as discussed earlier, assumes exogenous timing and endogenous magnitude of price adjustment: this model focuses on the intensive margin of price adjustment as key for inflation. The other option is an as-yet-undeveloped alternative model that assumes exogenous price adjustment size and endogenous timing, which I will call the simple state-dependent pricing (SDP) model to

draw its connection to the literature: this alternative model would focus entirely on the extensive margin of price adjustment as key for inflation.

Neither of these setups would have a chance of explaining all of the dimensions of the micro data, of course, but we can still ask: which model would we choose for understanding inflation and why?

Investigating the Margins of Adjustment

To answer this question, it is useful to look at figures 6.9 and 6.10, which are drawn from the unpublished research of Nakamura and Steinsson (2008b) on the U.S. consumer price index. For 1988 to 2004, they calculate these following four statistics: m^+ , the average size of price increases; m^- , the average size of price decreases; f^+ , the fraction of firms increasing prices; and f^- = the fraction of firms decreasing prices. Their findings, as displayed in Figures 6.9 and 6.10, are highly revealing.

The magnitude of adjustment does not move strongly with inflation. The size of price changes—particularly price increases when inflation is




Figure 6.9

Magnitude of Regular Price Changes: The Size of Average Price Increases (m^+) and Average Price Decreases (m^-) in the U.S. CPI, 1988–2005

Source: Nakamura and Steinsson (2008b).




Figure 6.10

Frequency of Regular Price Changes: The Frequency Price Increases ($f+$) and Price Decreases ($f-$) in the U.S. CPI, 1988–2005

Source: Nakamura and Steinsson (2008b).

positive—is the central variable in the Calvo model. Approximating $\pi_t + 1 = P_t/P_{t-1} = [\theta(P_t^*/P_{t-1})^{1-\varepsilon} + (1-\theta)]^{1/1-\varepsilon}$ around a zero inflation steady-state, we find that

$$\pi_t \approx \theta \sum_{j=1}^{\infty} (1-\theta)^j (\log(P_t^*) - \log(P_{t-j}^*))$$

with the bracketed term being exactly the average size of price changes in the Calvo model. That is, the Calvo model predicts strong co-movement between the average size of price changes and inflation. However, as figure 6.9 shows, there is no strong relationship between inflation and the average size of price increases (m^+) or the average size of price decreases (m^-).

Now, the fact that there is no important co-movement of the size of price changes with inflation is revealing about a broader class of models: it should extend to essentially any time-dependent pricing model, not just those with an exogenous and constant adjustment hazard.

Adjustment frequency moves strongly with inflation. The fraction of firms that choose to increase prices, f^+ , is strongly positively associated with inflation, as shown in figure 6.10. The fraction of firms reducing prices, f^- , is roughly constant.

Thus, the joint message of figures 6.9 and 6.10 is that understanding the timing of price adjustments is central to macroeconomics. We need to understand the “extensive margin” of adjustment, not the “intensive margin.”

Further Information on Adjustment Timing

The information underlying figures 6.9 and 6.10 is based on a particular set of price adjustment definitions and the results are reported at the annual frequency. Klenow and Kryvtsov (2008) also explore the four statistics developed by Nakamura and Steinsson, working with somewhat different definitions of price changes and examining co-movement at higher frequencies. Like Nakamura and Steinsson, Klenow and Kryvtsov find that it is the fraction of firms raising prices that correlates most strongly with inflation ($\text{corr}(f^+, \pi) = .69$), but they also find that the fraction of firms lowering prices is negatively correlated with inflation ($\text{corr}(f^-, \pi) = -.41$). Finally, they find that there is much smaller correlation of inflation with the magnitudes of price increases ($\text{corr}(m^+, \pi) = .19$) or decreases ($\text{corr}(m^-, \pi) = -.19$). Although there is some action on the intensive margin, these more detailed findings suggest that understanding the timing of price adjustments is central.

Thus, the simple model that we presently teach in our first-year classes badly misses out on the key co-movement, which is between inflation and adjustment frequency, and instead highlights a less important mechanism, which is a link between the magnitude of price changes and inflation.

Conclusion

The new data on micro prices provides discipline on quantitative macroeconomic model building and also provides challenges to currently popular views about nature of the DSGE models that must be constructed.

The standard model of Calvo (1983), as variously elaborated to provide empirical underpinning for price blocks in quantitative macroeco-

conomic models, fares badly vis-a-vis the micro price data. The available evidence is that price adjustment is relatively frequent, which limits the extent to which the price equations of a macroeconomic model can readily rationalize monetary non-neutrality. A standard extension of the basic model allows firms to costlessly index frequently to the past inflation rate, but not to make frequent fully optimal adjustments. While that dynamic indexation model can readily produce both inflation persistence and larger non-neutralities, it is dramatically inconsistent with the micro price data: there is just no evidence that firms actually adjust prices in the manner suggested by the dynamic indexation approach.

Further, there is relatively weak co-movement of the magnitude of price changes—the intensive margin of price adjustment—with inflation in recent empirical studies of U.S. micro price data, as would be suggested by most currently popular pricing models. However, there is strong co-movement of the fraction of firms that raise prices with the inflation rate. This evidence suggests that it is important to understand when firms choose to adjust prices—i.e., that a central focus for macroeconomic research should be to better understand the extensive margin of price adjustment. When we teach sticky price models to our first-year graduate students, they would be better served by our using an as-yet-undeveloped model that focuses solely on the extensive margin of price adjustment rather than the Calvo model, which focuses solely on the intensive margin.

Notes

1. Of course, there is one basic problem in trying to use *any* model in this class to explain the micro data, as Lucas and Golosov (2007) stress. If there is positive inflation, there is never any reason for a negative price change. But let's suppose that there might be some relatively easy way to fix this, by adding in microeconomic shocks and allowing for adjustment to these.

References

- Barro, Robert J. 1977. "Long-term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3(3): 305–316.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.

- Calvo, Guillermo A. 1983. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–398.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.
- Dennis, Richard. 2006. "The Frequency of Price Adjustment and New Keynesian Business Cycle Dynamics." Working Paper 2006–22. San Francisco: Federal Reserve Bank of San Francisco.
- Eichenbaum, Martin S., and Jonas Fisher. 2007. "Estimating the Frequency of Price Reoptimization in Calvo Style Pricing Models with Dynamic Indexation." *Journal of Monetary Economics* 54(7): 2032–2047.
- Fuhrer, Jeff, and George Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110(1): 127–160.
- Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.
- Golosov, Mikhail, and Robert E. Lucas Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115(2): 171–199.
- Hall, Robert, E. 1980. "The Rigidity of Wages and the Persistence of Unemployment." *Brookings Papers on Economic Activity* 2: 301–335.
- Johnston, Michael K. 2007. "Should Macroeconomists Discount Sales?" Working paper. Available at <http://michaelkjohnston.com/perm/mkj-sales2007.pdf>.
- King, Robert G., and Alexander L. Wolman. 1996. "Inflation Targeting in a St. Louis Model of the 21st Century." *Review* (Fall): 83–107. St. Louis: Federal Reserve Bank of St. Louis. Available at <http://research.stlouisfed.org/publications/review/96/05/9605rk.pdf>.
- Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(3): 863–904.
- Mankiw, N. Gregory. 1990. "A Quick Refresher Course in Macroeconomics." *Journal of Economic Literature* 28(4): 1645–1660.
- Nakamura, Emi, and Jón Steinsson, 2008a. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415–1464.
- Nakamura, Emi, and Jón Steinsson. 2008b. "More Facts about Prices—Supplement to 'Five Facts about Prices.'" Available at <http://www.columbia.edu/~en2198/papers.html>.
- Rotemberg, Julio J. 1987. "The New Keynesian Microfoundations." In *NBER Macroeconomics Annual* 2, ed. Stanley Fischer, 69–104. Cambridge, MA: The MIT Press.
- Sargent, Thomas J. 1978. "Estimation of Dynamic Labor Demand Schedules Under Rational Expectations." *Journal of Political Economy* 86(6): 1009–1044.

Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1(5): 1123–1175.

Yun, Tack. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37(2): 345–370.

Comments on “Implications of Microeconomic Price Data for Macroeconomic Models” by Bartosz Maćkowiak and Frank Smets

Virgiliu Midrigan

The paper by Bartosz Maćkowiak and Frank Smets provides a very careful survey of the recent work that uses micro price data in order to study the role of sticky prices in shaping the monetary transmission mechanism. It has three thematic parts. The first one is a survey of the facts uncovered in recent work using new price data made available by the Bureau of Labor Statistics (BLS) and through the Inflation Persistence Network, as well as from scanner price data. The picture that emerges is one of considerable complexity and heterogeneity in the price-setting behavior of retail firms. Dealing with this complexity has been the goal of a recent surge of work inspired by this data, some of which I will review in greater detail. The second theme, pretty much independent of how one deals with this price-setting complexity, is the overriding conclusion that prices change much too frequently, both in Europe and in the United States, for menu costs alone to be able to account for the sluggishness in prices at the aggregate level. The third part of the paper deals with the challenge of bridging the gap between the flexibility observed in the micro data and the inertia in the aggregates. In particular, the authors emphasize the real rigidities and informational frictions as promising avenues for building models in which firms do not respond to nominal disturbances even though they frequently adjust prices. I would like to add inventory-theoretic models of money demand that emphasize frictions that segment asset markets to the list of potential solutions to this challenge.

I will start by describing some of the challenges the wealth of micro price data has posed to macroeconomists and some of the progress we have made to address these issues. The first challenge is the large heterogeneity in the frequency of price changes across different sectors.

To see the extent of this heterogeneity, the frequency of price changes in the *Bils and Klenow (2004)* study ranges from coin-operated laundry machines that undergo price changes every 80 months on average to regular unleaded gasoline that undergoes price changes about every two weeks on average. There is also considerable skewness in this distribution: according to *Nakamura and Steinsson (2008a)*, the median frequency of price changes is 8.7 percent per month, while the mean frequency is 21.1 percent per month. Thus a question naturally arises: what is the effect of this heterogeneity on the real effects of money in sticky price models? What frequency of price adjustment should one use in a homogeneous-firm economy that abstracts from this heterogeneity? *Carvalho (2006)* studies this question in the context of a time-dependent model. He finds that the heterogeneity in price adjustment amplifies the real effects of money because of the convexity of the relationship between the real effects of money and the frequency of price changes. In fact, a heterogeneous-sector economy calibrated to match the distribution of frequency of adjustment in the micro data generates three times larger real effects than a homogeneous-sector economy calibrated to the mean of this distribution. *Nakamura and Steinsson (2008b)* study a similar question in the context of a menu-cost setting. Again, they find that heterogeneity amplifies the real effects of money. Moreover, although a homogeneous-agent economy calibrated to the mean (21.1 percent) frequency of price adjustment in the data understates the real effects of money, a similar homogeneous-agent economy calibrated to the median (8.7 percent) frequency of price changes generates similar real effects as the heterogeneous-agent economy.

A second challenge discussed in the paper by *Maćkowiak and Smets* is the fact that an important number of price changes in the data are temporary discounts (sales). These are special types of price changes because sales are short-lived and because firms frequently revert to the price they were charging prior to the sale. If one includes sales in the definition of price changes, prices appear to change frequently, every 4.3 months in the study by *Bils and Klenow (2004)*. If one excludes sales from the definition of price changes, prices are much stickier and change every 9 to 11 months, as shown in the study by *Nakamura and Steinsson (2008a)*. *Kehoe and Midrigan (2008)* take up the issue of how to deal with sales

in the data. They emphasize that the special feature of sale prices is that these are temporary, and Kehoe and Midrigan formulate a model in which firms have two technologies for price adjustment available. The firm can pay a fixed cost to change its regular price, as well as another fixed cost to undertake a temporary price change. If the firm changes its regular price, it buys the option to sell at this price forever into the future. In contrast, if the firm changes its price temporarily, it is able to charge this price for one period only, after which, absent another intervention, the firm charges its old regular price. The model Kehoe and Midrigan write thus is able to match the salient features of the micro price data, including the predominance of temporary price changes. They find that the approach of leaving temporary price changes out of the data overstates the real effects of money by 40 percent. In contrast, the approach of leaving sales in the data and treating these like any other price change predicts real effects of money that are one-fifth those in the model in which sales are temporary. Kehoe and Midrigan also discuss a number of shortcuts that can be used in economies without temporary price changes to replicate the real effects of money in an economy that has both regular and temporary price changes.

A third challenge is the complicated mapping from the frequency of price adjustment to the degree of monetary non-neutrality in economies that explicitly model price stickiness as arising from menu costs. The question really is does the Caplin-Spulber (1987) well-known neutrality result in a menu cost economy survive in a more general setup that is consistent with the micro price data? In the Caplin-Spulber model, money is neutral for any degree of price stickiness at the firm level because the firms that charge prices at any given point are exactly those that need the largest price change. Thus, even though few firms adjust, those that do so change prices by an amount that is sufficiently large to ensure that the aggregate price level responds one-for-one to changes in the growth rate of money. Key to this result, however, is the self-replicating uniform distribution of desired price changes of individual firms—thus leaving open the question of how general is this finding? A recent paper by Golosov and Lucas (2007) argues that the result survives in a much more general setting, provided that one allows for large idiosyncratic shocks that permit the economy to match the roughly 10 percent average size of price

changes Klenow and Krystov (2007) document in the BLS data. The intuition is similar to that of Caplin and Spulber: the firms that choose to adjust in times of, say, a monetary expansion, are exactly those firms that stand to gain most from increasing their prices. Therefore money has a strong selection effect, in that it affects the identity of adjusting firms, thus rendering the aggregate price level much more flexible than individual prices. Midrigan (2008) shows, however, that the strength of this selection effect is much smaller if one calibrates the economy to match the large dispersion (in addition to the mean targeted by Golosov and Lucas) in the size of price changes. An economy able to replicate the large number of small price changes as well as the fat tails of this distribution implies a much smaller selection effect (money has much less effect on who gets to adjust prices) than in the Golosov and Lucas model. Midrigan (2008) predicts real effects of money that are four-fifths of those found in the Calvo model where the selection effect is, by assumption, absent.

My own reading of the literature that the authors survey is that we have made considerable progress toward understanding the role played by menu costs in light of the evidence from the micro data on prices. However, what has also become clear is that menu costs alone are far from sufficient in gathering a sizable monetary transmission mechanism. Here I use a narrow definition of what menu costs are: restrictions on the price-setting technology that prevents firms from changing prices. To see why menu costs are not enough to account for changes in prices, recall that the frequency of price changes in the micro data is two to three quarters on average. In contrast, empirical evidence suggests that there is much more inertia in the response of the aggregate price level to nominal shocks. Models in which the only frictions are technological restrictions on the firm's ability to reprice can replicate this aggregate inertia only if prices change as infrequently as once every ten quarters. Two other pieces of evidence point out that the physical costs of price adjustment are not sufficient to explain price stickiness in the aggregate. First, the fact that after a sale prices return to the presale level, cent for cent, makes it clear that it is not the physical menu cost in itself that renders prices sticky. The return from a sale involves changing the price and paying the menu cost; firms nevertheless choose to return to the price level in effect prior to the sale. A second piece of evidence comes from work by Kehoe and

Midrigan (2007) on cross-sectional international real exchange rates. They find that the persistence and volatility of real exchange rates for disaggregated sectors of the economy is pretty much independent of the stickiness of prices in that sector: flexible-price goods like unprocessed foods or fuels have international relative prices that track nominal exchange rates as closely as very sticky-priced goods like services. This evidence makes it clear that the reasons prices do not change with nominal exchange rates is not because firms *cannot* respond to these movements, but rather because firms *choose not to* respond.

Maćkowiak and Smets list two mechanisms that may be able to bridge the gap between the flexibility in the micro price data and the inertia in the aggregate: real rigidities and information frictions. To this list, I would like to add inventory-theoretic models of money demand, as in Alvarez, Atkeson, and Edmond (forthcoming), in which changes in the velocity of money in the aftermath of nominal shocks impart sluggishness to the aggregate price level despite flexibility in micro prices. Considerable progress has been made in understanding the properties of economies that feature these mechanisms. I argue, however, that more empirical work, using both microeconomic and macroeconomic data, is needed in order to distinguish among these competing models and to assess their quantitative properties. The distinction between these models is crucial for formulating policy prescriptions. Even though these different mechanisms may have similar implications for the slope of the Phillips curve, these will, in general, have different implications about the behavior of other objects (for instance, relative price variability and other distortions caused by inflation) that may be of interest to policymakers.

For example, take the issue of real rigidities. Two types of real rigidities have been emphasized by Ball and Romer (1990). The first type of rigidity hinges on strategic interactions among competing firms that increase an individual firm's losses from having its price deviate from that of competing firms. With kinked demand curves or upward sloping marginal cost curves at the individual firm level, price setters have a strong incentive to keep prices close to those of their competitors. Adjusting firms thus *choose not to* respond to changes in the stance of monetary policy even when they do pay the menu costs and adjust their nominal prices. Recent work by Dotsey and King (2006) and Klenow and Willis (2007), as well

as Burstein and Hellwig (2007), has called into question the qualitative relevance of this mechanism. Intuitively, strategic complementarities of this type make it optimal for firms that adjust prices by small amounts to respond little to monetary shocks because it is costly for a firm's price to be very different from its competitors that do not reprice. But price changes tend to be large (10 percent on average): implausibly large menu costs and idiosyncratic shocks are thus needed to reconcile this first type of real price rigidity with the micro price data. Although the conclusions of these recent studies may be challenged by explicitly modeling the source of the strategic complementarity (for example, search frictions or customers' switching costs) and allowing for multiproduct price setters (which would render the relevant price as the price of the bundle of goods the firm charges, rather than the price of individual items it sells), these studies are examples of well-executed and careful attempts to quantify the importance of this first type of real rigidity.

A second type of real rigidity hinges on the slow responsiveness of aggregate real marginal costs to output fluctuations. Sticky wages, intermediate goods prices, and assumptions on preferences or technology that keep the marginal costs of producing low during monetary expansions dampen the response of prices to monetary shocks and lead to more inertia in the aggregate price level. Measuring the behavior of real marginal costs over the business cycle is the subject of a large body of work (summarized by Rotemberg and Woodford 1999), but is also a difficult task. Mapping observed factor prices into marginal costs requires making assumptions about technology, the nature of the relationship between buyers and suppliers, household preferences for intertemporal substitution of leisure, and so on. Bilal and Khan (2000) argue that one can infer much about the behavior of marginal costs from the cyclical behavior of inventories. If marginal costs are strongly procyclical and rise sharply in the aftermath of a monetary expansion, firms should reduce their level of inventories. Similarly, if real marginal costs moves little over the business cycle, inventories should increase. In recent work Kryvtsov and Midrigan (2009) show that the elasticity of real marginal costs with respect to output needed to account for the strongly countercyclical inventory-sales ratio observed in the data is only slightly lower than the inverse of the intertemporal elasticity of substitution for consumption. This finding,

reminiscent of some recent results in the lumpy investment literature,¹ suggests that one can use inventories, in addition to other aggregate variables, to pin down the size of this second type of real rigidity in estimated dynamic stochastic general equilibrium models.

To conclude, the exciting recent work using micro price data has proven very useful in sharpening our understanding of the role of menu costs in accounting for the slope of the inflation-output trade-off. This work has also convincingly reinforced Ball and Romer's (1990) insights that menu costs alone are far from sufficient in accounting for the extent of monetary non-neutrality observed in the data. More empirical work aimed at distinguishing among several competing mechanisms that can amplify the effect of menu costs on aggregate price inertia should answer many open questions in the future.

Notes

1. See Thomas (2002); Khan and Thomas (2007); Bachman, Caballero, and Engel (2008); and House (2007).

References

- Alvarez, Fernando, Andrew Atkeson, and Chris Edmond. Forthcoming. "Sluggish Responses of Prices and Inflation to Monetary Shocks in an Inventory Model of Money Demand." *Quarterly Journal of Economics*.
- Bachman, Ruediger, Ricardo J. Caballero, and Eduardo M.R.A. Engel. 2008. "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model." Manuscript. Available at <http://econ-www.mit.edu/files/2864>.
- Ball, Laurence, and David Romer. 1990. "Real Rigidities and the Non-Neutrality of Money." *Review of Economic Studies* 57(2): 183–203.
- Bils, Mark, and James A. Khan. 2000. "What Inventory Behavior Tells Us About Business Cycles." *American Economic Review* 90(3): 458–481.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.
- Burstein, Ariel T., and Christian Hellwig. 2007. "Prices and Market Shares in a Menu Cost Model." Working Paper 13455. Cambridge, MA: National Bureau of Economic Research.
- Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102(4): 703–726.

Carvalho, Carlos M. 2006. "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks." *Frontiers of Macroeconomics* 2(1): Article 1. Available at: <http://www.bepress.com/bejm/frontiers/vol2/iss1/art1>.

Dotsey, Michael, and Robert G. King. 2006. "Pricing, Production, and Persistence." *Journal of the European Economic Association* 4(5): 893–928.

Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115(2): 171–199.

House, Christopher L. 2007. "Fixed Costs and Long-Lived Investments." Working Paper 14402. Cambridge, MA: National Bureau of Economic Research.

Kehoe, Patrick J., and Virgiliu Midrigan. 2007. "Sticky Prices and Sectoral Real Exchange Rates." Research Department Working Paper 656. Minneapolis: Federal Reserve Bank of Minneapolis. Available at <http://www.minneapolisfed.org/research/WP/WP656.pdf>.

Kehoe, Patrick J., and Virgiliu Midrigan. 2008. "Temporary Price Changes and the Real Effects of Monetary Policy." Working Paper 14392. Cambridge, MA: National Bureau of Economic Research.

Khan, Aubhik, and Julia K. Thomas. 2007. "Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies." *American Economic Review* 97(4): 1165–1188.

Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(3): 863–904.

Klenow, Peter J., and Jonathan L. Willis. 2007. "Sticky Information and Sticky Prices." *Journal of Monetary Economics* 54 (Supplement 1): 79–99.

Kryvtsov, Oleksiy, and Virgiliu Midrigan. 2009. "Inventories, Markups, and Real Rigidities in Menu Cost Models." Working Paper 14651. Cambridge, MA: National Bureau of Economic Research.

Midrigan, Virgiliu. 2008. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations." Mimeo. Available at http://homepages.nyu.edu/~vm50/Main_RR.pdf.

Nakamura, Emi, and Jón Steinsson. 2008a. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415–1464.

Nakamura, Emi, and Jón Steinsson. 2008b. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model." Working Paper 14001. Cambridge, MA: National Bureau of Economic Research.

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics* Volume 1, ed. John B. Taylor and Michael Woodford, 1051–1135. Amsterdam: Elsevier.

Thomas, Julia K. 2002. "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy* 110(3): 508–534.

7

Is the Phillips Curve Vertical in the Long Run?

Hysteresis in Unemployment

Laurence Ball

1. Introduction

Much of mainstream macroeconomics is based on an “accelerationist” Phillips curve, which was described by Friedman (1968) in his 1967 Presidential Address to the American Economic Association. A simple form is

$$(1) \quad \pi = \pi_{-1} + \alpha(U - U^*), \quad \alpha < 0.$$

Inflation depends on lagged inflation, often interpreted as a proxy for expected inflation. It also depends on the deviation of unemployment from the natural rate or the non-accelerating inflation rate of unemployment (NAIRU), U^* .¹

According to Friedman, shifts in aggregate demand coming from either monetary policy or other sources have short-run impacts on unemployment. In the long run, however, U always returns to U^* . And U^* is not influenced by aggregate demand. Instead, it is determined by the supply side of the economy, especially frictions in labor markets. This means the classical dichotomy holds in the long run: monetary policy cannot cause long-run changes in unemployment.

Practically speaking, most economists think monetary policy can push U away from U^* for a few years. Paul Volcker, for example, managed to raise unemployment over 1980–1983. However, many economists hold that changes in unemployment over a decade or more are determined by changes in the natural rate. They contend that the fact that unemployment was substantially higher in France in 2000 than in 1980 has little or nothing to do with monetary policy.

My essay questions this conventional wisdom. I accept equation (1), but not the view that only supply-side factors influence U^* . I believe in the concept of hysteresis advocated by Blanchard and Summers (1986): the natural rate can be influenced by the path of actual unemployment. If U rises above U^* , for example, there exist mechanisms that pull U^* upward. Since aggregate demand influences U , hysteresis means that demand also influences U^* .

Hysteresis is central to long-run unemployment movements in many countries. If we want to know why unemployment rose in much of Europe in the 1980s, or why it fell in the United Kingdom in the 1990s, or why it has remained relatively stable in the United States, we need to understand hysteresis.

This essay addresses two broad issues. The first is whether there is clear evidence of hysteresis effects. To put it differently, can we reject the hypothesis that the NAIRU, and hence the long-run behavior of unemployment, is independent of aggregate demand?

The answer to this question is **YES!** I review past evidence on hysteresis and present some new evidence.

The second broad issue is the nature of hysteresis. Through what mechanisms do short-run unemployment movements influence the NAIRU? What determines the strength of these effects in different countries and time periods? What are the implications for monetary policy?

My discussion of these topics is speculative. In my view, it is clear that *some* form of hysteresis exists, but it is not clear why. The relationships among unemployment, the natural rate, and inflation appear to be nonlinear, but it is hard to pin down the nonlinearities precisely. As a result, policy implications are not crisp.

In sum, hysteresis is an important phenomenon, but one that is not well understood. This means more research is needed. The topic of hysteresis has been neglected in recent years, and that should change.

2. The Phillips Curve and the Changing NAIRU

Friedman says, “There is always a temporary trade-off between unemployment and inflation; there is no permanent trade-off. The temporary trade-off comes not from inflation per se, but from unanticipated infla-

tion, which generally means, from a rising rate of inflation” (1968). Eventually, says Friedman, unemployment returns to the natural rate.

Friedman’s theory is summarized by equation (1). Today economists use NAIRU (for non-accelerating inflation rate of unemployment) as a synonym for the natural rate, because the natural rate is the unemployment level consistent with stable inflation. Forty years after Friedman wrote, equation (1) is a foundation for much of applied macroeconomics.

Friedman says the natural rate depends on features of the labor market such as minimum wages, labor unions, and frictions in matching the unemployed with job vacancies. He says that monetary policy cannot affect the natural rate. Friedman suggests that the natural rate may change over time, and experience has shown that it does. In the United States, the NAIRU has varied by moderate amounts; according to estimates (detailed later), it fell from 7.1 percent in 1980 to 4.9 percent in 2007. In Europe, the NAIRU has changed by larger amounts; in Spain, it rose from 6.5 percent in 1980 to 14.4 percent in 1995, then fell to 7.5 percent in 2007.

A large literature has tried to explain changes in the NAIRU. Some researchers focus on changes in labor-market imperfections of the type discussed by Friedman (for instance, Nickell, Nunziata, and Ochel 2005). Others examine interactions between such institutions and economic shocks, such as the productivity slowdown and globalization (Blanchard and Wolfers 2000; Blanchard 2005). While the stories vary, they almost always involve the supply side of the economy. They presume that aggregate demand does not affect the NAIRU.

Much work focuses on Europe, where the NAIRU rose dramatically between 1960 and 2000. Mankiw (2007) tells undergraduates that there is a “leading theory” of this experience, one from the class of shock/institution theories. In this story, proposed by Krugman (1994) and others, the shock is a decrease in the demand for low-skill labor caused by technological change. The institutions are labor-market distortions that create wage rigidity. The equilibrium wages of low-skill workers have fallen but their actual wages have not, so unemployment has risen.

This story gets much of its appeal from the fact that it fits two data points, the United States and aggregate Europe. The United States has more flexible labor markets than Europe and has not experienced a rise

in the NAIRU. We will see, however, that the story works less well when we extend the sample from two economies to twenty.

Departing from most of the literature, this paper will argue that NAIRU changes are caused largely by shifts in aggregate demand. Demand influences actual unemployment, U , which in turn influences the natural rate through hysteresis channels.

What are these channels? When Blanchard and Summers (1986) introduced the idea of hysteresis, they emphasized the “insider-outsider” theory of wage bargaining. When workers become unemployed, the remaining employed workers increase their wage targets, preventing the unemployed from getting their jobs back. In my view, however, there is little evidence for this kind of hysteresis effect.

There is more evidence for stories in which the long-term unemployed become detached from the labor market. These workers are unattractive to employers, or they do not try hard to find jobs. These stories fit evidence that hysteresis effects are stronger in countries with long-lived unemployment benefits. However, as discussed later, we have at best a hazy understanding of hysteresis mechanisms.

Allowing for hysteresis can greatly change our explanations for unemployment movements and our prescriptions for monetary policy. However, I do not view hysteresis as a radical departure from mainstream economic theory. It is not a rejection of Friedman’s model, but a generalization of it. We expand the set of factors that cause the U^* term in equation (1) to change over time: these factors include movements in actual unemployment as well as supply-side variables.

To study movements in the NAIRU, we need to estimate this variable. One simple method is to smooth the series for actual unemployment with the Hodrick-Prescott (HP) filter, based on the idea that the NAIRU is the long-term trend in unemployment. In this paper, I use a somewhat more sophisticated technique based on Ball and Mankiw (2002). This procedure modifies the results from a univariate smoother based on the behavior of inflation. During a period of falling inflation, for example, the Ball-Mankiw method produces lower NAIRU estimates than a univariate smoother, because falling inflation suggests that U^* is below U .

The appendix to this paper details my procedure for estimating the NAIRU. As an example of the results, figure 7.1 shows the estimated




Figure 7.1
Unemployment in France, 1980–2007
Source: OECD and author's calculations.

NAIRU for France from 1980 to 2007 and compares it to a univariate unemployment trend (based on the HP filter with $\lambda = 100$). The estimated NAIRU is below the univariate trend over 1980–1997, reflecting the fact that inflation was falling. After that the two series converge, as inflation was stable.

I estimate NAIRU series for 20 countries: all the countries with populations above one million that were members of the Organisation of Economic Co-operation and Development (OECD) in 1985. This group includes 15 countries in Western Europe, Canada and the United States, Australia, New Zealand, and Japan.

I focus on the period since 1980. The NAIRU rose in many countries during the 1970s, but it is harder to detect hysteresis effects in that period. The large supply shocks make it harder to estimate Phillips curves and NAIRUs. Also, there was a major change in the real economy—the productivity slowdown—that probably increased the NAIRU in many countries. Hysteresis effects may have been secondary in the 1970s. Since 1980, however, hysteresis is a big part of the unemployment story.

3. Previous Evidence of Hysteresis

This paper will confess to major gaps in our understanding of hysteresis, but argue that it clearly exists in some form. That is, there is strong evidence against the hypothesis that movements in the NAIRU are independent of aggregate demand.

Here I discuss evidence for hysteresis in previous work. I emphasize two papers from some time ago: Ball (1997) and Ball (1999). The reason I focus on my own past work is not narcissism (or, at least, that is not the only reason). Beyond my work, there is not much literature to review, as most researchers in the twenty-first century have ignored hysteresis. However, the Boston Fed's invitation to write about the topic has rekindled my hope that economists will take it seriously.²

Disinflation in the 1980s

My 1997 paper examines changes in the NAIRU from 1980 to 1990. It uses estimates of the NAIRU produced by the OECD with a method that is similar in spirit to the Ball-Mankiw method. According to these estimates, the NAIRU rose over the 1980s in 17 of the 20 countries in the sample. NAIRU changes ranged from -1.4 percent in the United States and Portugal to $+9.3$ percent in Ireland.

I argue that NAIRU increases in the 1980s were caused largely by monetary tightenings aimed at reducing inflation. This conclusion is based on the following evidence:

- Measures of labor market distortions—the degree of unionization, the severity of firing restrictions, and so on—are generally uncorrelated across countries with changes in the NAIRU. The only exception is a weak effect of the duration of unemployment benefits. Overall, this is evidence against the Krugman story about the interaction of labor demand shifts with labor market rigidities. That story predicts greater increases in the NAIRU where rigidities are greater.
- Of the 20 countries, 19 reduced inflation over the 1980s. There is a significant relationship across countries between the size of the inflation decrease and the change in the NAIRU. My interpretation is that larger disinflations required larger monetary tightenings, therefore raised unemployment more, and therefore raised the NAIRU more through hysteresis.

- The change in the NAIRU is related not only to how much inflation fell, but also to the length of time over which disinflation occurred. Holding constant the total fall in inflation, a quick disinflation raises the NAIRU less than one that is drawn out over time. This result suggests mechanisms for hysteresis, as discussed later.
- While measures of labor-market distortions are generally uncorrelated with NAIRU changes, one of these variables—the duration of unemployment benefits—interacts significantly with the size and length of disinflation. That is, a given disinflation is associated with a larger rise in the NAIRU if unemployment benefits are available indefinitely. Once we control for this interaction, there is no direct effect of benefit duration. Again, this result is suggestive about hysteresis mechanisms.

Policy Responses to Recessions

My 1999 paper examines the disinflations of the 1980s from another angle. Countries that reduced inflation generally experienced recessions and short-run rises in unemployment. However, the aftermath of disinflation varied: in some countries unemployment fell again after a few years, while in others the NAIRU rose and unemployment stayed high.

I argue that these differences are largely explained by the conduct of monetary policy. Some central banks tightened policy to reduce inflation, but reversed course when recessions occurred. They eased policy, pushing unemployment back down. Other central banks tightened policy and kept it tight, so high unemployment persisted.

To make this argument, the 1999 paper measures policy responses to recessions with changes in nominal and real interest rates (following Romer and Romer 1994). The Federal Reserve is one central bank that cut rates sharply when recessions began, even though inflation had not yet fallen significantly. Many European central banks, by contrast, did not respond aggressively to recessions. They were reluctant to ease policy until inflation was clearly defeated. In addition, exchange-rate concerns deterred some central banks from cutting interest rates.

The paper also measures the degree of hysteresis in each country by comparing increases in the NAIRU to short-run increases in unemployment during disinflation. I find that hysteresis effects are larger when central banks respond less strongly to recessions.

By itself, the fact that persistently tight policy causes persistently high unemployment is consistent with conventional macroeconomics, specifically the IS curve. Where the early-1980s experience deviates from conventional models is in the behavior of inflation. If a monetary tightening does not affect the NAIRU, then equation (1) says inflation should fall as long as unemployment remains high. In many countries, inflation fell for a few years but then leveled off with unemployment still high. This meant by definition that the NAIRU rose.

Success Stories

The NAIRU started falling in some countries in the mid-1980s. Another part of my 1999 paper asks why. I focus on four countries that reduced the NAIRU (as estimated by the OECD) by at least 2 percentage points between 1985 and 1997. One is the United Kingdom, which reversed the NAIRU run-up of the early 1980s; the others are Ireland, the Netherlands, and Portugal.

Many observers attribute these NAIRU decreases to reductions in labor-market distortions (for example, Siebert 1997). But this interpretation does not withstand scrutiny. Countries where the NAIRU fell did implement some labor-market reforms, but these were modest. For example, the Netherlands slightly reduced the replacement ratio for unemployment insurance (UI), and the United Kingdom increased job counseling for UI recipients; neither country put a time limit on benefits. Many other countries had similar or more extensive labor-market reforms (Spain reduced the replacement ratio by the same amount as the Netherlands, and Belgium introduced a job-placement program similar to the United Kingdom's). The four countries where unemployment fell do not stand out as aggressive labor-market reformers.

Instead, these countries stand out for their macroeconomic histories: they experienced demand expansions during the period when the NAIRU fell. The demand expansions reduced unemployment, which reduced the NAIRU through hysteresis.

The United Kingdom, for example, departed from the monetary policy of other European countries when it dropped out of the Exchange Rate Mechanism in 1992 and lowered interest rates. Before that, in the late 1980s, the United Kingdom experienced the "Lawson boom," named

after the Chancellor of the Exchequer who pursued fiscal expansion at the same time that financial liberalization raised consumption and investment. The Bank of England was slow to tighten policy to choke off the boom, and inflation rose by more than 5 percentage points.

A substantial run-up in inflation also accompanied the NAIRU decreases in Portugal and the Netherlands (although not in Ireland). As in the United Kingdom, central banks did not raise inflation intentionally, but they failed to offset expansionary shocks. In my view, the coincidence of rising inflation with a falling NAIRU suggests that hysteresis is at work, meaning that a demand expansion is driving the NAIRU down. I return to this point later.

An important nuance is that the inflation run-ups in the United Kingdom and elsewhere were not permanent. A period of overheating and rising inflation was needed to reduce the NAIRU, but eventually inflation went back down. And when that happened, the NAIRU did *not* go back up.

4. Some New Evidence for Hysteresis

Here I present new evidence of hysteresis effects. I try to capture these effects in a simple way using data from 1980 through 2007.

My strategy is to focus on *large* changes in the NAIRU. I define large to mean a rise or fall of at least 3 percentage points. While my method for estimating the NAIRU is imprecise, an estimated change of 3 percent almost certainly indicates a substantial change in the true NAIRU.

I restrict attention to episodes in which the NAIRU changed by at least 3 percent within a period of ten years. This ten-year rule means I ignore changes in the NAIRU that are substantial but very gradual. It is harder to identify the sources of gradual changes than of relatively abrupt changes.

Usually the ten-year periods I identify lie within longer periods in which the NAIRU moves in the same direction. I define a NAIRU-change episode as the entire period in which the NAIRU moves in one direction. This implies that episodes start and end at peaks and troughs in the NAIRU series, or at the start and end of the 1980–2007 period.

In France, for example, the NAIRU increased from 1980, when it was 5.4 percent, to 1996, when it peaked at 9.4 percent (see figure 7.1). This

period qualifies for my set of episodes because the NAIRU rose by more than 3 percent over the ten years from 1980 to 1990.

For the 20 countries in the sample, there are eight episodes of NAIRU increases that meet my criteria and nine episodes of NAIRU decreases. Table 7.1 lists the episodes, their dates, and the changes in the NAIRU over the episodes.

For each episode, I examine the behavior of inflation. This seems a natural way to distinguish between conventional stories about NAIRU changes and hysteresis theories. In hysteresis theories, changes in the NAIRU are driven by demand movements that initially push U away

Table 7.1
Large Changes in the NAIRU, 1980–1997

INCREASES

Country	Period of Change	Size of Change
Finland	1980–1996	9.7 percent
France	1980–1996	4.0 percent
Germany	1980–2007	5.6 percent
Ireland	1980–1989	5.2 percent
Italy	1980–1996	4.9 percent
New Zealand	1980–1994	4.9 percent
Spain	1980–1995	7.8 percent
Sweden	1983–1999	4.1 percent

DECREASES

Australia	1994–2007	–4.0 percent
Finland	1996–2007	–4.3 percent
Ireland	1989–2007	–11.0 percent
Italy	1996–2007	–3.9 percent
The Netherlands	1988–2007	–3.8 percent
New Zealand	1994–2007	–4.1 percent
Portugal	1981–1992	–3.3 percent
Spain	1995–2007	–6.9 percent
United Kingdom	1987–2007	–4.4 percent

from U^* . Assuming equation (1) holds, we should see rising inflation if strong demand is pushing the NAIRU down, and falling inflation if the NAIRU is rising. That is, inflation and the NAIRU should move in opposite directions.

The implications for inflation are different if real factors, such as changes in productivity growth or in labor-market distortions, cause changes in the NAIRU. In this case, one possibility is that the central bank adjusts U to keep it near U^* . If that happens, inflation remains stable as U^* changes.

Another possibility is that actual unemployment lags behind changes in the NAIRU. In this case inflation moves in the same direction as the NAIRU, the opposite of the co-movement predicted by hysteresis theories. Orphanides (2000) argues that this happened in the United States in the 1970s. The NAIRU rose but policymakers did not recognize the change, so they tried to hold unemployment at the old NAIRU. With U below U^* , inflation rose.

In examining inflation behavior, as with unemployment, I look for large changes. I identify major disinflations, defined as a fall of at least 3 percent in “trend inflation.” Following Ball (1994, 1999), trend inflation is measured by a nine-quarter centered moving average of inflation. Similarly, I identify major inflation run-ups, defined as increases in trend inflation of at least 3 percent. I ask whether episodes of large changes in the NAIRU are associated with large disinflations or inflation run-ups.³

For each of the 17 episodes of major NAIRU changes, table 7.2 shows the disinflations and inflation run-ups that occurred within the episode or overlapped with it significantly. The table gives the dates and sizes of the inflation movements. A given NAIRU-change episode includes from zero to three inflation-change episodes.

What do we learn from table 7.2? Let’s first examine the episodes of increasing NAIRUs. In six of these eight episodes, there was a significant disinflation, and no inflation run-up. The other two cases, Sweden and New Zealand, have the pattern of a disinflation followed by an inflation run-up followed by another disinflation. In both of these cases, each of the disinflations is larger than the intervening run-up, and the total change in inflation over the three periods is highly negative (–9.2 percent in Sweden and –14.7 percent in New Zealand). I interpret these two

Table 7.2
Major Inflation Changes During Changes in the NAIRU

EPISODES OF NAIRU INCREASES

NAIRU-Change Episode	Major Changes in Inflation
Finland 1980–1986	-8.2 percent, 1981–1986
France 1980–1996	-10.4 percent, 1981–1987
Germany 1980–2007	-5.9 percent, 1981–1986
Ireland 1980–1989	-16.7 percent, 1981–1987
Italy 1980–1996	-14.4 percent, 1980–1987
New Zealand 1980–1994	-8.9 percent, 1980–1983 +8.6 percent, 1983–1985 -14.4 percent, 1985–1992
Spain 1980–1995	-5.2 percent, 1989–1997
Sweden 1983–1999	-8.2 percent, 1980–1986 +5.6 percent, 1986–1990 -6.6 percent, 1990–1993

EPISODES OF NAIRU DECREASES

Australia, 1994–2007	-3.1 percent, 1995–1998 +3.8 percent, 1998–2001
Finland, 1996–2007	
Ireland, 1989–2007	+3.9 percent, 1998–2001 -3.2 percent, 2001–2004
Italy, 1996–2007	
The Netherlands, 1988–2007	+4.4 percent, 1986–1989 +3.8 percent, 1997–2000
New Zealand, 1994–2007	
Portugal, 1981–1992	+8.8 percent, 1980–1984 -17.5 percent, 1984–1987 +4.0 percent, 1987–1989
Spain, 1995–2007	
United Kingdom, 1987–2007	+5.6 percent, 1986–1989 -7.1 percent, 1989–1993

countries as having disinflationary regimes overall, despite an interruption in disinflation.⁴

I therefore count all eight episodes of NAIRU increases as involving disinflations. One way to put the result is that a major NAIRU increase is sufficient to tell us that a country experienced a major disinflation:

NAIRU Increase → Disinflation,

where the arrow does not indicate causality, but rather sufficiency in the sense that if you find an episode with a NAIRU increase, it is always an episode with a major disinflation. To put the same result a different way, a major disinflation is a necessary condition for a NAIRU increase. Note that the reverse result does not hold: a disinflation is *not* sufficient for a NAIRU increase (equivalently, a NAIRU increase is not necessary for disinflation). Many countries in the sample experienced major disinflations without the NAIRU rising by 3 percent. In some countries, such as the United States and Norway, disinflation occurred with almost no change in the NAIRU.

Now let's examine decreases in the NAIRU. Here the story is more complex.

Of the nine NAIRU-decrease episodes, five include at least one inflation run-up. One of these five episodes, in the Netherlands, includes two run-ups and no disinflations. The other four include a disinflation as well as a run-up. However, in contrast to the cases of Sweden and New Zealand, the inflation run-ups and disinflations are similar sizes. In Australia, Ireland, and the United Kingdom, the inflation run-up and disinflation (which are always contiguous) sum to 0.7 percent, 0.7 percent, and -1.5 percent, respectively. Portugal is a special case of volatile inflation: there are *two* inflation run-ups with a large disinflation in-between. The total inflation change over these episodes is -4.7 percent.⁵

Overall, I interpret these five episodes as consistent with hysteresis theories. In each case, the fall in the NAIRU produced a major inflation run-up at some point, suggesting demand expansions. These demand expansions reduced the NAIRU because they were not overwhelmed by much larger disinflations, as in Sweden and New Zealand.

The evidence shows, however, that reducing the NAIRU does not require a *permanent* increase in inflation. This is most clear in Ireland

and the United Kingdom, where an inflation run-up was followed by a disinflation of similar magnitude. A successful theory of hysteresis will need to explain this pattern.

Four countries have decreases in the NAIRU with neither inflation run-ups nor disinflations: Finland, Italy, New Zealand, and Spain. Notice that, in all four cases, the episodes of falling NAIRUs followed large NAIRU increases, and only partly reversed these increases. The decreases look like some kind of mean reversion. One interpretation is that hysteresis effects are long-lived but not permanent. Tight monetary policy causes a rise in unemployment that lasts a long time, but eventually unemployment starts falling even if inflation is stable.

Note that four of the NAIRU decreases in table 7.1 were *not* preceded by large NAIRU increases. These four episodes are among the five in which a NAIRU decrease was accompanied by a run-up in inflation. So the data suggest that a rise in inflation is necessary for reducing the NAIRU if mean reversion is not at work. We can summarize the results with

NAIRU Decrease \rightarrow Previous NAIRU Increase *or* Inflation Run-up,
capturing the fact that all NAIRU decreases involve at least one of the factors on the right of the arrow.

We can also look at the inflation run-up/NAIRU relationship from the other direction. Table 7.3 lists all episodes of inflation run-ups since 1980—those included in table 7.2 and those not included in table 7.2—because they did not coincide with major changes in the NAIRU. The episodes are ranked by the size of the inflation increase.

I want to argue that inflation run-ups are associated with decreases in the NAIRU. That is not true for all of the run-ups in table 7.3, but I have good excuses for discounting some of these cases. The two with asterisks are the Swedish and New Zealand episodes in which inflation run-ups interrupt regimes that are disinflationary overall. In the two cases with double asterisks, in Japan and Switzerland, a 3 percent decrease in the NAIRU was impossible because the NAIRU was less than 3 percent when inflation started to rise.

That leaves nine inflation run-ups, and seven of them occurred during periods of NAIRU decreases. The two that did not are the two smallest inflation run-ups on the list—early run-ups in Australia and Finland. So, among inflation run-ups that were not sandwiched between big disinfla-

Table 7.3
All Inflation Run-ups, 1980–2007

Portugal 1980–1984	8.8 percent
New Zealand 1983–1985	8.6 percent *
United Kingdom 1986–1989	5.6 percent
Sweden 1986–1990	5.6 percent*
Switzerland 1986–1990	4.7 percent**
The Netherlands 1986–1989	4.4 percent
Portugal 1987–1989	4.0 percent
Ireland 1998–2001	3.9 percent
Japan 1987–1990	3.9 percent**
Australia 1998–2001	3.8 percent
The Netherlands 1997–2000	3.8 percent
Australia 1984–1986	3.2 percent
Finland 1986–1989	3.2 percent

* Preceded and followed by larger disinflations

** Initial NAIRU <3 percent

tions, and where the NAIRU was not below 3 percent initially, the seven largest run-ups occurred during episodes of NAIRU decreases. To a first approximation we can say

Inflation run-up → Decrease in NAIRU.

With some qualifications, an inflation run-up is sufficient for a NAIRU decrease (or a NAIRU decrease is necessary for an inflation run-up).

To summarize, the patterns we see in these data are complex. It appears, however, that there are relationships of some type among large rises and falls in the NAIRU and large rises and falls in inflation. These relationships generally go in the direction predicted by hysteresis theories. The data are inconsistent with purely real theories of NAIRU changes, which predict either no relationship between NAIRU changes and inflation or a positive relationship.

5. Open Questions

While there is evidence that hysteresis exists, there are many open questions about the nature of the phenomenon.

What Mechanism?

Why might hysteresis exist? In introducing the concept, Blanchard and Summers explained it with an insider-outsider model of wage bargaining. These models have not been popular in recent years, however, and there may be good reason for this. There is not much empirical evidence for insider-outsider models. In particular, these models suggest that the degree of hysteresis should depend on wage-setting institutions, and that does not seem to be the case. For example, my 1997 paper finds no link between hysteresis and a country's level of unionization.

A more promising idea, which Blanchard and Summers discuss but de-emphasize, involves the behavior of the long-term unemployed. The key idea is that these workers become detached from the labor market, both because they appear unattractive to employers and because they do not search vigorously for jobs. Consequently, while a high level of short-term unemployment puts downward pressure on wage inflation, a high level of long-term unemployment does not.

If this effect is strong, then it potentially explains hysteresis. One story is that a decrease in aggregate demand initially causes a rise in short-term unemployment, but this turns into long-term unemployment if the slump continues. The initial short-term unemployment causes inflation to fall, but then inflation stabilizes. At that point the NAIRU is higher because of the large pool of long-term unemployed.

This story is lent plausibility by evidence (in both my 1997 and 1999 papers) that a long duration of unemployment benefits magnifies hysteresis. Presumably it is more likely that the long-term unemployed become detached from the labor market if they can live on the dole indefinitely.

The story is also consistent with Llaudes (2008), who estimates Phillips curves with separate terms for long-term and short-term unemployment. For many countries, Llaudes finds that long-term unemployment has smaller effects on inflation. This result is stronger in countries with long-lived unemployment benefits.

Yet current stories about hysteresis mechanisms are speculative. More research is needed. In particular, researchers should directly examine the idea that the long-term unemployed become detached from the labor market. One method would be interviews of the type in Bewley

(1999). Researchers could ask employers about their attitudes toward the long-term unemployed, and ask the unemployed about their search behavior.

Nonlinearities and State Dependence

In explaining the idea of hysteresis to students, I sometimes combine the Phillips curve, equation (1), with

$$(2) \quad U^* = (1 - \mu)U^*_{-1} + \mu U_{-1}.$$

Here, the NAIRU is pulled toward actual unemployment. The parameter μ measures the degree of hysteresis.

Empirically, however, it is clear that no such linear relationship exists. Changes in U sometimes cause changes in U^* and sometimes do not. It seems to depend on the past history of U^* and the length of time that U is pushed away from U^* . Hysteresis also appears asymmetric (for example, an inflation run-up means it is very likely U^* is falling, while disinflations often occur without U^* rising).

As usual, it is difficult to measure nonlinearities precisely. And our hazy understanding of hysteresis mechanisms means theory does not give us much guidance. However, there are promising avenues for research.

In particular, there should be more work examining the time-series behavior of short-term and long-term unemployment. Suppose, as suggested by Llaudes's work, that long-term unemployment puts less pressure on inflation than does short-term unemployment. Then we can learn about the varying effects of U on U^* by examining the evolution of U of different durations. For example, we can directly check whether NAIRU increases are tied to shifts from short-term to long-term unemployment.

We also might better understand why some countries reduce the NAIRU without significant effects on inflation, while inflation rises in other cases. Perhaps in some countries a demand expansion cuts into long-term unemployment without much effect on short-term unemployment. Elsewhere, a NAIRU decrease involves falling short-term unemployment, either because there is less long-term unemployment initially or because demand expands more rapidly. In this case, the effects on inflation are likely to be larger.

Policy Implications

If hysteresis exists, a broad lesson is that it is dangerous for central banks to focus policy too heavily on inflation, either through explicit inflation targeting or otherwise. If the natural rate is independent of monetary policy, then focusing on inflation can at worst exacerbate short-run unemployment movements. With hysteresis, by contrast, a given inflation target is consistent with more than one level of unemployment, even in the long run. A central bank might achieve its inflation target but create needlessly high unemployment in the process.

A closely related point is that policy should ease when a recession occurs. This principle might seem like common sense, and the Federal Reserve has followed it (Romer and Romer 1994), but not all central banks have. Ball (1999) finds that inadequate responses to recessions have contributed to hysteresis in some countries.

One can dream up more novel ideas for policy based on the types of hysteresis effects that seem to exist. For example, maybe central banks facing high unemployment should expand demand, accepting a rise in inflation to reduce the NAIRU. Then they should tighten policy to reduce inflation, but reverse the tightening quickly, before a temporary rise in unemployment can push the NAIRU back up.

However, central banks generally presume that steady policies are better than tricky plans for first overheating and then underheating the economy. We would need much greater confidence in our understanding of hysteresis to give contrary advice.

6. Conclusion

In the last decade, mainstream economists have not paid much attention to the idea of hysteresis. Likely reasons include the theoretical appeal of long-run neutrality and our weak understanding of hysteresis mechanisms. In addition, many economists interpret the 1960s and 1970s as showing that it is dangerous for central banks to target unemployment. Hysteresis stories evoke negative reactions because they seem like a step back toward the bad old days.⁶

Yet there is considerable evidence that hysteresis is an important factor in unemployment behavior. And there are clear avenues for research, for example using data on short-term and long-term unemployment. I hope hysteresis becomes a more popular topic in the future.

■ *I am grateful for research assistance from Sandeep Mazumder and for comments from V.V. Chari, Jordi Galí, Engelbert Stockhammer, two anonymous referees, and conference participants.*

Notes

1. This is an old-fashioned backward-looking Phillips curve, replaced in much modern research by the forward-looking New Keynesian Phillips curve. This paper is based on the premise that the old Phillips curve is a useful framework. The relative merits of old and new Phillips curves can be debated elsewhere.
2. Another promising sign is Stockhammer and Sturn (2008), which updates and extends the analysis in Ball (1999).
3. An alternative would be to measure the total change in trend inflation over the NAIRU-change episode. One problem with this approach is that the results would be sensitive to the dating of starts and ends of episodes. These dates are hard to pin down with confidence, as they depend on how the series for unemployment is smoothed.

In addition, simply examining total inflation changes would hide the fact that significant fluctuations in inflation can occur within a NAIRU-change episode. As shown in table 7.2, some episodes include both a disinflation and an inflation run-up.

4. In New Zealand's case, the seesaw pattern of inflation may reflect wage and price controls, which were introduced in 1982 and lifted in 1984.
5. A referee suggests that Australia's inflation run-up was caused by the introduction of a sales tax. However, the tax was introduced in July 2000, and most of the run-up occurred before then. My measure of trend inflation rose from 0.6 percent in 1998:Q1 to 4.2 percent in 2000:Q2. Over the same period, a backward-looking four-quarter average of inflation rose from -0.2 percent to 3.2 percent.
6. Another factor is that Blanchard and Summers have been poor stewards of their hysteresis idea. Summers has been busy with other things. Blanchard has written extensively about unemployment since 1985, but much of his work explicitly or implicitly denies the existence of hysteresis. For example, Blanchard and Wolfers (2000) take it as given that shifts in aggregate demand affect actual unemployment but not "equilibrium" unemployment. When even the creator of an idea does not seem to believe it, the idea loses credibility.

References

- Ball, Laurence. 1994. "What Determines the Sacrifice Ratio." In *Monetary Policy*, ed. N. Gregory Mankiw, 155–193. Chicago: University of Chicago Press.
- Ball, Laurence. 1997. "Disinflation and the NAIRU." In *Reducing Inflation: Motivation and Strategy*, ed. Christina D. Romer and David H. Romer, 167–194. Chicago: University of Chicago Press.

- Ball, Laurence. 1999. "Aggregate Demand and Long-Run Unemployment." *Brookings Papers on Economic Activity* 2: 189–236.
- Ball, Laurence, and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives* 16(4): 115–136.
- Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.
- Blanchard, Olivier. 2005. "European Unemployment: The Evolution of Facts and Ideas." Working Paper No. 11750. Cambridge, MA: National Bureau of Economic Research.
- Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual*, ed. Stanley Fischer, 15–90. Cambridge, MA: The MIT Press.
- Blanchard, Olivier, and Justin Wolfers. 2000. "The Roles of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence." *Economic Journal* 110(462): C1–C33.
- Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review* 58(1): 1–17.
- Krugman, Paul. 1994. "Past and Prospective Causes of High Unemployment." *Federal Reserve Bank of Kansas City Economic Review* 79(4): 23–43.
- Llaudes, Ricardo. 2008. "Unemployment and Inflation in Europe." Ph.D. Dissertation, Johns Hopkins University.
- Mankiw, N. Gregory. 2007. *Macroeconomics* 6th ed. New York: Worth Publishers.
- Nickell, Stephen, Luca Nunziata, and Wolfgang Ochel. 2005. "Unemployment in the OECD Since the 1960s: What Do We Know?" *Economic Journal* 115(500): 1–27.
- Orphanides, Athanasios. 2000. "The Quest for Prosperity without Inflation." *Journal of Monetary Economics* 50(3): 633–663.
- Romer, Christina D., and David H. Romer. 1994. "What Ends Recessions?" In *NBER Macro Annual*, ed. Stanley Fischer and Julio J. Rotemberg, 13–80. Cambridge, MA: The MIT Press.
- Siebert, Horst. 1997. "Labor Market Rigidities: At the Root of Unemployment in Europe." *Journal of Economic Perspectives* 11(3): 37–54.
- Stockhammer, Engelbert, and Simon Sturn. 2008. "The Impact of Monetary Policy on Unemployment Hysteresis." IMK Working Paper 15-2008. Available at http://www.boeckler.de/pdf/p_imk_wp_15_2008.pdf.

Appendix: Estimating the NAIRU

To estimate the NAIRU, Ball and Mankiw (2002) first estimate the parameter α in

$$(3) \quad \pi = \pi_{-1} + \alpha(U - U^*) + \varepsilon,$$

which is equation (1) with an error ε , which we interpret as a short-run supply shock. We estimate α by OLS, treating U^* as a constant.

Rearranging equation (3) gives us

$$(4) \quad U^* - (1/\alpha)\varepsilon = U - (1/\alpha)(\pi - \pi_{-1}).$$

We construct the right side of this equation from the estimated α and data on unemployment and inflation, giving us the left side. This expression, $U^* - (1/\alpha)\varepsilon$, is the NAIRU minus a term proportional to the supply shock. We smooth this series with the Hodrick-Prescott filter to get NAIRU estimates.

The Ball-Mankiw procedure is internally inconsistent because it estimates a time-varying U^* , but assumes a constant U^* to estimate α . Here I resolve this inconsistency with an iterative procedure. Once I have a series for U^* , I use that series to reestimate equation (3), yielding a new estimate of α . I use the new α to estimate a new series for U^* , and so on until the results converge to an α and a U^* series that are consistent.

This procedure is applied to data from 1975 through 2007. (I only use NAIRU estimates for 1980–2007, but I start the estimation in 1975 to minimize endpoint problems.) I use a λ parameter of 100 in the HP filter.

Comments on “Hysteresis in Unemployment” by Laurence Ball

V.V. Chari

I'll start with a general remark that paraphrases something Bob Lucas said that I liked. The idea goes something like this: the observation that prices rise more rapidly in periods of economic expansion than they do in periods of economic contraction is an ancient one. The policy idea that permanent inflation induces a permanent economic high is no doubt equally ancient. The interesting question for historians of economic thought is the mysterious transformation of this policy idea from obvious fallacy to the cornerstone of economic policy.

Larry Ball's paper contains two basic ideas. The first is a second-generation Phillips curve, which relates *changes* in the inflation rate to the *level* of the unemployment rate—deviations from the natural rate or the non-accelerating inflation rate of unemployment (NAIRU). The second idea is that monetary policy has extremely persistent effects on the NAIRU, well beyond effects that take place over the relevant business cycle.

So the messages of Ball's paper are crystal clear and very interesting. He wants to argue that monetary policy has large, long-run real effects and he wants to argue that there is a long-run Phillips curve that is policy-invariant. The problem, as Ball understands it, is that the standard models that we write down have exactly the opposite implication. Monetary policy typically has small long-run effects, not zero. There are always distortions caused by inflation. These distortions could raise or lower long-term unemployment or output. Therefore, the long-run Phillips curve is not policy-invariant.

Both of Ball's ideas come from statistical relationships that he thinks he finds in the data, together with a presumption that these statistical

relationships are invariant with respect to policy. This is the kind of paper that forces us back to some basic facts about the data and requires us to restate why mainstream macroeconomists use as a starting point models with certain basic features: 1) business cycles are fluctuations around a balanced growth path, 2) sustained high inflation has only modest effects on the level of output in the balanced growth path, and 3) there is no exploitable short-run Phillips curve.

I'm going to focus my discussion on a simple question: where do these standard messages come from? Do these come from the fact that we just don't look at the data correctly, or we just don't think about things hard enough? I want to argue that these standard messages, which come from classic dynamic general equilibrium models, come about because those dynamic general equilibrium models have certain features that are built into the model not by accident, not by chance, but rather because the data seemed to suggest that our models need to incorporate these features. I will focus on three features that are present in our models, and I will argue that a lot of these features are driven by what we think we see in the data.

My first observation is that there seem to be fairly small, if at all permanent, effects of fluctuations. The data seem to suggest that the effects of inflation in the long run are modest, and that there is no stable short-run Phillips curve. This is kind of basic, it is in Greg Mankiw's intermediate macroeconomics textbook, and all of you know this. But it is useful to start with the basics.

Figure 7.2 shows a plot of U.S. gross domestic product (GDP) going back to 1929. I didn't use Kendrick's data to go back further, but it is very similar if you go further back. When I look at this picture, I see two things. First, I see obviously big fluctuations like the Great Depression and World War II. Second, I see remarkably stable long-run growth. If you put a straight 45-degree line through this growth chart, it just looks remarkably stable. Figure 7.3 focuses entirely on the post-World War II era, and it does not look like business cycle fluctuations have permanent or persistent effects.

Why do we like models which have these kinds of features, like a balanced growth path and a deterministic steady state and fluctuations




Figure 7.2
 U.S. Real GDP Per Capita, 1929–2006 (logarithmic scale, chained to 2000 dollars)
 Source: U.S. Bureau of Economic Analysis.




Figure 7.3
 Postwar U.S. Real GDP Per Capita, 1947:Q1–2007:Q1 (logarithmic scale, chained to 2000 dollars)
 Source: U.S. Bureau of Economic Analysis.

around that balanced growth path? Figure 7.4, a plot of the capital-output ratio, tells you why—this is a remarkably stable relationship. Figures 7.2 to 7.4 all illustrate well-known facts, but I think it is worthwhile to remind ourselves of well-known basic facts.

Figure 7.5 is a plot of output per hour and compensation per hour. Output per hour is the black line. The labor compensation per hour is the grey line. What I want you to get out of this is that you cannot really see the sector depicted by the black line or the grey line, as these sectors track each other remarkably well. Figure 7.6 is a plot of the U.S. employment rate, the fraction of people over the age of 16 years who are employed relative to all people above the age of 16. What I see here is some stability and some increase, particularly starting in the 1980s, when the employment rate rose fairly markedly.

All of these figures show that to a first approximation, it makes sense to think of the United States as being on a balanced growth path with modest (other than the Great Depression) fluctuations around this path. Along this path, the capital-output ratio is remarkably constant and the employment rate (measured as the number of people working as a fraction of the population) has shown some modest changes. Labor compensation per hour tracks output remarkably well. Not displayed, but another well-known fact about the U.S. economy is that the real return to capital (measured from the National Income and Product Accounts) is remarkably constant. All this has occurred over decades in the face of remarkably different inflation experiences.

Is there striking evidence regarding the association between inflation and real activity? The second observation I will make is that in most of our models, for good or for ill, the effects of inflation in the long run are modest. Why?

Figure 7.7 is a plot of five-year averages of inflation and real output growth for a sample of 15 countries over the past 180 years, from 1820 through 2000. This is a plot for essentially all the currently industrialized countries or a large fraction of them. What I see here is no systemic relationship between inflation and real activity. Note that this period from 1820 to 2000 encompasses dramatically different monetary regimes, ranging from the gold standard to fixed exchange rate regimes to partially floating exchange rate regimes to pure floating regimes.




Figure 7.4
U.S. Capital-Output Ratio, 1951–2006
Source: U.S. Bureau of Economic Analysis.




Figure 7.5
U.S. Output Per Hour and Labor Compensation Per Hour, 1959:Q1–2004:Q4
Source: U.S. Bureau of Economic Analysis, U.S. Census Bureau.




Figure 7.6
U.S. Employment Rate, 1947–2006
Source: Economic Report of the President.




Figure 7.7
Five-Year Averages of Inflation and Real Output Growth in 15 Industrialized Countries, 1820–2000
Source: Adapted from Atkeson and Kehoe (2004).

My third point focuses on whether there is striking evidence that when inflation rates are higher than average, unemployment rates are systematically below average. Figure 7.8 (labeled the first-generation Phillips curve) provides no evidence of such an association. This figure depicts four representative industrialized countries.

The starting point for modern macroeconomics is some version of a Solow, Cass-Koopmans, Kydland-Prescott growth model. It is not blind resistance to obvious fact, or stubborn adherence to shopworn ideology that makes such a model a starting point. Rather, it is an attempt to distill what we know from this collection of figures and tables into an abstraction that is suitable for policy analysis. The key feature of modern business cycle models is the concept of business cycle fluctuations as persistent fluctuations that ultimately revert to a balanced growth path. In many such models, monetary policy plays an important role over the business cycle, but has only a small effect on the growth path itself.

Every intelligent student in my first-year graduate course in macroeconomic theory starts off being disappointed with the class of business cycle and growth models we all study as being too limiting a class. Every graduating Ph.D. student ends up admiring the great economists who created these powerful abstractions for having the ability to put a collection of disparate facts under the lens of this class.

What about the statistical relationships that Ball documents? I start with the relationship between changes in inflation rates and unemployment rates. This statistical relationship belongs to a class that I call second-generation Phillips curves—that is, statistical relationships of the form

$$(1) \quad \pi_t - \pi_{t-1} = \alpha(\mu_t - \mu^*).$$

Figure 7.9 plots the left side of equation (1) against its right side for the United States over the postwar period. The coefficient α appears to be negative, and indeed a regression of $\pi_t - \pi_{t-1}$ on μ_t yields a (statistically) significant negative value of the coefficient α . By now, we should be wary of statistical relationships that are thought to be invariant to policy. One way of asking whether the relationship is invariant to policy is to examine whether this relationship holds for other countries, many of which seem to pursue very different monetary policies than does the United

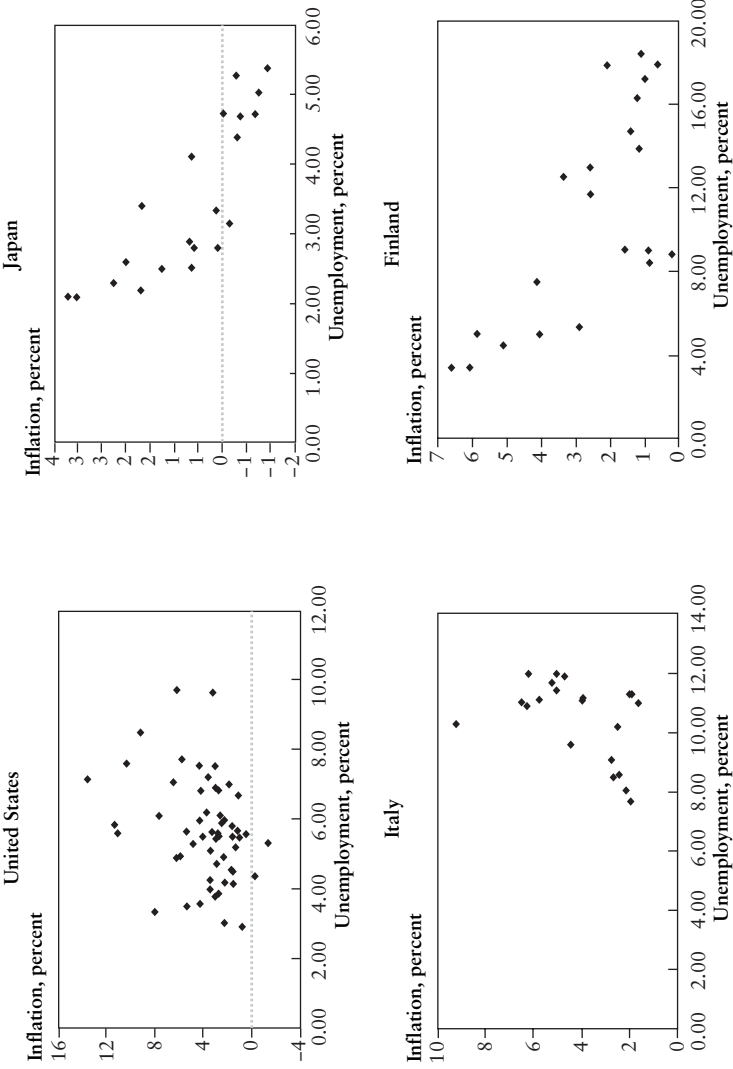


Figure 7.8 Changes in Inflation versus Unemployment Under First-Generation Phillips Curve, 1980–2004
Source: International Monetary Fund, International Financial Statistics.




Figure 7.9

U.S. Postwar Changes in Inflation versus Unemployment Under Second-Generation Phillips Curve

Source: International Monetary Fund, International Financial Statistics.

States. Figure 7.10 repeats a plot of $\pi_t - \pi_{t-1}$ against unemployment for essentially all the industrialized countries. I see no stable relationship here between inflation and unemployment.

A different way of asking whether rising inflation raises economic activity is to ask whether this relationship passes the smell test. If this relationship were robust, Zimbabwe today or Brazil and Argentina during their hyperinflations should have become incredibly prosperous. Most economists think of ever-rising inflation rates as a threat, not a boon, to prosperity, precisely because they are well aware of how damaging hyperinflations have been in a bewildering number of countries.

So what is the bottom line? As contemporary macroeconomists have a bunch of standard kinds of models, which are variants of models that the heroes of our profession, people like Solow, Cass and Koopmans, and Kydland and Prescott, developed after much painstaking work.




Figure 7.10
 Change in Inflation versus Unemployment, 1980–2004
Source: International Monetary Fund, International Financial Statistics.

When I teach my first-year graduate macro class, the smart students in the class, when first introduced to these models, say “God, what a limiting class.” By the time they graduate with their doctorates, they say “God, the guys who wrote down these models really knew how to write down abstractions that capture the essential features of the data in a compact, wonderful way.” And these models work in a way that is useful for policy analysis. These are the standard models we have, and we like them, not because we are blind to obvious fact, not because we are stubborn ideologues, but because we are humble applied theorists.

References

- Atkeson, Andrew, and Patrick J. Kehoe. 2004. “Deflation and Depression: Is There an Empirical Link?” *American Economic Review* 94(2): 99–103.
- Economic Report of the President*. 2007. Washington, DC: U.S. Government Printing Office.

Comments on “Hysteresis in Unemployment” by Laurence Ball

Jordi Galí

Ball's paper deals with a very important subject, namely, the possibility that monetary policy may have permanent (or nearly permanent) effects on unemployment. Ball casts this question in terms of the NAIRU (the non-accelerating rate of unemployment) and the degree to which the latter may vary in response to changes in unemployment itself, a phenomenon known as “hysteresis,” and originally put forward by Blanchard and Summers (1986) as a possible explanation for the behavior of European unemployment. (One may want to view this idea as one particular aspect of a more general question, namely, that regarding the long-run neutrality of monetary policy.)

Ball expresses some frustration at the little attention given to hysteresis in recent years. It is hard to disagree with him on this point. There has been little empirical or theoretical work on unemployment hysteresis. The workhorse New Keynesian model widely used for monetary policy analysis does not allow for hysteresis or other long-run non-neutralities of monetary policy. In fact, standard versions of that model do not even incorporate involuntary unemployment. Given that the presence of any hysteresis effects is likely to have important implications for monetary policy design, their absence in those models is worrisome if those effects are empirically relevant. One possible explanation is that those models have been originally developed as representations of the U.S. economy, where hysteresis effects do not seem so important. But to the extent that those models are being developed and used in the euro area and other economies, this absence should be more of a concern. A piece of good news is that progress has been made in recent years in incorporating

unemployment explicitly into optimizing monetary models with nominal frictions. Using those models to understand the mechanisms that can generate the highly persistent fluctuations observed in unemployment should be a fruitful research avenue in the upcoming years.¹

The objective of Ball's paper is to provide some additional evidence on the presence and importance of hysteresis effects in OECD countries. As a framework of reference for that exploration Ball adopts a traditional, accelerationist Phillips curve of the form

$$(1) \quad \pi_t = \pi_{t-1} - \alpha (u_t - u_t^*)$$

where π_t denotes inflation, u_t is the unemployment rate, and u_t^* is the natural rate of unemployment. This model is used to estimate a time series for the natural rate for each country, which is then used to identify episodes of *large* changes in the natural rate. Under the conventional view, tracing back to Friedman and Phelps, changes in u_t^* are the result of structural changes in labor markets unrelated to aggregate demand. Thus, we would expect any change in u_t^* to be accompanied by a change in inflation in the same direction, at least under the maintained assumption that u_t does not vary in that case by more than the natural rate itself, which seems reasonable (though not strictly necessary!). On the other hand, the hysteresis hypothesis implies that changes in u_t^* may be a consequence of a change in u_t resulting from variations in aggregate demand conditions, in which case we would expect changes in inflation and the natural rate to have the opposite sign (again, under the plausible assumption that u_t^* varies less than u_t in this case). Ball's analysis of 17 episodes of large changes in the natural rate and the corresponding changes in inflation point to a clear prevalence of co-movement signs that one would expect if hysteresis is the main factor behind large changes in the natural rate. That finding would seem to warrant Ball's call for further research on the nature and mechanisms behind the hysteresis phenomenon.

My comments to Ball's paper are organized in three parts. First, I review and discuss Ball's measure of the natural rate of unemployment. After that, I suggest possible alternative approaches one could take to evaluate the relevance of the hysteresis hypothesis. I conclude with some final thoughts.

Ball's Natural Rate Measure

Even if we take (1) as an accurate description of the relationship between unemployment and inflation, a basic identification problem plagues any attempt to come up with measures of the natural rate, since $u_t^* = u_t + \frac{1}{\alpha} \Delta\pi_t$ cannot be directly computed using data on unemployment and inflation if α is not known (which is the case). Some assumptions have to be made in order to overcome this identification problem. Ball's approach, described in an appendix, also relies on some assumptions. Unfortunately, the latter are not discussed explicitly in the text. I will fill in that gap next.

Ball (plausibly) assumes that (1) doesn't really hold exactly, but instead we have

$$(2) \quad \pi_t = \pi_{t-1} - \alpha (u_t - u_t^*) + \varepsilon_t,$$

where ε_t represents a (possibly serially correlated) supply shock. Then, and given an estimate of α (denoted by $\hat{\alpha}$ and obtained as described below), Ball computes the natural rate of unemployment as the Hodrick-Prescott (HP) trend of $u_t + \frac{1}{\hat{\alpha}} \Delta\pi_t$. Note, however, that (2) implies $u_t + \frac{1}{\hat{\alpha}} \Delta\pi_t = u_t^* + \frac{1}{\hat{\alpha}} \varepsilon_t$. Thus, by taking the HP trend of $u_t + \frac{1}{\hat{\alpha}} \Delta\pi_t$ as a measure of u_t^* , Ball is implicitly assuming that the latter variable evolves much more smoothly than the error term ε_t . This could very well be true, but it is not an immediate implication of the theory. If the opposite were true, Ball's natural rate would effectively be measuring supply shocks. Interestingly, however, Ball's findings may justify *ex post* that interpretation since they point to a negative co-movement between (large) changes in the natural rate measure and the change in inflation, which we would not expect if the HP-trend was proxying ε_t instead.

Ball estimates α through an iterative procedure, which consists of regressing $\Delta\pi_t$ on $u_t - u_t^*$, given a series for u_t^* and using the resulting $\hat{\alpha}$ estimate to compute a new series for u_t^* , as described earlier. But in order for this approach to yield a consistent estimate for α , the regressor $u_t - u_t^*$ should be orthogonal to the supply shock, which also seems an arbitrary assumption. Unfortunately no discussion of that assumption and its plausibility is offered in the paper.

Since Ball's subsequent analysis relies heavily on his natural rate measure, its credibility is not independent of that of the above assumptions.

Alternative Approaches

Ball's analysis of the co-movement between his natural rate measure and inflation during episodes characterized by large changes in the former provides a way of testing the hysteresis hypothesis only under the maintained assumption that the Phillips curve (1) is a good representation of the joint dynamics of inflation and unemployment. An alternative, more direct approach would involve a comparison of the size of the changes in unemployment u_t and those in the natural rate u_t^* during the Ball episodes. Under Ball's logic, if changes in the natural rate during those episodes are driven by the changes in the unemployment rate itself we would expect the inequality $|\Delta u| > |\Delta u^*|$ to hold. As shown in figure 7.11, that condition is satisfied in 14 out of 17 of Ball's episodes, meaning in more than 80 percent of the cases. That evidence, which is not distorted by changes




Figure 7.11

Changes in Actual versus Natural Unemployment

Source: Author's calculations based on Ball's data.

in inflation by factors other than those captured by (1), appears to reinforce Ball's findings and conclusions.

Unfortunately, Ball's evidence pointing to significant unemployment hysteresis effects in OECD economies does not shed any light on the mechanisms that may underlie that phenomenon. One of the potential mechanisms put forward by Ball involves the behavior of the long-term unemployed, who may become detached from the labor market and stop searching vigorously for jobs. As a result they may stop putting downward pressure on wages, and this may lead to a permanent increase in measured unemployment. One way to assess the validity of that hypothesis would consist of redoing Ball's analysis from scratch after excluding the long-term unemployed from the unemployment data. If there is no longer evidence favorable to hysteresis once the adjusted unemployment data are used, one would have to conclude indeed that it is through changes in long-term unemployment that hysteresis comes about.

Final Thoughts

I sympathize with Ball's assessment of the insufficient attention that the profession has given to the topic of hysteresis in unemployment. One can think of several reasons for this. First, it is a fact of life that economic research is largely driven by developments in the real world. The relative stability of unemployment fluctuations in the United States and Europe over the past two decades (albeit fluctuations about very different means) may partly explain the diminished interest. But that era of mild fluctuations is likely to come to an end as a result of the current crisis, with unemployment rates bound to skyrocket to levels much higher than the ones we had become used to. When the current downturn comes to an end and growth resumes, natural questions will be raised as to how long it will take to bring the unemployment rate back to the levels that prevailed in recent years (about 5 percent in the United States and 8 percent in the euro area), or about the possibility that unemployment will remain for a long period above those levels. Those questions, spurred by unfolding events, will likely trigger a renewed interest in the subject of hysteresis and the related literature.

Secondly, empirical work on hysteresis is bound to be plagued with all sorts of difficulties. Some of the difficulties are conceptual (for example, how do we define the natural rate of unemployment?). Others are statistical, including the need to disentangle exogenous variations in the natural rate (for instance, whether these are due to demographic factors or exogenous changes in relevant labor market parameters) from those that may have been induced by a change in unemployment itself as a consequence of other, nonlabor-market-related, shocks. Such difficulties make it hard to avoid taking some shortcuts or relying on often questionable assumptions. But the importance of the topic, given the likely large welfare consequences of persistent unemployment fluctuations (and, even more so, of permanent effects of transitory shocks on the level of activity), may make us more tolerant and open to experimentation. Ball's paper in the present volume is a good example of research in that spirit.

Note

1. See, for example, Blanchard and Galí (2008).

References

- Blanchard, Olivier, and Jordi Galí. 2008. "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment." Working Paper No. 13897. Cambridge, MA: National Bureau of Economic Research.
- Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual*, ed. Stanley Fischer, 15–90. Cambridge, MA: The MIT Press.

8

**Lessons for Central Bankers:
A Panel Discussion**

Israeli Monetary Policy and the Phillips Curve

Stanley Fischer

This conference has been extremely valuable for the policymakers—some of us former academics—who are participating, for it provides an opportunity to catch up with the recent literature on the Phillips curve and its implications for monetary policy. The topic is central to our policy decisions, and there has in the last decade been considerable progress in analysis and understanding of the key issues. Thus we are even more than usually grateful to the Boston Fed, to President Eric Rosengren, to Research Director Jeff Fuhrer, and the other organizers for choosing the topic and the presenters of papers, and for the conference's extremely efficient organization.

I will talk about Israeli monetary policy and its relationship to the Phillips curve. The bottom line will be that we have had to deal with almost every difficult issue that has come up in this conference. In particular, we have had to contend lately with the difficulties of pinning down both the expected inflation rate and the output gap, issues that have had an impact on our decisions. In addition, we have suffered from the implications of the Lucas critique as the monetary mechanism has changed in a striking way as the extent of indexation of contracts to the exchange rate has declined sharply as a result of the real appreciation of the shekel in the last year.

First, some background: the Israeli economy is small, with a GDP of about \$200 billion and a population of over 7 million. We are coming off a period of very good economic performance after a deep recession at the beginning of this decade. The economy has grown by over 5 percent each year since 2004, remarkably steadily—indeed the growth rate averaged over 5 percent for the five years from mid-2003 to mid-2008. This

represents a considerable achievement, not least in 2006, the year of the second Lebanese War, which included a quarter of negative growth. To our surprise, growth in the first quarter of this year (2008) was 5.6 percent, despite our having expected to see an impact from the global growth slowdown. We expect growth this year to be 4 percent, a level that is virtually guaranteed because of base period effects.¹ The economy is very open, with trade—the average of imports and exports of goods and services relative to GDP—at about 45 percent.

The economy has a history of very high inflation, which reached nearly 450 percent in 1984. Inflation was stabilized in mid-1985, remained in the range of 20 percent and lower from 1986 on, and during the 1990s was brought down to advanced country levels. An inflation targeting regime was introduced in 1991–1992, soon after Jacob Frenkel became the governor, but the exchange rate, which is regarded as (and is) the key relative price in the economy, continued to be managed within a band until 1997. The band was abandoned—after a struggle with and within the government—because the exchange rate had been pushing on the appreciated side, forcing the Bank of Israel to buy dollars at a time when it did not have an effective means of sterilization, and thus producing real appreciation through inflation. Formally the band was abolished only in 2005, though it had had no practical effect since 1997. Over the decade of the 1990s and into the present decade, capital controls were gradually dismantled, to the point where capital flows in both directions are perfectly free. There is no formal exchange rate target.

As of now the inflation target, set by the government, is specified as a range, 1–3 percent, and the target is headline consumer price index (CPI) inflation. For this decade the average inflation rate has been lower than 2 percent (see figure 8.1), but the 12-month inflation rate has frequently been outside the target range—both below and above—and is difficult to control with any degree of precision. At present (June 2008), we share the problem of most countries of being well above the target range, with inflation running above 4 percent.

We have examined and rejected the case for using a core inflation measure as the target, mainly because we prefer to target a concept that the public understands, which is the CPI, and then to try to explain why




Figure 8.1
 Rate of Inflation in Last 12 Months and Inflation Targets, 1992–2008
 Source: Bank of Israel data.

we have difficulty in hitting it, rather than using a concept that is less understandable and relevant to the public but that we can hit more accurately.

In his introductory remarks for this session, Allan Meltzer asked how we aggregate opinions in making policy decisions. The Bank of Israel law (of 1954 vintage) specifies that the governor makes the monetary policy decision, and we are thus an example of a single decisionmaker model. However, I make the decision in consultation with a committee of relevant department heads of the Bank of Israel, and we publish minutes, including the voting record, though without identifying individual votes. The department heads are very independent, and I have not noticed that we suffer from their being excessively deferential to the views of the governor. On a few occasions I have made a decision opposed by a majority of the department heads, but that is not done lightly. We are currently working on a new Bank of Israel law, which if passed will, among other things, set up a monetary policy committee, with an equal number of inside and outside members, and with the governor having a double vote in the event of a tie.

Why the difficulty in hitting the inflation rate? It is largely a result of the formerly close tie between the exchange rate and prices. As a result of the inflationary history of the Israeli economy, many contracts remained denominated in U.S. dollars long after inflation was stabilized—including, most importantly, apartment and housing rental prices. The monthly payment is specified in dollars as unit of account, and the actual payment is made in shekels at the exchange rate prevailing at the time of the payment. Appropriately, the CPI registers the price as the shekel payment. With housing and rentals accounting for 20 percent of the CPI, there was an almost immediate pass-through of 20 percent from the exchange rate to the price level. Add to that a pass-through of about 15 percent from the prices of imports, and there was a close and rapid link between the exchange rate and the price level, amounting to about one-third within a quarter. In addition, it is likely that there is a later, second round of pass-through, as the prices of imported raw materials and other factors of production begin to affect local prices in a way similar to that in other countries—however, the extent and dynamics of the later effect has been hard to pin down accurately.

As is well known, the exchange rate is very hard to predict, and that situation is no different in Israel. Furthermore, with rapid exchange rate pass-through, unanticipated movements in the exchange rate tend to affect inflation rapidly, and before monetary policy has an opportunity to respond. Until recently, the inflation rate tended to be low or negative when the exchange rate was appreciating, and inflation tended to be high when the currency was depreciating against the dollar.² For instance, as can be seen in figure 8.1, the inflation rate was low or negative in 2004–2005, a period of exchange rate appreciation. Once monetary policy recognized that inflation was moving outside the target range, it tended to adjust in order to deal with inflation, but that sometimes gave the impression that monetary policy was directed at the exchange rate rather than the inflation rate.

Monetary policy tended to work relatively fast on the inflation rate when the pass-through was high, since a change in the interest rate affected the exchange rate, which affected the inflation rate. During the last year the transmission mechanism from monetary policy to inflation has changed as a result of a very rapid reduction in the percentage of rental contracts tied to the dollar. More than 90 percent of rental contracts were specified in dollars about a year ago, and the percentage now is down to about 30 percent. Accordingly the housing price index is affected much less immediately by the exchange rate, and the rapid impact of the interest rate on inflation has been greatly attenuated. During the last year, inflation has risen as it has in the rest of the world, despite a very rapid appreciation of the shekel against the dollar, with much of the impact being traceable to the apparent failure of the appreciation of the shekel to keep the rental price index for housing from rising.

This is the Lucas critique aspect of monetary policy issues in Israel today—an issue which is difficult to deal with in part because the monetary mechanism has changed only recently and we do not yet have sufficient data to generate robust estimates of the monetary policy transmission mechanism.

We use a variety of estimates of the expected inflation rate in setting monetary policy: breakeven inflation rates from the money and bond markets, analyst forecasts, company and consumer surveys, and our own internal forecasts. Figure 8.2 shows the breakeven inflation rate for




Figure 8.2
Inflation Expectations, October 2000–June 2008
Source: Bank of Israel data.

the next 12 months, the average of forecasters' predictions of inflation for the next 12 months, and the five-year breakeven inflation rate. The breakeven inflation rate and the average of the forecasters' expectations are not identical, with the forecasters recently predicting higher inflation than the breakeven rate.³ The facts that the forecasters both differ among themselves and do not regard the breakeven inflation rate as equal to the relevant expected rate of inflation provide an indication of the uncertainties about this variable. Nonetheless we need a measure of expected inflation in making the monetary policy decision, and we examine all the different measures. If we were forced to use only one, we would probably use the breakeven inflation rates derived from the financial markets. However that calculation is not straightforward, because at the end of their lives, in the absence of daily price indexes, price-indexed bonds become nominal bonds for a period of up to eight weeks. Thus the calculation of the breakeven 12-month-ahead inflation rate from the financial markets has to be based on assumptions about the inflation expected for the next month or two.

As can be seen in figure 8.2, the inflation rate expected over the next 12 months has been within the target inflation range of 1–3 percent almost all the time since 2004, though it did briefly fall below the lower bound of the target range early in 2007, when the inflation rate itself was below the target range.⁴ The five-year-ahead breakeven inflation rate rose well above the target range when inflation was high around 2002–2003, gradually declining to within the target range, and apparently stabilizing at around 2.5 percent. We regard this number as including a risk premium, and thus as an indication that the five-year expected inflation rate is close to the 2 percent center of the target range.

We place a heavy weight on the expected rate of inflation in making the interest rate decision. At one point, International Monetary Fund missions used to suggest to the Israeli authorities that they in essence were using a single variable—the expected inflation rate—Taylor rule. Further, since the expected inflation rate was taken to be the breakeven rate observed *in the markets*, the Fund's economists were concerned that the entire inflation process might be unanchored. They urged the Bank of Israel to use econometric models to ensure the consistency of inflation expectations with current and expected policy decisions.

The governor at that time, Jacob Frenkel, argued that there is no indeterminacy when the central bank uses the expected inflation rate as calculated from market behavior in deciding on the interest rate—on the grounds that in using the breakeven inflation rate as a measure of expected inflation, the Bank of Israel was in effect drawing on the variety of models being used by market participants in determining their own estimates of future inflation. As a formal matter, this argument is probably correct provided at least one—or a significant fraction—of the market participants is bringing some information other than the breakeven inflation rate into their calculation of the estimated inflation rate.

In addition to the expected inflation rate, the Bank of Israel also paid attention to the interest rate gap with the United States, under the belief that monetary policy had to maintain a sizable interest rate gap relative to the United States because—as it used to be said—there was a mountain of shekels waiting to get out of the country. Gradually the interest rate was reduced to the point where in 2006, the Bank of Israel interest rate fell below the Fed's rate, as can be seen in figure 8.3. Since then the Bank of Israel interest rate has moved reasonably independently of the Fed rate, though it remains true that in setting the interest rate, we take account of the interest rate gap with the United States. At present the money mountain that is creating problems for monetary policy is not the shekel mountain but the dollar mountain, since capital inflows have helped create a major appreciation of the shekel.

I will now turn to the Phillips curve. The draft of the new Bank of Israel law sets out the aims of monetary policy in a manner very similar to that of the Bank of England and other central banks that have recently modernized their laws. Our primary responsibility is to maintain price stability as defined by the government (currently 1–3 percent), but subject to achieving price stability, we are also expected to contribute to attaining the other goals of government policy—particularly growth and employment—and to financial stability. In making our monetary decisions, we use two models, one a dynamic stochastic general equilibrium (DSGE) model and the other a more standard old-fashioned Keynesian model. Each embodies a Phillips curve-like mechanism, but the DSGE model uses a Hodrick-Prescott filter to calculate the gap, and the




Figure 8.3
 Bank of Israel Interest Rate, Inflation Expectations*, and the Federal Reserve's Interest Rate, 2004–2008
 Source: Bank of Israel data.
 *Note: For 12 months, as derived from the capital market.

Keynesian-type model uses a production function and unemployment rate to determine the full employment labor force.

The two models gave fairly consistent results so long as unemployment was high and clearly well above the full employment level. Now the unemployment rate is close to our estimate of the full employment rate, and we are in a situation in which the DSGE model sees the economy as fluctuating around the full employment rate of unemployment, and the Keynesian-style model has for some time indicated that the economy is at more than full employment. This difference, plus small differences in the way the exchange rate—which remains a key factor in determining inflation—affects inflation in the two models, means that the two models now give significantly different estimates of the interest rate path that is consistent with returning inflation to within the target range within a year. We give ourselves one year—which is short relative to other central banks—to return to the target range, because the Israeli inflation rate is very volatile and we do not have much confidence in our forecast two years out.

These differences are now important. For some months, as the inflation rate rose, we were inclined to say that our inflation was due to the increasing prices of food and energy, and that it was not primarily a demand-side problem. Nonetheless the research department kept warning that we were in a situation of excess aggregate demand, which would soon find expression in the overall inflation rate. In April—the last price level reading we have—the inflation rate, month-on-month, was a major surprise, at 1.5 percent, and the inflation was visible in almost every price group. From this evidence we concluded that the relative price pressures from energy and food had spread through the economy, their path eased by the high level of demand, and decided that we had to begin raising the interest rate.

Which is to say, that the Phillips curve is alive, that it matters how you measure expected inflation, that it matters how you calculate the output gap, and that you have to keep watching to make sure that the Lucas Critique is not creeping up on the structure of your economic models. These are all lessons we have been reminded of by the papers and discussions presented at this excellent conference.

■ *Data and comments in the text use only information available up to the time of the conference, June 2008, though in one or two places later information is noted in footnotes.*

Notes

1. Revised data (published in September 2008) for the second quarter show quarter-to-quarter growth of about 4 percent, and the Bank of Israel's current (September) forecast for the year is about 4.5 percent growth.
2. The dollar exchange rate—and not other exchange rates—was critical, because contracts were indexed to it and not to other exchange rates.
3. In the light of hindsight, the forecasters were more right than the calculated breakeven rate.
4. In addition, the 12-month-ahead expected inflation rate rose above the target range briefly in August 2008, but has subsequently returned to within the range.

Lessons for Central Bankers from a Phillips Curve Framework

Donald L. Kohn

An economic model of inflation is an indispensable input to monetary policy deliberations. A model in the Phillips curve tradition remains at the core of how most academic researchers and policymakers—including this one—think about fluctuations in inflation; indeed, alternative frameworks seem to lack solid economic foundations and empirical support. But the modern Phillips curve differs substantially from versions in use several decades ago; policymakers and academics alike are now attuned to the importance of expectations, the possibility of structural change, and the uncertainty that surrounds our understanding of the dynamics of wage and price adjustment. Moreover, the link between inflation and resource utilization often emphasized in a Phillips curve framework accounts for only a modest part of inflation fluctuations. My comments will focus on how the lessons from recent research on the Phillips curve are helping me think about the influence of fluctuations in the prices of commodities, such as oil, on the outlook for inflation, and the appropriate policy responses to such developments.¹

1. Policy Objectives and a Framework for Analyzing Inflation Fluctuations

The Federal Reserve has been charged with the pursuit of price stability and maximum employment. Price stability is uniquely in the control of the central bank over long periods, and it is a prerequisite for the economy performing efficiently over time. The welfare costs of inflation result from many factors: the potential costs to households and firms

that result from efforts to insure themselves against inflation or from confusion regarding real and nominal prices; distortions to the financial system related to inflation; imperfect indexation of taxes, especially with respect to capital income, and the related distortions to economic activity; and the costs associated with a slow adjustment of nominal prices and wages. The costs of inflation imply that central banks should aim for low measured inflation. Moreover, many of the costs of inflation—such as those associated with misconceptions regarding inflation, efforts to insure against inflation, and distortions to the financial system—are associated with the rate of change in the entire set of prices of goods and services facing households or firms, suggesting that measured inflation should be gauged by the rate of change in a broad set of prices. Accordingly, the Federal Open Market Committee (FOMC) has been emphasizing that it gauges price stability over the long term by the behavior of the overall personal consumption expenditures price index.

The economic framework that helps me think about fluctuations in inflation is based on the substantial body of research on models of price dynamics that has been developed over the past five years. At its heart, the framework is based on the importance of sluggish adjustment in (some) nominal wages and prices to changing economic conditions. This sluggishness undoubtedly reflects a number of factors, which include the costs of adjusting some nominal prices and wages, imperfect information regarding shifts in economic conditions, and learning by firms and households about the structure of the economy, including the setting of monetary policy.

Regardless of its source, the presence of sluggish nominal adjustment brings to the fore three key elements driving wage and price dynamics: inflation expectations, supply shocks, and resource utilization. Because some prices and wages are adjusted only infrequently, both firms and households anticipate the future erosion of real prices and wages by incorporating the expectations they have for inflation into their current price settings and wage demands. As a result, inflation expectations play a critical role in the formation of monetary policy. Moreover, the tendency of some prices to adjust very quickly to changing circumstances in conjunction with sluggish adjustment in other prices and wages implies

that large, sharp price movements, such as a change in the price of oil, lead to relative price distortions throughout the economy; these distortions imply that relative price shocks have important implications for the functioning of the economy.² Finally, fluctuations in resource utilization, through their effects on the costs of production and on firms' desired markups over these costs, are a significant determinant of price and wage decisions. The link from resource utilization to inflation provides a major channel through which monetary policy settings influence inflation: adjustments in the policy interest rate bring about changes in resource utilization, which then influence current and expected inflation.

2. Inflation Forecasting and Commodity Price Shocks

The economic outlook is the prime focus of monetary policy. Because the stance of policy influences economic activity and inflation only with a lag, policymakers must adjust policy to minimize the costs from fluctuations in activity and inflation in the future. The Phillips curve framework is an important input into the forecast for inflation. I will focus specifically on how a shock to the price of oil affects the inflation outlook within this framework. This topic is especially salient of late—commodities prices in general, and perhaps most glaringly the price of oil, have risen sharply over the past year; indeed, oil prices have risen sharply for more than four years.

Consider a sharp rise in the price of oil that primarily reflects a shift in the balance between demand and supply in the global market for oil.³ In the first instance, higher oil prices lead to an increase in the overall level of consumer prices. When thinking about the outlook for future inflation, it is useful to distinguish between oil and other prices. The most significant factor determining oil prices is the current and prospective balance between demand and supply. The aggregate behavior of most other prices, consisting of a large set of nominal prices and wages that adjust slowly, is driven by the factors that enter the Phillips curve—inflation expectations, resource utilization, and supply shocks (in this case, a shock to the price of oil).

In a forecasting context, the Phillips curve framework motivates reduced-form regressions of the rate of inflation for consumer prices, excluding food and energy, against proxies for each key factor. Lagged values of inflation typically proxy for inflation expectations. The deviation of output from potential or of the unemployment rate from its sustainable rate serves as a proxy for resource utilization. And changes in relative prices for energy, food, and imports are traditionally included as measures of supply shocks.⁴ This type of regression is among the most useful tools for forecasting inflation. Nonetheless, its forecast record is far from spotless, and hence I consider the forecasts from such regressions as just one input that helps inform my outlook for inflation.

The results of such exercises imply that, over recent history, a sharp jump in oil prices appears to have had only modest effects on the future rate of inflation. This result likely reflects two factors. First, commodities like oil represent only a small share of the overall costs of production, implying that the magnitude of the direct pass-through from changes in such prices to other prices should be modest, all else equal. Second, inflation expectations have been well anchored in recent years, contributing to a muted response of inflation to oil price shocks. But the anchoring of expectations cannot be taken as given: indeed, the type of empirical exercises I have outlined reveal a larger effect of the price of oil on inflation prior to the last two decades, a period in which inflation expectations were not as well anchored as they are today.⁵

Of course, oil prices have jumped repeatedly in recent years. The (relatively) continuous rise in energy prices since 2003 has been a surprise to me and to most others, at least as best as I can gauge by looking at prices that have been embedded in futures contracts over this period. These contracts currently suggest that the price of oil will flatten out in the period ahead.

Nonetheless, repeated increases in energy prices and their effect on overall inflation have contributed to a rise in the year-ahead inflation expectations of households, especially this year. Of greater concern is that some measures of longer-term inflation expectations appear to have edged up since last year. Any tendency for these longer-term inflation expectations to drift higher or even to fail to reverse over time would have troublesome implications for the outlook for inflation.

3. The Structural Phillips Curve, Commodity Price Shocks, and Monetary Policy

The central role of inflation expectations implies that policymakers must look beyond this type of reduced-form exercise for guidance. After all, the lags of inflation in reduced-form regressions are a very imperfect proxy for inflation expectations. As emphasized in Robert Lucas's critique of reduced-form Phillips curves more than 30 years ago, *structural* models are needed to have confidence in the effect of any shocks on the outlook for inflation and economic activity.⁶

The importance of structural relationships as inputs to the monetary policy process poses a challenge; for instance, there are many "structural" models of nominal price and wage adjustment. Each of these models emphasizes different frictions or imperfections and therefore can have different policy implications, and empirical work has reached different conclusions regarding the merits of alternative models.⁷ As a result, policymakers must look to lessons that are common across alternative specifications and base policy on our current understanding of the most likely important structural factors. Fortunately, I think that many of the models of nominal price and wage adjustment imply similar conclusions regarding the influence of commodity prices on the inflation outlook and the appropriate response of monetary policy.

I will again focus on a sharp jump in the price of oil, reflecting supply and demand in the market for oil. Because many nominal prices and wages are costly and slow to adjust, the efficient allocation of resources is impeded during a transition period in which relative price signals are distorted. For example, the prices of energy-intensive goods need to rise relative to those of less-energy-intensive goods, but this adjustment follows a gradual and asynchronous pattern. Similarly, the equilibrium real wage—the relative price of labor—will tend to be depressed by an oil price shock due to the accompanying adverse movements in the terms of trade and reduction in labor productivity, but the needed wage or price adjustments proceed gradually.⁸ An efficient monetary policy should attempt to facilitate the needed economic adjustments so as to minimize distortions to economic efficiency on the path to achieving, over time, its dual objectives of price stability and maximum employment.⁹

In particular, an appropriate monetary policy following a jump in the price of oil will allow, on a temporary basis, both some increase in unemployment and some increase in price inflation. By pursuing actions that balance the deleterious effects of oil prices on both employment and inflation over the near term, policymakers are, in essence, attempting to find their preferred point on the activity/inflation variance-tradeoff curve introduced by John Taylor 30 years ago.¹⁰ Such policy actions promote the efficient adjustment of relative prices: since real wages need to fall and both prices and wages adjust slowly, the efficient adjustment of relative prices will tend to include a bit of additional price inflation and a bit of additional unemployment for a time, leading to increases in real wages that are temporarily below the trend established by productivity gains.

I should emphasize that the course of policy I have just described has taken inflation expectations as given. In practice, it is very important to ensure that policy actions anchor inflation expectations. This anchoring is critical: as demonstrated by historical experiences around the world and in the United States during the 1970s and the 1980s, efforts to bring inflation and inflation expectations back to desirable levels after these have risen appreciably involve costly and undesirable changes in resource utilization.¹¹ As a result, the degree to which any deviations of inflation from long-run objectives are tolerated, in order to allow the efficient relative price adjustments that I have described, needs to be tempered so as to ensure that longer-term inflation expectations are not affected to a significant extent.

4. Global Demand, Trending Commodity Prices, and Monetary Policy

My remarks so far have concentrated on the factors guiding the monetary policy response to a shock in the prices of commodities like oil that stems from a shifting balance of supply and demand in the specific market for these commodities. Some might think that this focus misses the point in the current context, for at least two reasons. First, it has been suggested that the run-up in the prices of a broad range of commodity prices reflects, in part, global excess demand rather than sector-specific forces. And second, some have suggested that important commodity prices, like

that of oil, may be on a more significant upward trend than is currently embedded in future prices.

It seems highly likely that, over the period since 2003, the rise in commodity prices has reflected strong global economic growth as well as some sector-specific factors, such as geopolitical tensions and other disruptions to the supply of oil.¹² In this regard, I share the views expressed by Chairman Bernanke at this conference, in which he discussed a range of factors that have likely influenced relative commodity prices.¹³

However, the fact that rising relative commodity prices have likely reflected many factors does not, by itself, change the analytical framework that I used to frame policy deliberations. As I highlighted at the beginning of my remarks, the most important drivers of inflation in the model of inflation dynamics I use are relative price shocks, inflation expectations, and the balance between aggregate demand and supply in the United States, as measured by some notion of resource utilization. If a shift in global demand affects both commodity prices and the demand for U.S. goods, the model I have in mind accounts for these influences on inflation through relative price shocks and resource utilization. For example, the rise in the price of oil this year has lowered consumption demand by pinching households' real incomes and likely dampened the growth in labor productivity by trimming energy input; both of these factors have probably contributed to a lower equilibrium real wage, as I described earlier. If the impact on demand from these factors was accompanied by stronger global demand that boosted demand for U.S. goods generally, the forecasts of inflation would need to take this into account. In any event, resource utilization has been slackening, judging from the rise in the unemployment rate and the slow pace of economic growth in the United States, on average, over the past six to nine months.

Some have suggested that the price of oil is on a more significant upward trend than currently appreciated.¹⁴ Such an unanticipated shift in trend would not be embedded in the anticipated rate of change in slowly adjusting nominal prices and wages, implying that an adjustment period with distortions to relative prices would follow. Moreover, there would likely be upward pressure on overall inflation during this period, reflecting the slow response of the rate of change in some nominal prices to the

new trend in the price of oil. This tendency for higher overall inflation could risk a rise in inflation expectations.

An appropriate monetary policy response would share many of the characteristics I discussed earlier. In particular, bringing overall inflation immediately back to the low rate consistent with price stability could be associated with a much higher rate of unemployment for a short time. It may be efficient to allow some adjustment period in which both overall inflation exceeds its desired low level and the unemployment rate is higher than its long-run sustainable level; as before, setting policy in a manner that balances the undesirable effects of a shock to the system on both inflation and employment will tend to be more efficient than setting policy so as to deliver more extreme outcomes in either inflation or unemployment.¹⁵

However, two additional considerations are likely important when considering a shift in trend. First, developments in inflation expectations following a significant shift in the relative price trend of a commodity like oil should be monitored carefully, as our understanding of changes in long-run inflation expectations is limited and shifts in trends are infrequent, potentially implying a greater chance of confusion between relative price trends and overall inflation. Second, it is very important to remember that the costs of inflation in excess of the low rate of measured inflation consistent with price stability over any extended period are significant and reflect a broad range of factors. As I emphasized earlier, economic research into the many costs of inflation has suggested that these costs are associated with the rate of change in a broad set of prices. As a result, a trend in any individual relative price should not, in itself, lead to a change in the desirable rate of measured inflation over the long run.

5. Summary

To reiterate, the Phillips curve framework is one important input to my outlook for inflation and provides a framework in which I can analyze the nature of efficient policy choices. In the case of a shock to the relative price of oil or other commodities, this framework suggests that policy-

makers should ensure that their actions balance the deleterious economic effects of such a shock in the short run on both unemployment and inflation.

Of course, the framework helps to define the short-run goals for policy, but it does not tell you what path for interest rates will accomplish these objectives. That issue is what we wrestle with at the FOMC and is perhaps a subject for a future Federal Reserve Bank of Boston conference.

Notes

1. Michael Kiley, of the Board's staff, contributed to these remarks. The views expressed are my own and not necessarily those of my colleagues on the Board of Governors or the Federal Open Market Committee.
2. Relative price shocks act like supply shocks in the Phillips curve framework for two reasons. First, some prices are sluggish and others are flexible, and the Phillips curve framework focuses on the adjustment of sluggish prices. Second, prices are more likely to adjust to very large shocks, and the skewness on the distribution of relative price disturbances can fluctuate substantially, giving rise to shocks to the Phillips curve. Ball and Mankiw (1995) discuss these issues and argue that the second type of relative price shock is quantitatively very important.
3. The rise in a broad range of commodity prices, as has occurred from time to time recently, would not affect the basic analysis, which rests on the contrasting behavior of flexible and sluggishly adjusting prices. However, the simultaneous rise in many commodity prices might suggest that strong global aggregate demand is playing an important role, which would affect the appropriate setting for monetary policy. See section 4 on Global Demand, Trending Commodity Prices, and Monetary Policy.
4. The amount of related literature is large. The article by Robert Gordon (1998) represents a good example and is relevant when considering the notion of supply shocks generally. Stock and Watson (1999) present a broad interpretation of the empirical Phillips curve in which a large number of macroeconomic indicators are used to forecast inflation.
5. These factors are not the only two that have contributed to a lower effect of oil prices on inflation. For example, the energy intensity of the economy has fallen over time. Research like that in Hooker (2002) and subsequent work has discussed various possibilities in more detail. Blanchard and Galí (2007) suggest that the seemingly muted effect of changes in the price of oil and inflation in recent years has been the result of falling energy intensity, more flexible labor markets, better monetary policy, and good luck.
6. See Lucas (1976). Robert Lucas had emphasized the importance of structural models of the Phillips curve well before his 1976 article; for example, see his

contribution at the conference on wage and price dynamics held at the Federal Reserve in 1970 (Lucas 1972). Michael Woodford (1994) presented an important critique of some research on commodity prices, inflation, and monetary policy in which the role of the Lucas critique was central. In particular, he re-emphasized that a tendency of commodity prices to forecast inflation may not be structural and could break down under alternative policy regimes—a tendency that seemed to be confirmed by Mark Hooker (2002) in his work documenting a break in the link between oil prices and inflation in recent decades.

7. For example, Rudd and Whelen (2007) and Kiley (2007) review a number of models and empirical studies as well as conduct their own empirical analyses; these two studies reach quite different conclusions on the merits of alternative specifications.

8. For example, see the article by Wei (2003) illustrating the effects of a rise in oil prices in a dynamic equilibrium model.

9. Erceg, Henderson, and Levin (2000) discuss the importance of relative price adjustments in the presence of sluggish nominal prices and wages in a general context; Mankiw and Reis (2003) present related results. Bodenstein, Erceg, and Guerrieri (2007) apply this reasoning to oil price shocks.

10. See Taylor (1979).

11. Ball (1994) presents estimates of the costs of disinflation for a variety of countries, including the United States. Kiley (2008) provides evidence on how survey measures of inflation expectations in the United States have responded to economic developments and presents a model that attempts to explain the patterns in the data. Orphanides and Williams (2005) present a model in which perpetual learning leads to fluctuations in inflation expectations at various horizons and provide examples of how alternative monetary policy settings can influence the course of inflation expectations, thereby illustrating the importance of the interaction between expectations formation and efficient monetary policy.

12. With regard to oil prices, Bodenstein, Erceg, and Guerrieri (2008) discuss the effects on the price of oil and the U.S. economy from shocks to various factors in a theoretical model; Kilian (forthcoming) provides an empirical analysis. More generally, this area remains a fertile field for future research.

13. See Bernanke in part 9 of this volume.

14. One might question whether significant price trends can reasonably be anticipated for a storable commodity like oil. The price of such a commodity should reflect expected demand, and “anticipated” increases should be limited to approximately the cost of storage, including the nominal interest rate. Still, the following discussion would also apply to a shift in the trend in any important subset of consumer prices.

15. An example of a change in the trend of other consumer prices could be seen in the relative prices of computers and other high-technology goods and services in the second half of the 1990s, which declined at an unexpectedly rapid rate

as productivity accelerated. This shock placed downward pressure on inflation and raised both employment and the equilibrium real wage. In the presence of nominal price and wage rigidities, an efficient policy response would facilitate the rise in the real wage by allowing some downward drift in price inflation and upward drift in employment and wage inflation, which is, in fact, about the result observed over this period.

References

- Ball, Laurence M. 1994. "What Determines the Sacrifice Ratio?" In *Monetary Policy*, ed. N. Gregory Mankiw, 155–182. Chicago: University of Chicago Press.
- Ball, Laurence M., and N. Gregory Mankiw. 1995. "Relative-Price Changes as Aggregate Supply Shocks." *Quarterly Journal of Economics* 110(1): 161–193.
- Blanchard, Olivier J., and Jordi Galí. 2007. "The Macroeconomic Effects of Oil Shocks: Why Are the 2000s So Different from the 1970s?" Working Paper 13368. Cambridge, MA: National Bureau of Economic Research.
- Bodenstein, Martin, Christopher Erceg, and Luca Guerrieri. 2007. "Oil Shocks and U.S. External Adjustment." International Finance Discussion Paper 897. Washington, DC: Board of Governors of the Federal Reserve System. Available at http://www.eia.doe.gov/conf_pdfs/Tuesday/Guerrieri.pdf.
- Bodenstein, Martin, Christopher Erceg, and Luca Guerrieri. 2008. "Optimal Monetary Policy in a Model with Distinct Core and Headline Inflation Rates." International Finance Discussion Paper 941: Washington, DC: Board of Governors of the Federal Reserve System. Available at <http://www.federalreserve.gov/pubs/ifdp/2008/941/default.htm>.
- Erceg, Christopher, Dale W. Henderson, and Andrew T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46(2): 281–313.
- Gordon, Robert J. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 2: 297–333.
- Hooker, Mark A. 2002. "Are Oil Shocks Inflationary? Asymmetric and Non-linear Specifications versus Changes in Regime." *Journal of Money, Credit and Banking* 39(2): 540–561.
- Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit and Banking* 39(s1): 101–125.
- Kiley, Michael T. 2008. "Monetary Policy Actions and Long-Run Inflation Expectations." Finance and Economic Discussion Series 2008–03. Washington, DC: Board of Governors of the Federal Reserve System. Available at <http://www.federalreserve.gov/pubs/feds/2008/200803/200803pap.pdf>.

Kilian, Lutz. Forthcoming. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*. Available at <http://www-personal.umich.edu/~kilian/aer-061308final.pdf>.

Lucas, Robert E., Jr. 1972. "Econometric Testing of the Natural Rate Hypothesis." In *The Econometrics of Price Determination*, ed. Otto Eckstein, 50–59. Washington, DC: Board of Governors of the Federal Reserve System.

Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1: 19–46.

Mankiw, N. Gregory, and Ricardo Reis. 2003. "What Measure of Inflation Should a Central Bank Target?" *Journal of the European Economic Association* 1(5): 1058–1086.

Orphanides, Athanasios, and John C. Williams. 2005. "Inflation Scares and Forecast-Based Monetary Policy." *Review of Economic Dynamics* 8(2): 498–527.

Rudd, Jeremy, and Karl Whelan. 2007. "Modeling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit and Banking* 39(s1): 155–170.

Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44(2): 293–335.

Taylor, John B. 1979. "Estimation and Control of a Macroeconomic Model with Rational Expectations." *Econometrica* 47(5): 1267–1286.

Wei, Chao. 2003. "Energy, the Stock Market, and the Putty-Clay Investment Model." *American Economic Review* 93(1): 311–323.

Woodford, Michael. 1994. "Nonstandard Indicators for Monetary Policy: Can Their Usefulness Be Judged From Forecasting Regressions?" In *Monetary Policy*, ed. N. Gregory Mankiw, 95–116. Chicago: University of Chicago Press.

The Phillips Curve and the European Central Bank

Jürgen Stark

I, too, would like to thank the Federal Reserve Bank for having organized this very timely conference on a very important topic against the background of the current global economic and financial situation.

I would like to structure my remarks in the following way: first, I will discuss briefly, in very general terms, the Phillips curve and what I have learned from it. Second, I will present the analytical framework of the European Central Bank (ECB). I think most of you are familiar with our monetary policy framework, but I would like to remind you of what we have learned so far and then briefly discuss the current monetary policy challenges, including a brief comparison between the 1970s and the current situation.

1. My Experiences with the Phillips Curve

Like others, the first time I heard about the Phillips curve was when I was a student of economics. This was more than 40 years ago and what I learned at that time was, to put it in a nutshell, in a very simplistic way, that a little more inflation was supposed to bring a little more real income as a permanent effect. I remember very well the impact of the discussion about the Phillips curve on policymaking in continental Europe in the 1970s, and the trade-off between unemployment and inflation. A very high-ranking European politician at that time argued that if he had to make a political choice between unemployment and inflation, he would prefer 5 percent inflation rather than 5 percent unemployment. It is really evident here that the Phillips curve was taken literally, offering a menu of

combinations of unemployment and inflation from which policymakers could choose at will—which is an exploitable trade-off.

Since the late 1970s, the Phillips curve has been much in the minds and in the charts (and hearts) of macroeconomists. Forty years ago, the Phillips curve was positively sloped. Then it became vertical as advocates of rational expectations drew attention to a principle that had been discovered long ago, but had somehow been forgotten: namely, monetary neutrality.

It is futile for a central bank to maneuver inflation in the hopes of systemically stimulating growth, as this policy is doomed to failure in the end and produces inflation as the only certain outcome. But these theoretical developments by themselves would most likely not have succeeded in displacing what was the consensus of economists at the time if it had not been for a dramatic concomitant real world development, namely the great deflation—followed by the great moderation—with the breakdown of empirical Phillips curves.

The revival of interest in the Phillips curve started in the 1990s. An amended Phillips curve that shifts with the state of expectations has been made the modeling centerpiece of innovative proposals to conduct monetary policy without explicit reference to monetary aggregates. In the New Keynesian Phillips curve there is no role for money. And there is no role for monetary analysis. The question is whether a central bank can content itself with such a stylized representation of the economy, leaving aside important features and important factors of the economy.

There are, in my view, at least two issues which have to be addressed: first, can this model be trusted to get the facts right in the long run and, second, can monetary shocks be ignored even at the business cycle frequencies?

I will not enter into details here. Let me only mention, among other things, that the New Keynesian Phillips curve model does not reproduce the lead-lag structure that links the two variables, monetary growth and inflation. It does not feature a money market or a financial sector. It simply assumes away the shocks that even in the short run can originate in those sectors. I will come back to that in a minute. But, unlike its predecessor, the consensus model centered on a reconstructed Phillips curve grants no free lunch to policymakers. So there is more than a

sense in which we all can say that the New Keynesian Phillips curve encapsulates fundamental tenets of prudent monetary policymaking. The inflation process is forward-looking and will quickly incorporate and perpetuate any deviations from price stability if and when central banks were to start experimenting with the economy. You know that inflation formation is a complex phenomenon in which expectations interact with current and past shocks to the cost structure of firms, as well as to the monetary and financial fabric of the economy in ways that the Phillips curve cannot fully account for or describe. The failure to adopt an all-encompassing view of the inflation process will lay down the potential for economic damage. The first conclusion I would like to draw from what I have said so far is that we should beware of models that short-circuit the workings of a highly complex market economy in a single equation.

2. The European Central Bank's Analytical Framework

Let me briefly remind you of the main features of the ECB's analytical framework. First, our primary objective is price stability, which is enshrined in the Treaty on the establishment of the European Community. Second, we have a definition of price stability that, expressed in terms of the harmonized index of consumer prices, is an inflation rate of below but close to 2 percent. Third, the analytical framework is based on two pillars that structure our analysis and our deliberations in the ECB's Governing Council. The broad economic and monetary analyses help to extract all information relevant for the assessment of the risks to price stability in the medium term. Under the economic pillar, the staff projections on growth and inflation are an input to our monetary policy deliberations. These projections are a key element, but only one among others, in the assessment of economic prospects and of the short- to medium-term risks to price stability. Let me add here that the Governing Council does not underwrite the results of the projections. Neither the Governing Council nor the Executive Board of the ECB are the "owners" of these staff projections. These forecasts are and remain the responsibility of staff.

In these projections, a broad range of models is used, including those based on the Phillips curve. The strategic assignment to expand our view

beyond the real sector has been instrumental in developing different avenues of monetary analysis, which to mention but a few of the avenues we follow include:

- monetary dynamics and the analysis of monetary aggregates
- the components of broad money
- the counterparts of M3
- sectoral developments.

At the ECB, the money and credit data have proved crucial signposts. These data reveal the underlying long-term inflationary trends in the economy and also signal the potential emergence of financial imbalances. Our strategy, in particular, the monetary analysis, has encouraged the macroeconomic modeling of the monetary sector. As a result—and only as one of the examples—we can currently use large-scale estimated dynamic stochastic general equilibrium (DSGE) models with a developed credit market that permit informative simulation experiments around the baseline projections.

We exploit the active channels of transmission from credit and money to activity and inflation, which are built into those models to quantify more precisely the risks to the projections that one could associate with the financial turmoil.

And let me say here that doing this would have been impossible if we had contented ourselves with economic models in which inflation and output move only because of innovations to consumption, to investment, or to cost-push forces. Today, we can count on a first-rate system of real-time monitoring of monetary facts, which we can resort to instantly, especially in times of emergency. Since the start of the financial tensions in early August 2007, monetary analysis has proved a crucial bulwark for the ECB's conduct of monetary policy. Apart from that, from the very outset of the financial market tensions, we have separated the ECB's monetary policy stance from its money market operations.

A close evaluation of the credit data has allowed us to conclude that the availability of bank loans to euro-area companies has not thus far been significantly impaired by the financial turmoil. While we continue to closely monitor all developments in financial markets and in the real

economy, such evidence has proved an important counter to the gloomy prognosis of a credit crunch that has driven much of the recent debate on the outlook for the euro-area economy.

And even more importantly, our mandate and our monetary policy strategy have maintained the necessarily medium-term orientation of monetary policy at a time when short-term forces threatened to overwhelm it. And it has helped to focus our attention on the inflation outlook at medium- to longer-term horizons over which central banks can exert control and for which they ultimately need to take responsibility. The ECB's mandate and our analytical framework gave us guidance in very difficult times when the ECB was assumed by many observers to be in a dilemma with declining economic activity and rising inflation. How important it is to have a robust, credible, and well-understood monetary policy framework is demonstrated by the lessons to be drawn from the Great Inflation of the 1970s. Those central banks and those economies with a very credible medium-term-oriented monetary policy framework succeeded better and produced better results than other central banks.

Let me conclude with the challenges we are facing at present. We are experiencing multiple adverse shocks, as has already been indicated by the previous speakers, with sharp commodity price increases and with financial markets in ongoing turmoil. There are signs of a global reemergence of inflation. However, the euro area has so far been quite resilient to these shocks. As far as the real economy is concerned, what has changed over the last ten years may be that the euro-area economy today is more flexible due to the restructuring of the corporate sector that was accompanied by economic reforms in many countries of the euro area. The economic fundamentals in the euro area are sound and there are no major imbalances. The impact of the financial market turmoil has been limited so far. We still see strong loan growth. The major problem we are facing at present is the high inflation rate. According to our staff projections, which were published ten days ago (early June 2008), we are confronted with an inflation rate that is clearly above what we define as price stability. The range that has been published recently is between 3.2 percent and 3.6 percent for 2008 and between 1.8 percent and 3 percent for 2009. This is an issue that we take very seriously.

In the euro area, inflation expectations are still anchored. However, the longer high inflation rates persist, the more likely inflation expectations will become unanchored. This is something that we have to prevent. On first-round effects, there is little or nothing central banks can do. Now we see first signs of so-called second-round effects, the pass-through from higher commodity prices to wage and price setting. However, we have communicated for quite some time that we are willing to prevent these kinds of second-round effects and it is against the background of what I have just described that the signal given by the Governing Council last week has to be understood. Thank you very much.

Monetary Policy at the Riksbank and the Phillips Curve

Lars E.O. Svensson

I'm very grateful to Eric and Jeff and the organizers of this conference for the opportunity to speak at this great occasion. I will say a few words about monetary policy at the Riksbank and the role of the Phillips curve there. At the Riksbank we conduct flexible inflation targeting, which means that we try to stabilize inflation around the inflation target, which is 2 percent for the consumer price index (CPI) in our case. We also attach some weight to stabilizing the real economy—that is, stabilizing resource utilization measured, for instance, by the output gap. This approach is consistent with minimizing a conventional quadratic loss function that equals the inflation gap between inflation and the inflation target squared plus the weight λ times the output gap squared. We do what can be called “forecast targeting”: we choose a repo-rate path (the repo rate is the Riksbank’s instrument rate) such that the forecast for inflation and the real economy looks good. “Looks good” means that inflation goes to the inflation target and resource utilization goes to a normal level at an appropriate pace, say within two to three years or so. We publish and explain a repo-rate path and our forecast for inflation and the real economy. We try to take the idea of managing expectations seriously. That is, we accept that the current repo rate matters very little or not at all. It is really expectations about the future repo rate and the expectations about inflation and the real economy that matter for the decisions made by the private sector.

Now to the Riksbank’s decisionmaking process. We have a six-member Executive Board. Each member is supposed to have the same information about the policy situation and an equal influence on policy

decisions. We make six policy decisions per year, so on average we have one every other month. At three of these policy meetings we publish a longer *Monetary Policy Report*. At the intervening three meetings we publish a shorter *Monetary Policy Update*. Both the *Report* and the *Update* contain a forecast of inflation, the real economy, and the repo rate. During each decision cycle, there is a series of meetings and a lot of interaction between the staff and the Executive Board. These meetings and interactions result in a main forecast of the repo rate, inflation, and the real economy and possible alternatives to these forecasts. At the final policy meeting in the decision cycle, we discuss and vote on the decision and the *Report* or *Update*. The *Report* or *Update* is published the day after the policy meeting when a press conference is also held. Two weeks later the meeting minutes are published. The minutes are attributed, meaning that each comment or statement is preceded by the name of the speaker. The minutes also include the result of the voting, any dissenting views, and the explanation for such disagreement.

The forecasts and the policy simulations are generated using a set of models. The main model is a state-of-the-art dynamic stochastic general equilibrium (DSGE) model called Ramses. It has been in operational use since 2005, so we have several years of experience in using this model in the decisionmaking process. We also have a Bayesian vector autoregression (VAR) model, and we have a few other models mostly for short-term forecasting, including indicator models and a few single-equation models. The results from these models are combined through a kind of informal model averaging. Quite a bit of judgment is also applied. The end result is our main forecast and a few alternatives to the main forecast. Our forecasts are mean forecasts, not mode forecasts. In practice, we rely on the mean forecasts for policy, so we implicitly assume that certainty equivalence is an acceptable approximation, so the mean forecasts provide enough information for our decisions. We also publish uncertainty intervals, but these serve mostly to remind people about the uncertainty of the forecast and that the forecast, especially the repo-rate path, is simply a forecast and not a firm prediction. Figure 8.4 shows a standard picture in our *Report* or *Update*. The mean and the uncertainty intervals are shown.

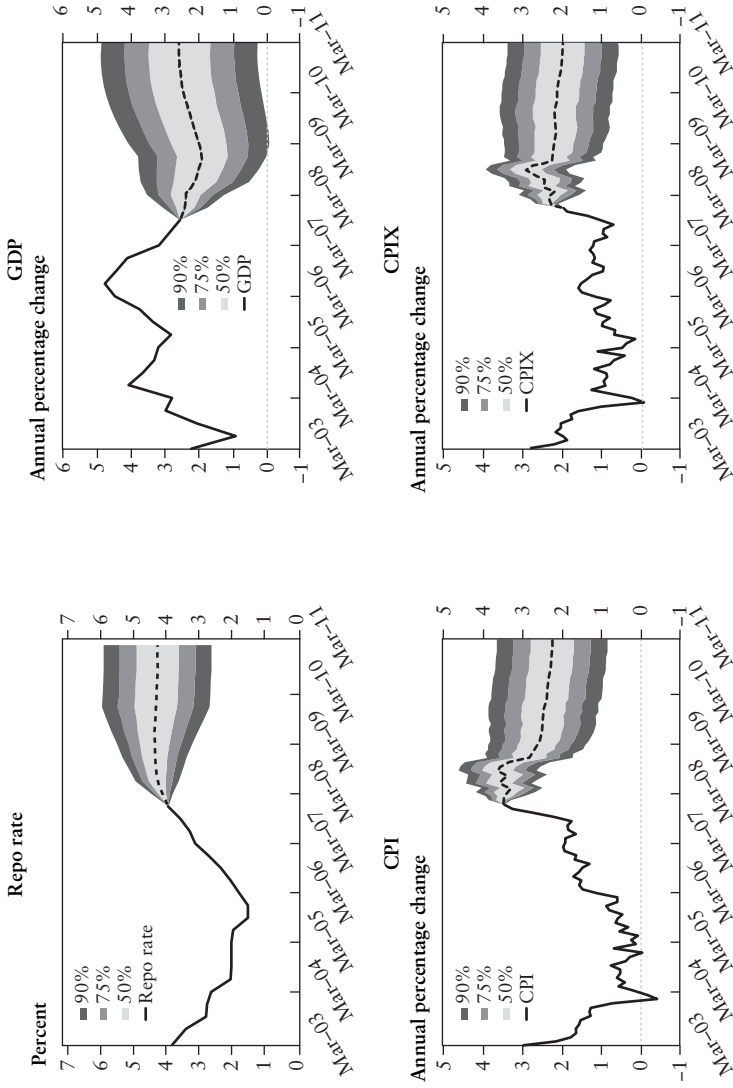


Figure 8.4
Mean Forecasts, Uncertainty Intervals
Source: Sveriges Riksbank.

Let me move on to discuss Ramses, our main model, and the role the Phillips curve plays in it. Ramses is a state-of-the-art open-economy DSGE model and is described in Adolfson, Laséen, Lindé, and Villani (2007). It is estimated with Bayesian methods. The model's structure is similar to the many other central bank DSGE models. There is an aggregate-supply bloc that contains state-of-the-art New Keynesian Phillips curves. There are different Phillips curves for domestic goods, imported consumer goods, imported investment goods, and exports. The aggregate-supply bloc provides the trade-off between the real economy and inflation in the model. There is an aggregate-demand bloc with state-of-the-art Euler conditions for consumption and investment. This bloc specifies how monetary policy affects the real economy. So far most simulations have been carried out with an estimated empirical reaction function, but we are working on implementing optimal policy in the framework, which means having a specific intertemporal loss function and solving the model and producing optimal projections that minimize the loss function (Adolfson, Laséen, Lindé, and Svensson 2008).

What are the implications for the policy discussion given the decision a year and a half ago to publish a repo-rate path? The Riksbank started to publish a repo-rate path in February 2007. This is something that, as an academic, I argued should be done for a long time. My colleagues on the Board actually decided to do this before my appointment to the Board in May 2007. As a consequence of publishing the repo-rate path, the discussion among the Board members is much more about the future repo-rate path than about the current repo rate—the decision about the current repo rate is just a consequence of the path that you have agreed on previously. I think publishing a repo-rate path is a healthy and good policy development. It means that we get a more medium- and long-term perspective on policy. Because the models, in particular Ramses, serve to some extent as a communication framework, we get much more of a general-equilibrium perspective in the policy discussion. We also get more systematic treatment of alternative assumptions about the development of exogenous variables, alternative assumptions about the transmission mechanism of monetary policy, and so on.

Figure 8.5 illustrates the implicit model averaging that occurs. Here the dashed line curve is from Ramses, the light gray curve is from our

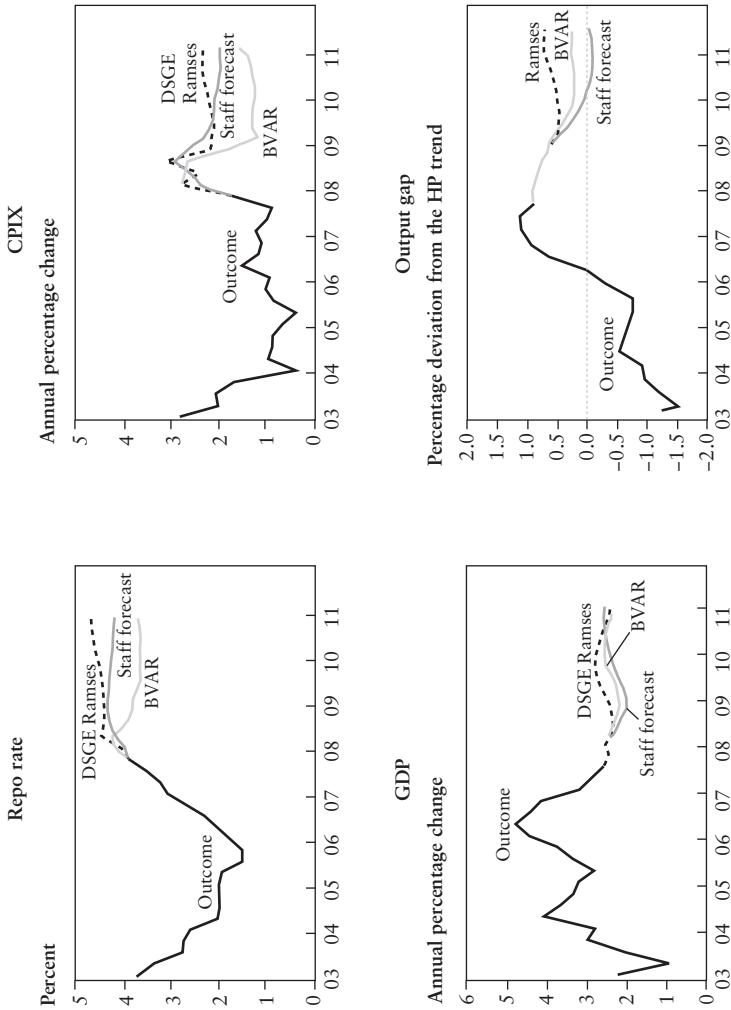


Figure 8.5
 Model and Staff Forecasts
 Source: Sveriges Riksbank.

Bayesian VAR, and the dark gray curve is the staff forecast, the result of implicit model averaging and quite a bit of judgment by the staff.

Figure 8.6 shows the result of different assumptions about exogenous variables. The dark gray curve is the main scenario, the black dotted curve is a simulation with higher international inflation, and the light gray dotted curve is a simulation with greater financial market turmoil.

The Riksbank is the third central bank to publish its own instrument-rate path. Previously, the Reserve Bank of New Zealand (RBNZ) since 1997 and Norges Bank since 2005 have published their own instrument-rate path. At the RBNZ, there is a single decisionmaker, the governor. At Norges Bank, the forecast and repo-rate path presented to the Board (which has five external members and two members from the Bank, the governor and the deputy governor) is actually the forecast of the Bank and the governor. The Board may or may not accept the Bank's forecast and instrument-rate path. Therefore, you can say that Norges Bank also has a single decisionmaker behind the instrument-rate path. This means that the Riksbank is the first central bank to publish an instrument-rate path with a genuine individualistic committee (in Alan Blinder's 2008 terminology) and not a single decisionmaker. This is of some interest, since some people have argued that it is more or less impossible for a genuine committee with several board members to agree on an instrument-rate path (Goodhart 2005). The Riksbank has now demonstrated that it is possible.

Some of you may remember that in previous academic work (Svensson 2007), I have presented an idea of how to aggregate preferences over instrument-rate paths. Figure 8.7 illustrates this.

Suppose that you have three board members. Each one has his or her own preferred instrument-rate path. How do you aggregate these to one path? My suggestion was to just take the median path. In the top panel of figure 8.7, you see the three members' preferred instrument-rate paths as three dotted curves. For each horizon, you then take the median, the solid black curve in the bottom panel of figure 8.7. Then you would start arguing and negotiating about that median. Of course, there is a problem here because that median comes from different paths, and it may not be completely consistent and is not exactly optimal.

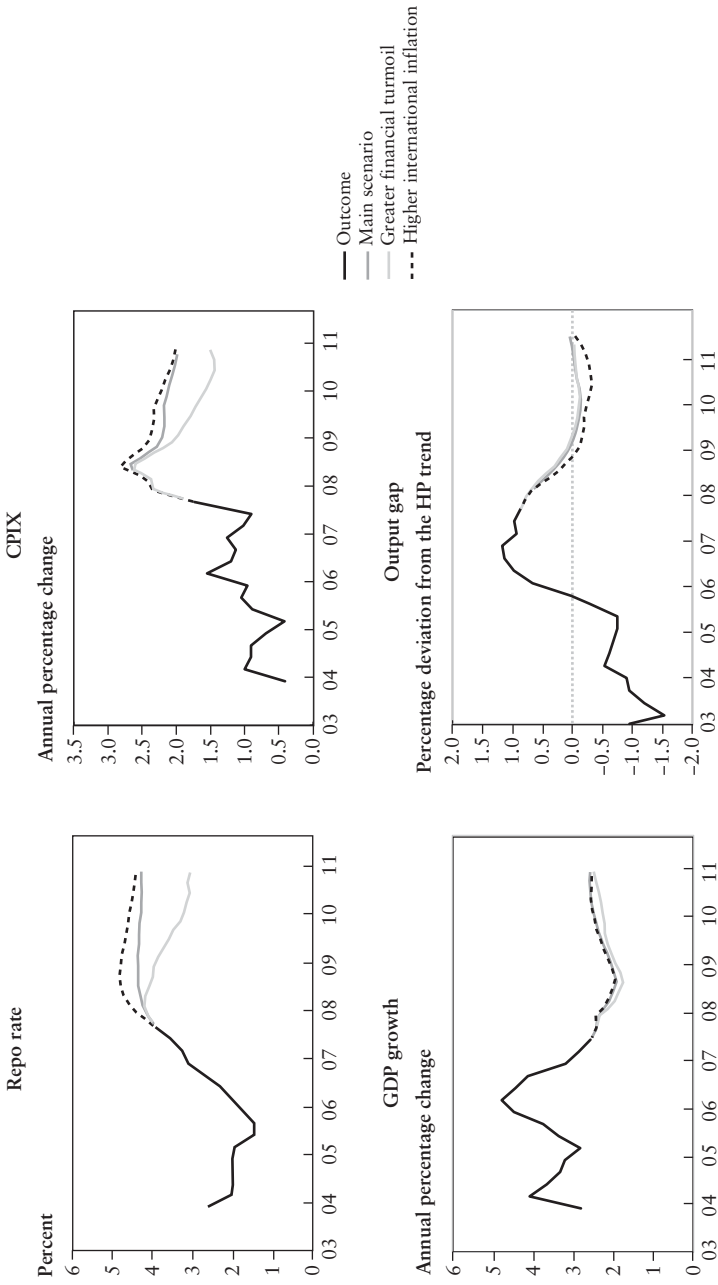


Figure 8.6
 Alternative Scenarios
 Source: Sveriges Riksbank.





Figure 8.7
Deciding on a Repo-rate Path: Median Path?
Source: Author's calculations.

In any case, this is not the way it has worked in practice at the Riksbank so far—it actually has been much easier. During the many meetings and interactions with the staff before the final meeting, we arrive at a main scenario of a repo-rate path and a forecast of inflation and the real economy that the staff deems the Board's majority is likely to prefer. In the process, we may also consider a number of alternatives. Figure 8.8 shows possible alternative repo-rate paths and corresponding forecasts for inflation, GDP growth, and the output gap. The dark gray curves show the main scenario, the black dotted path shows a higher repo-rate path, and the light gray curve shows a lower repo-rate path. The main scenario was chosen by the majority of the Board.

So the practice of choosing a repo-rate path has so far been much simpler than as an academic I thought it would be. In a genuinely individualistic committee, we can easily decide on a repo-rate path with six members. I think that one can do the same thing with a larger committee, say 9, 12, maybe 19. Who knows? I do not think that the number of committee members is crucial. However, the decisionmaking process may be easier if all of the members are full-time in-house members, as at the Riksbank. It remains to be seen. After the Riksbank, Sedlabanki Islands (the central bank of Iceland) and the Czech National Bank have started to publish instrument-rate paths. I look forward to seeing which central bank will be the next to do so.

Figure 8.9 shows our decisions so far, from February 2007 through April 2008. In February 2007, before I joined the Board, the interest rate path was pretty low and the Riksbank had a fairly low inflation forecast. In June 2007, my first policy meeting, the interest-rate path was raised to a higher level, since during the spring inflation pressure had increased quite a bit. Since June 2007 through April 2008, the path has been kept approximately unchanged. During this period, inflation pressure was increasing but the real-economy outlook was increasingly weak, so we thought that an unchanged repo-rate path remained the best compromise between stabilizing inflation and stabilizing the real economy. (The different histories for GDP growth and the output gap are due to data revisions.) At the time of writing, the next policy meeting is in early July 2008, and then we will reconsider the situation, the outlook, and our decision.

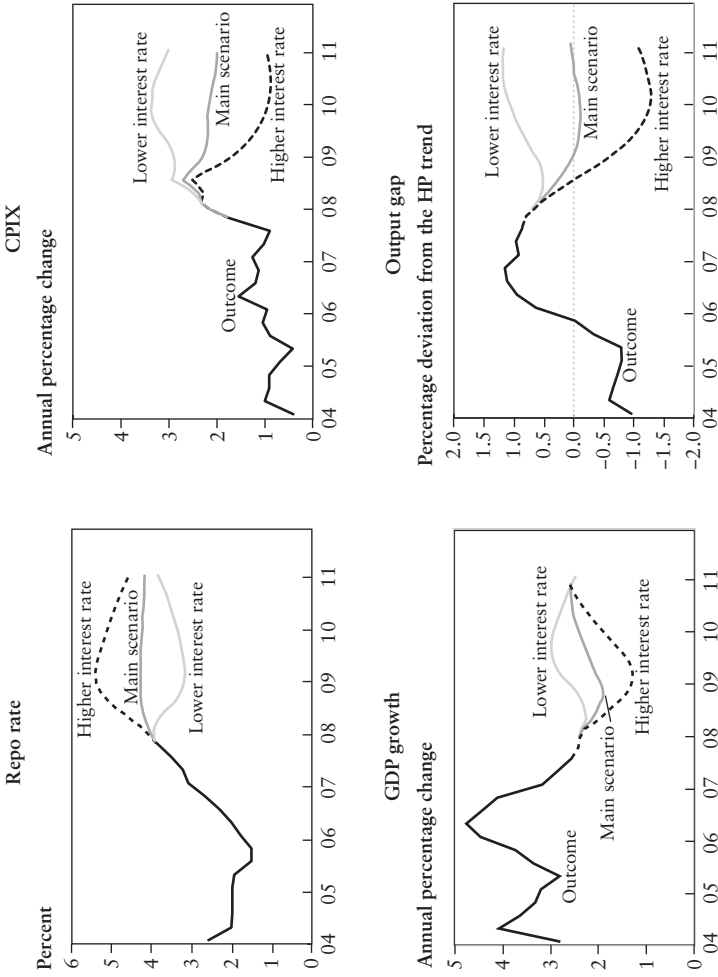


Figure 8.8
Deciding on a Repo-rate Path: Just Vote among a Few Alternatives
Source: Sveriges Riksbank.

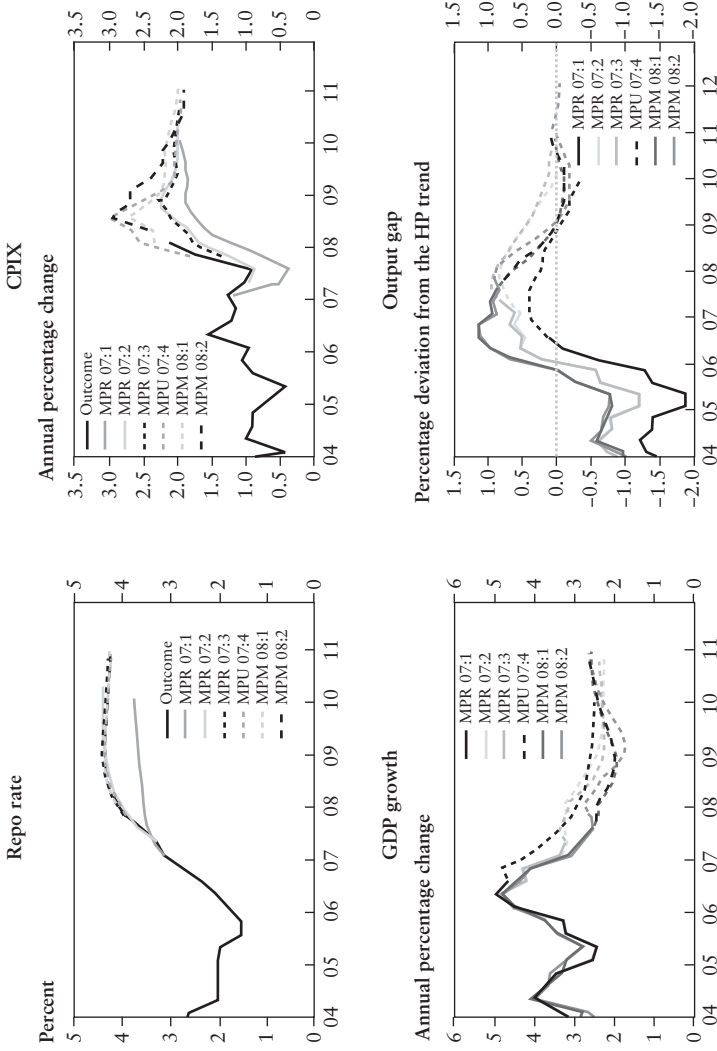


Figure 8.9
 Decisions February 2007–April 2008
 Source: Sveriges Riksbank.

■ *The views and conclusions are solely my responsibility and do not necessarily agree with those of other members of the Riksbank's Executive Board or staff.*

References

- Adolfson, Malin, Stefan Laséen, Jesper Lindé, and Mattias Villani. 2007. "RAMSES – A New General Equilibrium Model for Monetary Policy Analysis." *Sveriges Riksbank Economic Review* 2: 25–39.
- Adolfson, Malin, Stefan Laséen, Jesper Lindé, and Lars E.O. Svensson. 2008. "Optimal Monetary Policy in an Operational Medium-Sized DSGE Model." Working paper. Available at www.princeton.edu/svensson.
- Blinder, Alan S. 2008. "Making Monetary Policy by Committee." Paper presented at the Bank of Canada Economic Conference *International Experience with the Conduct of Monetary Policy under Inflation Targeting*, Ottawa, July 22–23. Available at <http://www.princeton.edu/~ceps/workingpapers/167blinder.pdf>.
- Goodhart, Charles A.E. 2005. "The Interest Rate Conditioning Assumption." Financial Markets Group Discussion Paper No. 547, London: London School of Economics.
- Svensson, Lars E.O. 2007. "Optimal Inflation Targeting: Further Developments of Inflation Targeting." In *Monetary Policy under Inflation Targeting*, ed. Frederic Mishkin and Klaus Schmidt-Hebbel, 187–225. Santiago: Banco Central de Chile.

**The Phillips Curve Going Forward:
What We Still Need to Learn**

Outstanding Issues in the Analysis of Inflation

Ben S. Bernanke

I am pleased to participate in the Federal Reserve Bank of Boston's 53rd annual economic conference, dedicated to the topic of inflation and the Phillips curve. Forecasting and controlling inflation are, of course, central to the process of making monetary policy. In this respect, policymakers are fortunate to be able to build on an intellectual foundation provided by extensive research and practical experience. Nonetheless, much remains to be learned about both inflation forecasting and inflation control. In the spirit of this conference, my remarks will highlight some key areas where additional research could help to provide a still-firmer foundation for monetary policymaking.

I will briefly touch on four topics of particular interest for policymakers: 1) commodity prices and inflation, 2) the role of labor costs in the price-setting process, 3) issues arising from the necessity of making policy in real time, and 4) the determinants and effects of changes in inflation expectations. Economists within the Federal Reserve System and other central banks have made and will continue to make important contributions in these areas. However, researchers in academia and elsewhere have long been essential partners in building the intellectual foundations for the conduct of monetary policy. One of my objectives in these remarks is to encourage the continuation of this fruitful collaboration.

1. Commodity Prices and Inflation

Rapidly rising prices for globally traded commodities have been the major source of the relatively high rates of inflation we have experienced in recent years, underscoring the importance for policy of both forecasting

commodity price changes and understanding the factors that drive those changes.

Policymakers and other analysts have often relied on price quotes from commodities futures markets to derive forecasts of the prices of key commodities. However, as you know, futures markets quotes have underpredicted commodity price increases in recent years, leading to corresponding underpredictions of overall inflation. The poor recent record of commodity futures markets in forecasting the course of prices raises the question of whether policymakers should continue to use this source of information and, if so, how.

Despite this recent record, I do not think it is reasonable, when forecasting commodity prices, to ignore the substantial amounts of information about supply and demand conditions that are aggregated by futures markets. Indeed, the use of some simple alternatives—such as extrapolating recent commodity price trends—would require us to assume that investors in commodity futures can expect to earn supernormal risk-adjusted returns, inconsistent with principles of financial arbitrage. However, it does seem reasonable—and consistent with the wide distributions of commodity price expectations implied by options prices—to treat the forecasts of commodity prices obtained from futures markets, and consequently the forecasts of aggregate price inflation, as highly uncertain.

These considerations raise several questions for researchers. First, is it possible to improve our forecasts of commodity prices, using information from futures markets but possibly other information as well? For example, the markets for longer-dated futures contracts are often quite illiquid, suggesting that the associated futures prices may not effectively aggregate all available information. Second, what are the implications for the conduct of monetary policy of the high degree of uncertainty that attends forecasts of commodity prices? Although theoretical analyses often focus on the case in which policymakers care only about expected economic outcomes and not the uncertainty surrounding those outcomes, in practice policymakers are concerned about the risks to their projections as well as the projections themselves. How should those concerns affect the setting of policy in this context?

Whatever the forecasting value of futures market quotes, these and other financial market prices provide limited information about the

structural relationships between commodity prices and their determinants. Absent a specification of those structural relationships, one cannot analyze the effects of alternative monetary policies or the implications of other shocks to the economy.

Empirical work on inflation, including much of the classic work on Phillips curves, has generally treated changes in commodity prices as an exogenous influence on the inflation process, driven by market-specific factors such as weather conditions or geopolitical developments. By contrast, some analysts emphasize the endogeneity of commodity prices to broad macroeconomic and monetary developments such as expected growth, expected inflation, interest rates, and currency movements. Of course, in reality, commodity prices are influenced by both market-specific and aggregate factors. Market-specific influences are evident in the significant differences in price behavior across individual commodities, which often can be traced to idiosyncratic supply and demand factors. Aggregate influences are suggested by the fact that the prices of several major classes of commodities, including energy, metals, and grains, have all shown broad-based gains in recent years. In particular, it seems clear that commodity prices have been importantly influenced by secular global trends affecting the conditions of demand and supply for raw materials. We have seen rapid growth in the worldwide demand for raw materials, which in turn is largely the result of sustained global growth—particularly resources-intensive growth in emerging market economies.¹ And factors including inadequate investment, long lags in the development of new capacity, and underlying resource constraints have caused the supplies of a number of important commodity classes, including energy and metals, to lag global demand. These problems have been exacerbated to some extent by a systemic underprediction of demand and overprediction of productive capacity for a number of key commodities, notably oil. Further analysis of the range of aggregate and idiosyncratic determinants of commodity prices would be fruitful.

I have only mentioned a few of the issues raised by commodity price behavior for inflation and monetary policy. Here are a few other questions that researchers could usefully address: first, how should monetary policy deal with increases in commodity prices that are not only large but potentially persistent? Second, does the link between global

growth and commodity prices imply a role for global slack, along with domestic slack, in the Phillips curve? Finally, what information about the broader economy is contained in commodity prices? For example, what signal should we take from recent changes in commodity prices about the strength of global demand or about expectations of future growth and inflation?

2. The Role of Labor Costs in Price Setting

Basic microeconomics tells us that marginal cost should play a central role in firms' pricing decisions. And, notwithstanding the effects of changes in commodity prices on the cost of production, for the economy as a whole, by far the most important cost is the cost of labor.

Over the past decade, formal work in the modeling of inflation has treated marginal cost, particularly the marginal cost of labor, as central to the determination of inflation.² However, the empirical evidence for this linkage is less definitive than we would like.³ This mixed evidence is one reason that much Phillips curve analysis has centered on price-price equations with no explicit role for wages.⁴

Problems in the measurement of labor costs may help explain the absence of a clearer empirical relationship between labor costs and prices. Compensation per hour in the nonfarm business sector, a commonly used measure of labor cost, displays substantial volatility from quarter to quarter and year to year, is often revised significantly, and includes compensation that is largely unrelated to marginal costs—for example, exercises (as opposed to grants) of stock options. These and other problems carry through to the published estimates of labor's share in the nonfarm business sector—the proxy for real marginal cost that is typically used in empirical work. A second commonly used measure of aggregate hourly labor compensation, the employment cost index, has its own set of drawbacks as a measure of marginal cost. Indeed, these two compensation measures not infrequently generate conflicting signals of trends in labor costs and thus differing implications for inflation.

The interpretation of changes in labor productivity also affects the measurement of marginal cost. As economists have recognized for half a century, labor productivity tends to be procyclical, in contrast to the

theoretical prediction that movements along a stable, conventional production function should generate countercyclical productivity behavior. Many explanations for procyclical productivity have been advanced, ranging from labor hoarding in downturns to procyclical technological progress. A better understanding of the observed procyclicality of productivity would help us to interpret cyclical movements in unit labor costs and to better measure marginal cost.

The relationship between marginal cost, properly measured, and prices also depends on the markups that firms can impose. One important open question is the degree to which variation over time in average markups may be obscuring the empirical link between prices and labor costs. Considerable work has also been done on the role of time-varying markups in the inflation process, but a consensus on the role of changing markups on the inflation process remains elusive.⁵ More research in this area, particularly with an empirical orientation, would be welcome.

3. Real-Time Policymaking

The measurement issues I just raised point to another important concern for policymakers, namely, the necessity of making decisions in real time, under conditions of great uncertainty—including uncertainty about the underlying state of the economy—and without the benefit of hindsight.

In the context of Phillips curve analysis, a number of researchers have highlighted the difficulty of assessing the output gap—the difference between actual and potential output—in real time.⁶ An inability to measure the output gap in real time obviously limits the usefulness of the concept in practical policymaking. On the other hand, to argue that output gaps are very difficult to measure in real time is not the same as arguing that economic slack does not influence inflation; indeed, the bulk of the evidence suggests that there is a relationship, albeit one that may be less pronounced than in the past.⁷ These observations suggest two useful directions for research. First, more obviously, there is scope to continue the search for measures or indicators of output gaps that provide useful information in real time. Second, we need to continue to think through the decision procedures that policymakers should use under conditions of substantial uncertainty about the state of the economy and underlying

economic relationships. For example, even if the output gap is poorly measured, by taking appropriate account of measurement uncertainties and combining information about the output gap with information from other sources, we may be able to achieve better policy outcomes than would be possible if we simply ignored noisy output gap measures. Of course, similar considerations apply to other types of real-time economic information.

Inflation itself can pose real-time measurement challenges. We have multiple measures of inflation, each of which reflects different coverage, methods of construction, and seasonality, and each of which is subject to statistical noise arising from sampling, imputation of certain prices, and temporary or special factors affecting certain markets. From these measures and other information, policymakers attempt to infer the “true” underlying rate of inflation. In other words, policymakers must read the incoming data in real time to judge which changes in inflation are likely to be transitory and which may prove more persistent. Getting this distinction right has first-order implications for monetary policy: because monetary policy works with a lag, policy should be calibrated based on forecasts of medium-term inflation, which may differ from the current inflation rate. The need to distinguish changes in the inflation trend from temporary movements around that trend has motivated attention to various measures of “core,” or underlying, inflation, including measures that exclude certain prices (such as those of food and energy), “trimmed mean” measures, and others, but alternative approaches are certainly worth consideration.⁸ Further work on the problem of filtering the incoming data so as to obtain better measures of the underlying inflation trend could be of great value to policymakers.

The necessity of making policy in real time highlights the importance of maintaining and improving the economic data infrastructure and, in particular, working to make economic data timelier and more accurate. I noted earlier the problems in interpreting existing measures of labor compensation. Significant scope exists to improve the quality of price data as well—for example, by using the wealth of information available from checkout scanners or finding better ways to adjust for quality change. I encourage researchers to become more familiar with the strengths and shortcomings of the data that they routinely use. Besides leading to bet-

ter analysis, attention to data quality issues by researchers often leads to better data in the longer term, both because of the insights generated by research and because researchers are important and influential clients of data collection agencies.

4. Inflation Expectations

Finally, I will say a few words on inflation expectations, which most economists see as central to inflation dynamics. But there is much we do not understand about inflation expectations, their determination, and their implications. I will divide my list of questions into three categories.

First, we need to better understand the factors that determine the public's inflation expectations. As I discussed in some detail in a talk at the National Bureau of Economic Research in July 2007, much evidence suggests that expectations have become better anchored than these were a few decades ago, but that these expectations nonetheless remain imperfectly anchored.⁹ It would be quite useful for policymakers to know more about how inflation expectations are influenced by monetary policy actions, monetary policy communication, and other economic developments such as oil price shocks.

The growing literature on learning in macroeconomic models appears to be a useful vehicle to address many of these issues.¹⁰ In a traditional model with rational expectations, a fixed economic structure, and stable policy objectives, there is no role for learning by the public. In such a model, there is generally a unique long-run equilibrium inflation rate that is fully anticipated; in particular, the public makes no inferences based on central bankers' words or deeds. But in fact, the public has only incomplete information about both the economy and policymakers' objectives, goals that may change over time. Allowing for the possibility of learning by the public is more realistic and tends to generate more reasonable conclusions about how inflation expectations change and, in particular, about how these can be influenced by monetary policy actions and communications.

The second category of questions involves the channels through which inflation expectations affect actual inflation. Is the primary linkage from inflation expectations to wage bargains, or are other channels impor-

tant? One somewhat puzzling finding comes from a survey of business pricing decisions conducted by Blinder, Canetti, Lebow, and Rudd, in which only a small share of respondents claimed that expected aggregate inflation affected their pricing decisions at all.¹¹ How do we reconcile this result with our strong presumption that expectations are of central importance for explaining inflation? Perhaps expectations affect actual inflation through some channel that is relatively indirect. The growing literature on disaggregated price setting may be able to shed some light on this question.¹²

Finally, a large set of questions revolve around how the central bank can best monitor the public's inflation expectations. Many measures of expected inflation exist, including expectations taken from surveys of households, forecasts by professional economists, and information extracted from markets for inflation-indexed securities. Unfortunately, only very limited information is available on expectations of price-setters themselves, namely businesses. Which of these agents' expectations are most important for inflation dynamics, and how can central bankers best extract the relevant information from the various available measures?

5. Conclusion

I have touched on only a few of the questions that confront policymakers as we deal with the challenges we face. The contributions of economic researchers in helping us to address these and other important questions have been and will continue to be invaluable.

Notes

1. According to one study, if the share of world trade and world gross domestic product for nonindustrial countries had remained at its 2000 levels, then by 2005, real oil prices would have been 40 percent lower, and real metals prices 10 percent lower, than these actually were (Pain, Koske, and Sollie 2006). Since 2005, continued strong growth in the demand for resources by emerging market economies has likely put further considerable upward pressure on commodity prices. In contrast, the demand for oil by members of the Organisation for Economic Co-operation and Development has been essentially flat since 2004.

2. Galí and Gertler 1999.
3. Rudd and Whelan 2007; Kiley 2007.
4. Gordon 1988.
5. Rotemberg and Woodford 1999.
6. Orphanides 2002.
7. For a counterargument, however, see Atkeson and Ohanian 2001.
8. Rich and Steindal 2007 provide a recent analysis of alternative measures of core inflation.
9. Bernanke 2007.
10. Orphanides and Williams 2007 and Kiley 2008 are good examples, but there are many others.
11. Blinder, Canetti, Lebow, and Rudd 1998.
12. For example, see Bils and Klenow 2004.

References

- Atkeson, Andrew, and Lee E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve Bank of Minneapolis *Quarterly Review* 25 (Winter): 2–11.
- Bernanke, Ben S. 2007. "Inflation Expectations and Inflation Forecasting." Speech delivered at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute. Cambridge, MA, July 10. Available at <http://www.federalreserve.gov/newsevents/speech/Bernanke20070710a.htm>.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.
- Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.
- Gordon, Robert J. 1988. "The Role of Wages in the Inflation Process." *American Economic Review* 78(2): 276–283.
- Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit and Banking* 39(1, Supplement): 101–125.
- Kiley, Michael T. 2008. "Monetary Policy Actions and Long-Run Inflation Expectations." Federal Reserve Board Finance and Economics Discussion Paper Series 2008-03. Available at <http://www.federalreserve.gov/Pubs/feds/2008/200803/200803pap.pdf>.

Orphanides, Athanasios. 2002. "Monetary Policy Rules and the Great Inflation." *American Economic Review* 92(2): 115–120.

Orphanides, Athanasios, and John C. Williams. 2007. "Robust Monetary Policy with Imperfect Knowledge." *Journal of Monetary Economics* 54(5): 1406–1435.

Pain, Nigel, Isabell Koske, and Marte Sollie. 2006. "Globalisation and Inflation in the OECD Economies." OECD Economics Department Working Paper Series 524. Paris: Organisation for Economic Co-operation and Development. Available at [http://www.oilis.oecd.org/olis/2006doc.nsf/LinkTo/NT0000731A/\\$FILE/JT03218198.PDF](http://www.oilis.oecd.org/olis/2006doc.nsf/LinkTo/NT0000731A/$FILE/JT03218198.PDF).

Rich, Robert W., and Charles Steindal. 2007. "A Comparison of Measures of Core Inflation." Federal Reserve Bank of New York *Economic Policy Review* 13. Available at <http://www.newyorkfed.org/research/epr/07v13n3/0712rich.pdf>.

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics* Volume 1, ed. John B. Taylor and Michael Woodford, 1051–1135. Amsterdam: Elsevier.

Rudd, Jeremy, and Karl Whelen. 2007. "Modeling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit and Banking* 39(1, Supplement): 155–170.

Contributors

Laurence M. Ball is a professor of economics at Johns Hopkins University, where he has taught macroeconomics since 1994. His theoretical work focuses on the microeconomic foundations of Keynesian macroeconomic models, while his applied research focuses on inflation and monetary policy, both in the United States and in high-inflation countries such as Brazil. In particular, Ball has made contributions to the literature on inflation targeting, productivity, NAIRU, and the relationship between monetary policy and expectations. He is a research associate at the National Bureau of Economic Research and was on the Academic Advisory Panel of the Federal Reserve Board in 2006. Ball held a Houblon-Norman Fellowship at the Bank of England in 2000–2001, and has been a visiting scholar at the Board of Governors (1993, 1996, and 1999), the Federal Reserve Bank of Kansas City (1995–2000), the Federal Reserve Bank of Philadelphia (1991–1993), and the central banks of Australia, Hong Kong, Japan, and Norway. Before joining Johns Hopkins, he was an assistant professor of economics at Princeton University from 1989 to 1994 and at the Graduate School of Business at New York University from 1985 to 1989. He earned a B.A. in economics from Amherst College, and a Ph.D. in economics from the Massachusetts Institute of Technology.

Ben S. Bernanke was sworn in on February 1, 2006, as chairman and a member of the Board of Governors of the Federal Reserve System. He also serves as chairman of the Federal Open Market Committee, the System's principal monetary policymaking body. Prior to his current appointment, from June 2005 to January 2006 Bernanke was chairman of the

President's Council of Economic Advisers. Bernanke's earlier service in the Federal Reserve System was as a member of the Board of Governors of the Federal Reserve System from 2002 to 2005; as a visiting scholar at the Federal Reserve Banks of Philadelphia (1987–1989), Boston (1989–1990), and New York (1990–1991 and 1994–1996); and as a member of the Academic Advisory Panel at the Federal Reserve Bank of New York (1990–2002). Before joining the Board of Governors, Bernanke had been a professor of economics and public affairs at Princeton University since 1985. From 1994 to 1996, Bernanke was the Class of 1926 Professor of Economics and Public Affairs, and from 1996 to 2002 he was the Howard Harrison and Gabrielle Snyder Beck Professor of Economics and Public Affairs and chair of the Princeton economics department. Previously, Bernanke had been at the Graduate School of Business at Stanford University, first as an assistant professor of economics (1979–1983) and then as an associate professor of economics (1983–1985). He was a visiting professor of economics at New York University in 1993 and at the Massachusetts Institute of Technology in 1989–1990. Bernanke is a Fellow of the Econometric Society and of the American Academy of Arts and Sciences. He served as the director of the Monetary Economics Program of the National Bureau of Economic Research (NBER) and as a member of the NBER's Business Cycle Dating Committee. Bernanke received a B.A. in economics from Harvard University and a Ph.D. in economics from the Massachusetts Institute of Technology.

Olivier Blanchard is currently the chief economist and director of the Research Department at the International Monetary Fund. He is on leave from the Massachusetts Institute of Technology, where he is the Class of 1941 Professor of Economics, and a former chair of the economics department. Blanchard is a macroeconomist who has worked on a wide set of issues spanning the role of monetary policy, the nature of speculative bubbles, the nature of the labor market and the determinants of unemployment, and economic transition in former communist countries. In the process, he has worked with numerous countries and international organizations. Blanchard is a research associate at the National Bureau of Economic Research, a fellow and council member of the Econometric Society, a past vice president of the American Economic Association, a

member of the American Academy of Sciences, and a member of the French Economic Advisory Council to the French Prime Minister. He taught at Harvard University from 1977 until 1982, when he joined the MIT economics department. A citizen of France, Blanchard has an M.S. in economics from the University of Paris-Nanterre, and a Ph.D. in economics from the Massachusetts Institute of Technology.

V.V. Chari is currently the Paul W. Frenzel Land Grant Professor of Liberal Arts at the University of Minnesota, where he has been an economics professor since 1994, serving as chair of the economics department from 1997 to 2000. His research interests focus on banking, fiscal and monetary policy, and economic development. Chari has written extensively on banking crises, exchange rate fluctuations, and international capital flows. Since 1994 Chari has been an advisor at the Federal Reserve Bank of Minneapolis, which he joined in 1986 as an economist. Chari was an assistant professor of managerial economics at Northwestern University's Kellogg Graduate School of Management from 1980 to 1986; in 1992 he returned to the Kellogg School as the Harold H. Hines, Jr. Professor of Risk Management. Chari is a research associate at the National Bureau of Economic Research and a fellow of the Econometric Society. He has a B.Tech. in chemical engineering from the Indian Institute of Technology, and a Ph.D. in economics from Carnegie-Mellon University.

William T. Dickens is Distinguished Professor of Economics and Social Policy at Northeastern University, Boston, and a nonresident senior fellow in the Economic Studies Program at The Brookings Institution. Dickens is currently codirector of a major international research project on wage rigidity, a collaborative effort that involves the Brookings Institution, the Federal Reserve Bank of New York, the European Central Bank, and economists from 13 country teams. His areas of expertise are labor economics, wage discrimination, unemployment, monetary policy, inner-city employment problems, effects of trade on employment and wages, poverty, income support, intelligence testing, and psychology and economics. Dickens was formerly a senior economist for the President's Council of Economic Advisers from July 1993 until June 1994; a professor of economics at the University of California, Berkeley from 1980

until 1995; and a visiting assistant professor at the Massachusetts Institute of Technology's Sloan School of Management. For the 2007–2008 academic year, he was named the Thomas C. Schelling Visiting Professor at the University of Maryland's School of Public Policy. He received his B.A. in social studies from Bard College and his Ph.D. in economics from the Massachusetts Institute of Technology.

Stanley Fischer has been governor of the Bank of Israel since May 2005. Prior to joining the Bank of Israel, he was vice chairman of Citigroup from February 2002 through April 2005, where he was also head of the Public Sector Group from February 2004 to April 2005, chairman of the Country Risk Committee, and president of Citigroup International. Fischer was the first deputy managing director of the International Monetary Fund (IMF) from September 1994 until the end of August 2001. Before he joined the IMF, Fischer was the Killian Professor and head of the Department of Economics at the Massachusetts Institute of Technology, which he joined in 1973 as an associate professor. From January 1988 to August 1990 he was vice president, development economics and chief economist at the World Bank. He was assistant professor of economics at the University of Chicago from 1970 until 1973. Fischer is a research associate at the National Bureau of Economic Research and a fellow of the Econometric Society and of the American Academy of Arts and Sciences. He earned a B.Sc. and M.Sc. in economics at the London School of Economics, and obtained his Ph.D. in economics from the Massachusetts Institute of Technology.

Jeff Fuhrer is executive vice president and director of research at the Federal Reserve Bank of Boston, where he is responsible for the Bank's research and monetary policy functions. He is an associate economist of the Federal Open Market Committee, and regularly attends this key U.S. policy meeting with the Bank's president. In June 1992 Fuhrer joined the Bank's research department as an assistant vice president and economist, and from 1995–2001 headed its open economy macro/international section. In 2000 Fuhrer was named senior vice president and monetary policy advisor, in 2001 he became director of research, and in 2006 he was named executive vice president. Fuhrer began his career at the Board

of Governors of the Federal Reserve System, first as a research assistant, and then in 1985 returned as a senior economist after earning his doctorate. He has been active in economic research for more than two decades, and has served as an associate editor for the *American Economic Review*. Fuhrer has published numerous scholarly papers on the interactions among monetary policy, inflation, consumer spending, and asset prices. He earned an A.B. in economics from Princeton University and an M.A. and Ph.D. in economics from Harvard University.

Jordi Galí directs the Center for Research in International Economics (CREI) at the University of Pompeu Fabra in Barcelona, where he has taught since 2001. Galí specializes in macroeconomic theory, monetary economics, and macroeconometrics. Much of his research centers on the causes of business cycles and optimal monetary policy, especially when using times series analysis. In 2005 Galí shared the Yrjö Jahnsson Award from the European Economic Association, given in recognition of his work on New Keynesian macroeconomics. Prior to his current position, Galí taught at Columbia University and New York University. In 2005–2006 he was a visiting professor in the economics department at the Massachusetts Institute of Technology and a visiting scholar at the Federal Reserve Bank of Boston. He is a fellow of the Econometric Society, a research associate at the National Bureau of Economic Research, and a research fellow at the Centre for Policy Research (CEPR) in London. Galí earned both an undergraduate degree in business and a master's in international management from the ESADE Business School in Barcelona, and a Ph.D. in economics from the Massachusetts Institute of Technology.

Michael T. Kiley became the assistant director of the Division of Research and Statistics at the Board of Governors of the Federal Reserve System in 2007. Prior to his current position, Kiley was the Board's chief of macroeconomics and quantitative studies from 2003 until 2007. His research focus spans the fields of macroeconomics, monetary economics, and econometrics. In particular, Kiley has worked on business cycle persistence, consumption growth, information issues, price stickiness, and wages. He worked as an economist at the Organisation for Economic Cooperation and Development from 2001 until 2003. Kiley first joined

the Board of Governors as an economist in 1996, a position he held until 2003. In 2000–2001 Kiley was a lecturer at Johns Hopkins University, and in 1998 he was a visiting assistant professor at the University of Michigan, Ann Arbor. He earned a B.A. in economics from the University of Delaware, and a Ph.D. in economics from the University of Maryland at College Park.

Robert G. King is a professor of economics at Boston University, which he joined in 2000. His broad interests lie in financial economics, monetary economics, and macroeconomics. King's particular research focus is on price stability and optimal monetary policy as these concerns apply to business cycles, and interactions between the money supply, credit, and prices. He has also contributed to the literature on inflation targeting. King is a research associate at the National Bureau of Economic Research and edits the *Journal of Monetary Economics*. Before his current academic appointment, King was at the University of Virginia, where from 1997 until 2000 he was the Robert P. Black Research Professor of Monetary Economics and the Carter Glass Professor of Banking Economics; from 1993 to 1997 he was the A.W. Robertson Professor of Economics. King taught at the University of Rochester from 1978 until 1993, and served as chair of the economics department from 1987 to 1988 and from 1990 to 1991. He also directed the Rochester Center for Economic Research from 1988 until 1991. King earned a B.A. in economics and mathematics at Brown University, which also granted his Ph.D. in economics.

Yolanda K. Kodrzycki is a senior economist and policy advisor at the Federal Reserve Bank of Boston, where she has worked since 1986. Kodrzycki specializes in regional, labor market, and public sector economics; her wide-ranging research has examined diverse topics, including housing prices, the long-term implications of job loss, the migration patterns of college graduates, the causes of regional differences in educational attainment, the privatization of government functions, and corporate tax policy at the national and state levels. Kodrzycki coedits *Massachusetts Benchmarks*, a quarterly journal on the commonwealth's economy published by the Donahue Institute in cooperation with the Federal Reserve

Bank of Boston. During 1991–1992 Kodrzycki was consulting for the U.S. Treasury advisory program in central and eastern Europe while on leave from the Boston Fed. Prior to joining the Federal Reserve Bank of Boston, she taught economics at Amherst College. Kodrzycki received a B.A. from Radcliffe College (Harvard University), and an M.A. and a Ph.D. in economics from the University of Pennsylvania.

Donald L. Kohn took office on August 5, 2002 as a member of the Board of Governors of the Federal Reserve System for a full term ending January 31, 2016. On June 23, 2006, Kohn was sworn in as vice chairman of the Board of Governors of the Federal Reserve System for a four-year term. Kohn has written extensively on issues related to monetary policy and its implementation by the Federal Reserve. He is chairman of the Committee on the Global Financial System, a central bank panel that monitors and examines broad issues related to financial markets and systems. Prior to becoming a member of the Board of Governors, Kohn served on its staff as adviser for monetary policy from 2001–2002, secretary of the Federal Open Market Committee from 1987–2002, director of the Division of Monetary Affairs from 1987–2001, and deputy staff director for Monetary and Financial Policy from 1983–1987. He also held several positions in the Board’s Division of Research and Statistics: associate director from 1981–1983, chief of capital markets from 1978–1981, and economist from 1975–1978. Kohn began his career as a financial economist at the Federal Reserve Bank of Kansas City, where he worked from 1970 to 1975 before joining the Board of Governors. Kohn received a B.A. in economics from the College of Wooster and a Ph.D. in economics from the University of Michigan.

Jane Sneddon Little is an economist in the research department of the Federal Reserve Bank of Boston, where she specializes in international macroeconomic issues. In the past few years she has written or coauthored papers on asset prices and economic stabilization, the evolution of the international monetary system, the offshoring of jobs, and the practice of inflation targeting overseas. Recent contributions include editing *Global Imbalances and the Evolving World Economy* and *Wanting It All: The Challenge of Reforming the U.S. Health Care System*. In

addition to her duties at the Boston Fed, where she has spent her entire career, starting as a research assistant, Little has worked on the Massachusetts Governor's Council on Economic Growth and Technology, and on the Task Force on the Health Care Industry. She has also been a lecturer for Simmons College. Little holds a B.A. from Wellesley College and M.A.L.D. and M.A. degrees from the Fletcher School of Law and Diplomacy at Tufts University.

Bartosz Maćkowiak is an economist in the monetary policy research Division at the European Central Bank (ECB), and a research affiliate in international macroeconomics at the Centre for Economic Policy Research. His research interests focus on macroeconomics, international economics, and applied time series econometrics. Maćkowiak has published on currency crises, price stickiness, and macroeconomic variables in emerging economies. Before joining the ECB in April 2007, he was an assistant professor at Humboldt University of Berlin from August 2002 to March 2007. He completed his dissertation as a visiting graduate student in the department of economics at Princeton. Maćkowiak has a B.A. in economics from Amherst College, and a Ph.D. in economics from Yale University.

N. Gregory Mankiw is the Robert M. Beren Professor of Economics at Harvard University, and a research associate at the National Bureau of Economic Research. His academic work includes research on business cycles, consumer behavior, economic growth, financial markets, monetary and fiscal policy, and price adjustment. Mankiw served as the Chairman of the Council of Economic Advisors from mid-2003 until early 2005. A prolific writer and a regular participant in academic and policy debates, his published articles have appeared in top academic journals and in more widely accessible forums, such as the *New York Times*, the *Washington Post*, the *Wall Street Journal*, and *Fortune*. He has written two very successful college textbooks—the intermediate-level *Macroeconomics* and the introductory *Principles of Economics*. Mankiw has been an adviser to the Federal Reserve Bank of Boston and the Congressional Budget Office. Currently, he is a visiting scholar at the American Enterprise Institute, where he is focusing on entitlements, fiscal policy, inter-

national economic policy, and trade. He earned an A.B. in economics at Princeton University, and a Ph.D. in economics at the Massachusetts Institute of Technology.

Virgiliu Midrigan joined New York University in 2007 as an assistant professor of economics. A specialist in macroeconomics and monetary economics, Midrigan has worked on price-setting, price stickiness, and how business firms make decisions in relation to monetary policy. Midrigan is currently a research faculty fellow at the National Bureau of Economic Research. From 2006 to 2007 he was a research economist at the Federal Reserve Bank of Minneapolis, and in 2005 Midrigan was a fellow at the Institute for Computational Economics at the University of Chicago/Argonne National Laboratory. He earned a B.A. in economics at American University in Bulgaria, and a Ph.D. in economics from Ohio State University in 2006.

Giovanni P. Olivei is a vice president at the Federal Reserve Bank of Boston, where he oversees the research department's macroeconomics/finance section. Olivei's current projects center on monetary economics, with particular emphasis on inflation dynamics. He joined the Boston Fed in 1998 as an economist, and later became a senior economist and policy advisor before being named vice president in late 2007. Olivei earned his B.A. in economics at Università Commerciale L. Bocconi, Italy, and earned his Ph.D. in economics at Princeton University.

Athanasios Orphanides has been the Governor of the Central Bank of Cyprus since 2007. Prior to this appointment, he was at the Board of Governors of the Federal Reserve System, starting there as an economist (1990–1996) before subsequent promotions as a senior economist (1996–2003), adviser (2003–2006), and then senior adviser (2006–2007). Orphanides's research interests concentrate on monetary policy and macroeconomics, with a special emphasis on inflation targeting, interest rates, and real-time data and decisionmaking. He was a visiting/adjunct professor at Johns Hopkins University from 1995 until 2004, and was a guest lecturer at the Kiel Institute for World Economics in November 2004 and at Goethe University in August 2002. Orphanides

served as an associate editor for the *Journal of Economic Dynamics and Control* from 2002 to 2007. He has an S.B in mathematics and a Ph.D. in economics from the Massachusetts Institute of Technology.

Adrian Rodney Pagan is a professor of economics at Queensland University of Technology, where he specializes in econometrics and macroeconomics. He is also a research associate at the Australian National University's Centre for Applied Macroeconomic Analysis, where he participates in the program on macroeconometric models and methods. Pagan's particular interests include business cycles and achieving coherence between theoretical and empirical macroeconomic models. Throughout his career Pagan has taught at many universities, either as a faculty member or a visiting professor. He was a professor of economics at the University of Rochester between 1986 and 1992, and has visited at Yale University, Monash University, Johns Hopkins University, University of California at Los Angeles, Nuffield College at Oxford University, and at the University of New South Wales. From 1995 to 2000, Pagan was a member of the Reserve Bank of Australia Board. He has been a fellow of the Econometric Society since 1985. He earned a B.Ec. from the University of Queensland, and a Ph.D. from Australian National University.

Christopher Pissarides holds the Norman Sosnow Chair in Economics at the London School of Economics (LSE), where he has spent most of his professional career since 1976. He is a macroeconomist with expertise in the economics of unemployment, labor market theory, and labor market policy research. Pissarides has served as head of the LSE economics department, and is an elected fellow of the British Academy, the Econometric Society, the European Economic Association, and the Society of Labor Economists. He is a research fellow at the LSE's Centre of Economic Performance, and former head of its research program in macroeconomics. Pissarides is also a research fellow at the Centre for Economic Policy Research (London) and at the Institute for the Study of Labor (IZA, Bonn). He is a nonnational senior associate, Forum for Economic Research in the Arab Countries, Iran and Turkey, and served as a member of the Monetary Policy Committee of the Central Bank of

Cyprus from 2000 to 2007. In 2005 he was awarded the IZA Prize in Labor Economics (jointly with Dale Mortensen) for his work on unemployment. In 2008 he received the Republic of Cyprus “Aristeion” for the Arts, Literature, and Science. Pissarides has a B.A. and M.A. in economics from the University of Essex, and a Ph.D. in economics from LSE.

Lucrezia Reichlin is a professor of economics at the London Business School. Prior to her current position, she served as the director general of research at the European Central Bank (ECB) from March 2005 to September 2008, and was a full professor of economics at the Université Libre de Bruxelles. Reichlin specializes in applied macroeconomics and time series; she has worked on developing models for signal extraction and forecasting key economic indicators from large panels of time series. Other research she has conducted centers on structural identification problems in vector autoregression and related models; these techniques have been widely applied in central banks around the world. She has consulted for the Bank of Italy, the Board of Governors of the Federal Reserve System, the ECB, and the Swiss National Bank. From 1999 to 2004, Reichlin codirected the Center for Economic Policy Research’s program in international economics, and is currently a research associate there. She has previously worked at the Observatoire Français des Conjonctures Economiques (OFCE) in Paris, the Graduate School of Business at Columbia University, and at the European University Institute. Reichlin has an undergraduate degree in economics from the University of Modena (Italy) and a Ph.D. in economics from New York University.

Paul A. Samuelson, professor emeritus of economics at the Massachusetts Institute of Technology, is the recipient of the 1970 Nobel Prize in economics. His contributions to economic theory are extensive, including, but not limited to, work in modern welfare economics, linear programming, Keynesian economics, economic dynamics, international trade theory, logic choice, and maximization. Samuelson has authored or coauthored hundreds of books and articles covering these and other topics, and he is the author of the all-time best-selling principals of economics textbook. He was awarded the David A. Wells Prize by

Harvard University in 1941, and the American Economic Association awarded Samuelson the John Bates Clark Medal in 1947 as the economist under 40 “who has made the most distinguished contribution to the main body of economic thought and knowledge.” He served as economic advisor to President John F. Kennedy and has been president of the International Economic Association, American Economic Association, and the Econometric Society. Samuelson received a B.A. from the University of Chicago and M.A. and Ph.D. degrees from Harvard University.

Christopher A. Sims is the Harold H. Helm ‘20 Professor of Economics and Banking at Princeton University, a post he has held since 2004. An econometrician and a macroeconomist, his current areas of research include econometric theory for dynamic models and macroeconomic theory and policy. Before joining Princeton University in 1999, Sims taught at Yale University from 1990 until 1999, at the University of Minnesota from 1970 until 1990, and at Harvard University from 1968 until 1970. He is a research associate at the National Bureau of Economic Research, a fellow and past president of the Econometric Society, a fellow of the American Academy of Arts and Sciences, and a member of the National Academy of Sciences. Sims earned an A.B in mathematics at Harvard College, and a Ph.D. in economics at Harvard University.

Frank R. Smets has been the deputy director general of the European Central Bank (ECB) since June 2005. Prior to assuming this post, he was a principal (December 1998 to November 2002) and then head of the ECB’s monetary policy research division (December 2002 to May 2005). Smets’s latest research focuses on international monetary policy coordination, while his previous contributions have been to monetary policy in relation to business cycles, technological shocks, learning, inflation persistence, and output gaps. Before joining the ECB, Smets was a senior economist in the monetary and economic department at the Bank for International Settlements from 1992 until 1998. In 2004 he was awarded the Hicks-Tinbergen Medal, a biannual prize for the best paper published in the *Journal of the European Economic Association*. Smets has an undergraduate degree in economics from Ghent University and a Ph.D. in economics from Yale University.

Robert M. Solow is the Institute Professor of Economics at the Massachusetts Institute of Technology (MIT), where he has spent his entire academic career since 1949. In 1961 Solow was awarded the John Bates Clark Medal. He served on the Council of Economic Advisers from 1961 to 1962, and was a consultant to the Council from 1962 until 1968. His academic research has been in employment, capital theory, and macroeconomic growth. Solow is best known for his work on the theory of economic growth, and his well-recognized contributions to this field culminated in his being awarded the Nobel Memorial Prize in Economics in 1987. Among other posts and honors, he is a fellow and past president of the Econometric Society, a past president of the American Economic Association, a member of the National Academy of Sciences, a fellow of the American Academy of Arts and Sciences, and was a member of the Board of the Federal Reserve Bank of Boston from 1975 to 1981. Currently he is the Foundation Fellow of the Russell Sage Foundation, where he has worked on explaining the productivity gains in the U.S. economy during the roaring 1990s, and is currently working on a comparative study on the nature and institutional background of low-wage work in the United States and several European countries. Solow earned an A.B., an M.A., and Ph.D. in economics from Harvard University.

Jürgen Stark has been a member of the executive board of the European Central Bank (ECB) since June 2006. In this capacity, he is responsible for economics and monetary analysis. Before joining the ECB, Stark was at the Deutsche Bundesbank, where he spent two four-year terms as vice president from September 1998 until May 2006. From October 1992 until joining the Bundesbank in 1998, he was at Germany's Federal Ministry of Finance, first as the deputy head of the department on national monetary policy and capital market policy (1992 to 1993), then as head of international monetary and financial relations (1993 to 1994) and then as state secretary of the Federal Ministry of Finance and personal representative of the Federal Chancellor in preparation for the G7/G8 economic summits (1995 to 1998). From 1978 until 1992, he held successive positions at the Federal Ministry of Economics, including work on the General Agreement on Tariffs and Trade and as head of the division on foreign trade and payments, money and foreign currency, and financial

markets. Stark studied economics at the University of Hohenheim and Eberhard Karls University of Tübingen, earning a doctorate in 1975.

James H. Stock has been a professor of economics at Harvard University since 2002, and has chaired the department since 2006. His research focuses on macroeconomic forecasting, monetary policy, and econometric methods for analyzing time series data. Stock's latest work includes an examination of the recent evolution of the U.S. business cycle and the impact of changes in monetary policy on that evolution. At Harvard University's John F. Kennedy School of Government, he was an assistant professor (1983 to 1988) and then associate professor (1988 to 1990) of public policy, a professor of economics at the University of California at Berkeley from 1990 to 1991, and then returned to the Kennedy School as the Roy E. Larsen Professor of Political Economy, a position he held from 1998 to 2002. Stock is a member of the academic advisory board of the Federal Reserve Bank of Boston. He is a fellow of the Econometric Society, a research associate at the National Bureau of Economic Research, and fellow of the American Academy of Arts and Sciences. Stock earned a B.S. in physics at Yale University, an M.A. in statistics at the University of California at Berkeley, and a Ph.D. in economics at the University of California at Berkeley.

Lars E.O. Svensson has been the deputy governor of Sveriges Riksbank since May 2007. Currently he is on leave as professor of economics at Princeton University, which he joined in fall 2001. Svensson's scholarly and policy interests center on macroeconomics, monetary economics, and international finance. Between 1984 and 2003 he was a professor of international economics at the Institute for International Economic Studies, Stockholm University. Svensson was active as advisor to Sveriges Riksbank during 1990–2007 and was a member of the Monetary Policy Advisory Board and the Economic Advisory Panel of the Federal Reserve Bank of New York until his appointment as deputy governor of the Riksbank. In 2000–2001 he undertook a review of monetary policy in New Zealand, commissioned by the New Zealand government, and in 2002 he chaired a committee reviewing monetary policy in Norway. Svensson was chair of the Prize Committee for the Alfred Nobel Memorial

Prize in Economic Sciences during 1999–2001, a committee member from 1993 to 2002, and committee secretary from 1988 to 1992. He is a member of the Royal Swedish Academy of Sciences, a foreign honorary member of the American Academy of Arts and Sciences, a fellow of the Econometric Society, a research associate of the National Bureau of Economic Research, and a research fellow of the Centre for Economic Policy Research, London. Svensson has a B.A. in economics from Stockholm University, an M.S. in mathematics from the Royal Institute of Technology (Stockholm), and a Ph.D. in economics from Stockholm University.

John B. Taylor is the Mary and Robert Raymond Professor of Economics at Stanford University, and the Bowen H. and Janice Arthur McCoy Senior Fellow at the Hoover Institution. He formerly served as the director of the Stanford Institute for Economic Policy Research, where he is now a senior fellow, and he was founding director of Stanford's Introductory Economics Center. Taylor specializes in macroeconomics, monetary economics, and international economics. He is known for his research on the foundations of modern monetary theory and policy, which has been applied by central banks and financial market analysts around the world. From 2001 to 2005, Taylor served as under secretary of the Treasury for International Affairs, where he was responsible for U.S. policies in international finance, including currency markets, trade in financial services, foreign investment, international debt and development, and oversight of the International Monetary Fund and the World Bank. Taylor received the Alexander Hamilton Award for his overall leadership in international finance at the U.S. Treasury. He is a fellow of the American Academy of Arts and Sciences and the Econometric Society. Before joining the Stanford faculty in 1984, Taylor was a professor of economics at Princeton University and at Columbia University. He received an A.B. in economics from Princeton University, and a Ph.D. in economics from Stanford University.

Mark W. Watson is the Howard Harrison and Gabrielle Snyder Beck Professor of Economics and Public Affairs at Princeton University, and is serving as the interim dean at the Woodrow Wilson School of Public and International Affairs. As a scholar, Watson focuses on time-series econo-

metrics, empirical macroeconomics, and macroeconomic forecasting. Before coming to Princeton in 1995, Watson served on the economics faculty at Harvard University and Northwestern University. He is a research associate at the National Bureau of Economic Research, a fellow of the American Academy of Arts and Sciences, and a fellow of the Econometric Society. He has served as a consultant for the Federal Reserve Banks of Chicago and Richmond. Watson did his undergraduate work at Pierce Junior College and California State University at Northridge, and completed his Ph.D. at the University of California at San Diego.

Author Index

- Abraham, Katharine G., 207, 209, 211
Adolfson, Malin, 436
Akerlof, George A., 6
Altig, David, 22, 25
Alvarez, Fernando, 355
Álvarez, Luis J., 306–307
Ang, Andrew, 119, 123, 126–128
Ascari, Guido, 125
Atkeson, Andrew, 101–102, 355
- Baharad, Eyal, 308
Ball, Laurence, 43, 207, 355, 357
 Chari comments and, 383–393
 Galí comments and, 395–400
 hysteresis and, 52–53, 361–381
 imperfect information and, 269
 non-accelerating inflation rate of
 unemployment (NAIRU) and,
 52–53
Bańbura, Marta, 124, 201
Banerjee, Anindya, 122
Barnes, Michelle, 125, 152
Barro, Robert J., 301, 335
Barsky, Robert B., 126
Basu, Susanto, 50, 322
Bates, J. M., 130
Beechey, Meredith, 274
Bekaert, Geert, 119, 123, 126–128
Berardi, Andrea, 127
Bernanke, Ben S., xii, 36, 117, 122,
 455–456
 commodity prices and, 447–451
 inflation forecasts and, 447–450
 labor costs and, 450–451
 natural rate hypothesis and, 84
 output gap and, 451–453
 rational expectations and, 282,
 453–454
 real time data and, 57–58
 role of labor costs and, 57–58
 vector autoregression (VAR) and,
 316
Beveridge, William, 209
Bewley, Truman F., 376–377
Bils, Mark
 Phillips curve and, 26–27, 87, 92
 price determination and, 300, 303,
 306, 320, 333, 344, 352, 356
Blanchard, Olivier, 45, 52
 comments on Dickens model,
 231–236
 continuous-time labor market
 model and, 212
 hysteresis and, 363–364
 insider-outsider wage bargaining
 theory and, 364, 376
 time-varying NAIRU and, 209,
 227
Bleakley, Hoyt, 209
Blinder, Alan S., 58, 308–309, 438,
 454
Bluestone, Barry, 95
Bodie, Zvi, 94
Boivin, Jean, 25, 28, 122, 292, 315–
 316
Bomfim, Antúlio, 269

- Brainard, William C., 109
 Brave, Scott, 118
 Brayton, Flint, 115
 Brunnermeier, Markus K., 261
 Bryan, Michael F., 128
 Burns, Arthur, xii
 Burstein, Ariel T., 322, 356
- Caballero, Ricardo J., 269, 321
 Cafiso, Gianluca, 325
 Calvo, Guillermo A., 15–16
 Camba-Mendez, Gonzalo, 115
 Cameron, Norman, 125
 Campbell, Jeffrey R., 312
 Canetti, Elie D., 308–309, 454
 Canova, Fabio, 109, 118–119, 123, 126
 Caplin, Andrew S., 49, 313, 353
 Carlton, Dennis W., 26
 Carroll, Christopher D., 269
 Carson, Joseph, 93
 Carvalho, Carlos M., 297, 352
 Cass, David, 389, 391
 Cecchetti, Stephen G., 26, 112, 114–115, 128, 291
 Chari, V. V., 53, 92–93, 379, 383–393
 Christiano, Lawrence J., 17, 21–22, 268–269, 341
 Chu, Rita S., 115
 Clark, Todd E., 120, 123
 Clements, Michael P., 123
 Cogley, Tim, 30–31, 93, 109, 259
 Cohen, Jessica, 207, 211, 227
- D'Agostino, Antonello, 122, 198, 203
 De Mol, Christine, 197–199
 Dennis, Richard, 341–343
 DeWachter, Hans, 127
 Dew-Becker, Ian, 152
 Dhyne, Emmanuel, 302–303
 Diamond, Peter A., 209, 212
 Dickens, William T.
 historical perspective on Phillips curve and, 6, 36, 44–45
 time-varying NAIRU and, 207–230
- Dicks-Mireaux, L. A., 237, 241
 Diebold, Francis X., 120
 Diron, Marie, 109, 120
 Dotsey, Michael, 307, 322, 355
 Dow, J. C. R., 237, 241
 Dupasquier, Chantal, 125
- Ebrahimy, Eshan, 240
 Eckstein, Otto, 267
 Eden, Benjamin, 308, 312
 Edge, Rochelle, 268
 Edmond, Chris, 355
 Eichenbaum, Martin, 17, 21–22, 268–269, 302, 311–312, 328n34, 341
 Eichmeier, Sandra, 122–123
 Eliasziw, Piotr, 122, 316
 Engel, Eduardo M. R. A., 321
 Erceg, Christopher J., 269–270
 Estrella, Arturo, 109
 Evans, Charles, 17, 21–22, 268–269, 341
 Evans, George W., 259
- Fabiani, Silvia, 308–309
 Fischer, Stanley
 historical perspective on Phillips curve and, 12–13, 24, 36, 54–56, 78
 Israeli monetary policy and, 403–413
 Fisher, Jonas D. M., 117–118, 126, 143, 191, 328n34
 Fisher, Timothy C. G., 306
 Ford, Gerald, 75
 Frenkel, Jacob, 404, 410
 Friedman, Milton, xiii, 7–11, 37
 accelerationist Phillips curve and, 361
 AEA presidential address of, 74–77, 361
 augmented expectations and, 79
 commodity prices and, 85
 hysteresis and, 361–363, 397

- natural rate hypothesis and, 51, 84
 rational expectations and, 89–90, 249
 reverse-direction Phillips curve and, 40
- Fuhrer, Jeff
 Beveridge curve and, 209
 historical perspective of Phillips curve, 3–68
 inflation forecasts and, 109
 large price changes and, 88
 persistent inflation and, 79–80, 341
 rational expectations and, 269
- Fujita, Shigeru, 241
- Gagnon, Etienne, 308
- Galí, Jordi, 379
 comments on hysteresis, 395–400
 Phillips curve and, 17–18, 53, 81
 price determination and, 319, 341–344
- Gavin, William T., 122
- Gertler, Mark, 17–18, 319, 341–344
- Geweke, John, 121
- Giacomini, Raffaella, 122
- Giannone, Domenico, 122, 124, 197–203
- Giannoni, Mark P., 25, 28, 292, 315
- Glassman, James E., 112
- Golosov, Mikhail
 historical perspective on Phillips curve and, 27–28, 49
 price determination and, 307–308, 311, 320–321, 324, 353–354
- Goodfriend, Marvin, 282
- Gordon, Robert, 11, 59nn7,8, 61n24
 non-accelerating inflation rate of unemployment (NAIRU) and, 129, 207–208
 step function lag restriction and, 128–129
 triangle model and, 102, 106–107
 (*see also* Triangle model)
- Granger, Clive W. J., 130
- Gray, Jo Anna, 12, 78
- Greenspan, Alan, xii–xiii, 95, 97
- Guerrieri, Luca, 80
- Gürkaynak, Refet S., 274
- Hall, Robert E., 335
- Hansen, Peter Reinhard, 121, 209
- Hellwig, Christian, 322, 356
- Hendry, David F., 123, 127
- Hernandez, Judy, 210
- Hernando, Ignacio, 306
- Honkapohja, Seppo, 259
- Hooper, Peter, 90
- Hubrich, Kirstin, 109, 127
- Inoue, Atsushi, 123–124
- Ireland, Peter, 93
- Jaditz, Ted, 112
- Jaimovich, Nir, 302, 311–312
- Johannsen, Ben, 274, 277
- Jorion, Philippe, 126
- Julliard, Christian, 261
- Kahneman, Daniel, 260
- Kashyap, Anil K., 26, 291
- Kehoe, Patrick J., 299–300, 352–355
- Kendrick, John W., 384
- Khan, James A., 356
- Kiley, Michael T., 47, 56, 92, 267–270
- Kilian, Lutz, 123–124
- King, Robert G., 50, 307, 322, 333–350, 355
- King, Stephen R., 11
- Klenow, Peter J., 50, 87, 92
 Bureau of Labor Statistics data set and, 26–27
 Calvo model and, 27
 CPI Research Database and, 295
 forced item substitution and, 297–300, 303, 310
 microeconomic price data and, 294–300, 303, 306, 310–311, 320, 322, 333–334, 336, 342, 344, 347, 352, 354–355

- Kliesen, Kevin L., 122
 Kodrzycki, Yolanda K., 3–68
 Kohn, Donald L., 36, 53–55, 415–426
 Konieczny, Jerzy D., 306, 308
 Konstant, Raymond A., 210
 Koopmans, Tjalling C., 389, 391
 Kozicki, Sharon, 29–31
 Krugman, Paul, 363
 Kryvtsov, Oleksiy, 50
 Calvo model and, 27
 forced item substitution and, 27, 297–300, 310
 microeconomic price data and, 294–300, 310–311, 323, 334, 336, 342, 347, 354, 356
 Kydland, Finn E., 389, 391
- Lach, Saul, 308
 Laforte, Jean-Philippe, 268
 Laséen, Stefan, 436
 Lebow, David E., 308–309, 454
 Leeson, Robert, 89–90
 Lenza, Michele, 200–201
 Levin, Andrew, 109, 269–270, 274, 277
 Linde, Jesper, 328n34, 436
 Lipsey, Richard, 237
 Little, Jane Sneddon, 3–68
 Liu, Chin Te, 117–118, 126, 143, 191
 Llaudes, Ricardo, 376
 Lucas, Robert, 10, 27–28, 51
 hysteresis and, 383
 imperfect information and, 270
 menu cost model and, 49
 microeconomic price data and, 307–308, 311, 320–324, 353–354
 New Keynesian Phillips curve and, 250–251
 rational expectations and, xiii, 9–12, 78, 267–269
 shock analysis and, 54
 Lyrio, Marco, 127
- Maćkowiak, Bartosz
 behavioral analysis and, 261
 historical perspective on Phillips curve and, 25, 48–51, 55
 King comments on, 333–348
 microeconomic price data implications and, 291–332
 Midrigan comments on, 351–357
 Mankiw, N. Gregory, xiii, 207, 384
 historical perspective on Phillips curve and, 23–25, 36, 38–39, 43
 Intermittent Observation theory and, 260
 learning from Phillips curve, 71–98
 microeconomic price data and, 313, 324, 335
 non-accelerating inflation rate of unemployment (NAIRU) and, 363
 Marcellino, Massimiliano, 122, 125
 Mariano, Roberto S., 120
 Marrocu, Emanuela, 125
 Matejka, Filip, 261
 Mazunder, Sandeep, 379
 McCallum, Andrew, 316
 McCracken, Michael W., 120, 123
 Medoff, James L., 209
 Meltzer, Allan, 95–97, 406
 Midrigan, Virgiliu, 50, 299–300, 321, 323, 325, 351–357
 Mihov, Ilian, 25, 28, 292, 315
 Mishkin, Frederic, 107, 126
 Moench, Emanuel, 48–49, 292, 313–315
 Mojon, Benoît, 109, 120
 Moore, George, 19, 79–80, 269, 341
 Moshiri, Saeed, 125
 Murry, Elizabeth, 325
 Muth, John, 9, 277n1
- Nakamura, Emi, 27, 50
 microeconomic price data and, 295–301, 309–311, 322, 324, 334, 345, 347, 352
 policy issues and, 271–272

- Nelson, Charles R., 30, 129
 Ng, Serena, 122
 Nickell, Stephen, 363
 Nixon, Richard M., xii
 Nunziata, Luca, 363
- Ochel, Wolfgang, 363
 Ohanian, Lee E., 101–102
 Okun, Arthur M., 309
 Olivei, Giovanni P., 3–68, 125
 Orphanides, Athanasios, 32, 56–57, 118
 comments on Sims by, 281–288
 New Keynesian Phillips curve and, 251
- Paciello, Luigi, 25
 Pagan, Adrian Rodney, 41–42, 187–193
 Perry, George L., 6, 109, 207
 Pesaran, M. Hashem, 190
 Petrongolo, Barbara, 241
 Phelps, Edmund, 8–9, 75, 78, 110–111, 397
 Phillips, Alban W.
 background of, 72
 basic correlation paper of, 5–6
 cost-of-living effect and, 5
 empirical statement of, 76, 247
 historical data set of, 39
 intuition of, 74–75
 reactions to paper of, 36–37
 speed limit effect and, 5
 Pick, Andreas, 190
 Piger, Jeremy M., 109
 Pissarides, Christopher A., 45–46, 232, 237–244
 Pivetta, Frederic, 109
 Posen, Adam, 207, 211, 227
 Prescott, Edward C., 389, 391
 Primiceri, Giorgio E., 6, 32–33, 109, 248
- Ramey, Garey, 241
 Rapping, Leonard A., 46, 248
- Rebelo, Sergio, 302, 311–312
 Reichlin, Lucrezia, 42, 124, 195–203
 Reis, Ricardo, 313
 historical perspective on Phillips curve and, 23–25, 49
 imperfect information and, 269, 270
 inflation forecasts and, 109, 127
 Intermittent Observation theory and, 260
 microeconomic price data and, 316, 324
- Ricketts, Nicholas, 125
 Roberts, John, 109–110, 115
 Robinson, Tim, 191
 Rodriguez-Palenzuela, Diego, 115
 Romer, Christina D., 32, 119, 128, 367, 378
 Romer, David H., 32, 119, 128, 355, 357, 367, 378
 Rossi, Barbara, 120
 Rotemberg, Julio
 historical perspective on Phillips curve and, 15–16, 19–20
 micro-founded Phillips curve and, 19–21
 price determination and, 336, 339, 356
 rational expectations and, 268–269
- Rowe, Tazewell S., 123
 Rudd, Jeremy B., 18–19, 308–309, 454
 Rünstler, Gerhard, 109
- Sahuc, Jean-Guillaume, 316
 Samuelson, Paul A., xi–xiii, 6–7, 36–37, 71, 73, 95
 Sargent, Thomas J. (Tom), xiii, 248
 inflation forecasts and, 109, 121
 lag principle of, 339
 Phillips curve and, 6, 10–14, 33
 rational expectations and, 259
- Sayers, Chera, 112
 Sbordone, Argia, 17, 30–31, 93, 251
 Scheinkman, José, 261

- Sekhposyan, Tatevik, 120
- Shannon, Claude, 260
- Shimer, Robert, 240–241
- Siebert, Horst, 368
- Sims, Christopher A.
 Atkeson-Ohanian model and, 119
 imperfect information and, 269–271
 inflation forecasts and, 109, 121, 128, 247–265
 Kiley comments on, 267–270
 microeconomic price data and, 313, 323
 Orphanides comments on, 281–288
 Phillips curve and, 25, 39, 46–47, 51
- Skrzypacz, Andrzej, 308
- Smets, Frank, 268
 DSGE models and, 293–294, 316–320 (*see also* Dynamic stochastic general equilibrium (DSGE) models)
 King comments on, 333–348
 microeconomic price data implications and, 291–332
 Midrigan comments on, 351–357
 Phillips curve and, 21–22, 48–51, 55
- Smith, Julie K., 128
- Solow, Robert M., xiii, 6–7, 389, 391
 learning from Phillips curve, 71–98
 New Keynesian Phillips curve and, 36–39
- Spulber, Daniel F., 49, 313, 353
- Staiger, Douglas, 43, 129, 207–208
- Stark, Jürgen, 36, 55–56, 427–432
- Startz, Richard, 227
- Stavrakeva, Vania, 227
- Steindel, Charles, 115
- Steinsson, Jón, 27, 50
 microeconomic price data and, 295–301, 309–311, 322, 324, 334, 345, 347, 352
 policy issues and, 272
- Stock, James H., 39, 207
 gap variables and, 108
 inflation forecasts and, 101–179
 natural rate hypothesis and, 84
 non-accelerating inflation rate of unemployment (NAIRU) and, 43, 92, 207–208
 Pagan comments and, 187–193
 rational expectations and, 284
 Reichlin comments and, 195–203
 unobserved components-stochastic volatility (UC-SV) model and, 40–41, 101, 105, 116, 131–141, 145–153, 157–158, 163, 168, 173, 187
- Stockhammer, Engelbert, 379
- Stockton, David J., 112
- Stone, Andrew, 191
- Summers, Lawrence H., 52, 364, 376
- Surico, Paolo, 198, 203
- Svensson, Lars E. O.
 Phillips curve and, 36, 55–56, 63n42
 Riksbank policy and, 433–444
- Swanson, Eric T., 274
- Taylor, John B., xii–xiii, 420
 data sources and, 291
 feedback rule and, 20
Handbook of Microeconomics and, 291
 historical perspective on Phillips curve and, 12–13, 15, 20, 29, 36–39, 46, 48, 81
 inflation forecasts and, 109
 learning from Phillips curve, 71–98
 microeconomic price data and, 325n1
 New Keynesian Phillips curve and, 37, 250–251
 rational expectations and, 268–269
 staggered contracts and, 12, 37, 325n1
- Thatcher, Margaret, 249
- Timmermann, Allan, 123, 130, 190

- Tinsley, Peter A., 29–31
 Tkacz, Greg, 125
 Tobin, James, xiii
 Tootell, Geoffrey M. B., 34–35
 Trichet, Jean-Claude, 282
 Tsiddon, Daniel, 308
 Tversky, Amos, 260
- Valletta, Robert, 209, 211, 227
 van Norden, Simon, 118
 van Zyl, Marileze, 191
 Vermeulen, Philip, 303, 306
 Villani, Mattias, 436
 Volcker, Paul, xii, 40, 95, 97, 252, 270, 361
- Wallace, Neil, 10–11, 13–14, 78
 Watson, Mark W., 39
 gap variables and, 108
 inflation forecasts and, 101–179
 natural rate hypothesis and, 84
 non-accelerating inflation rate of unemployment (NAIRU) and, 43, 92, 207–208
 Pagan comments and, 187–193
 rational expectations and, 284
 Reichlin comments and, 195–203
 unobserved components-stochastic volatility (UC-SV) model and, 40–41, 101, 105, 116, 131–141, 145–153, 157–158, 163, 168, 173, 187
 Webb, Roy H., 123
 Wei, Min, 119, 123, 126–128
- Whelan, Karl, 18–19
 Wiederholt, Mirko
 behavioral analysis and, 261
 historical perspective on Phillips curve and, 25, 48–49, 51
 microeconomic price data and, 292, 299, 309, 313–316, 323–324
 Wiggins, Rodney L., II, 128
 Williams, John C., 32, 47, 115, 269, 284
 Williams, Noah, 6, 33, 259
 Willis, Jonathan L., 299, 322, 355
 Wingard, Irvin F., 210
 Wolfers, Justin, 363
 Wolman, Alexander L., 307, 337
 Woodford, Michael, 19–21, 250, 272, 313, 356
 Wouters, Raf, 21, 22, 268, 341. *See also* Dynamic stochastic general equilibrium (DSGE) models
 Wright, Jonathan, 123–124, 274, 277
- Xiong, Wei, 261
- Yun, Tack, 337
- Zargorsky, Jay L., 210
 Zbaracki, Mark J., 308–310
 Zha, Tao, 6, 33, 109, 259
 Zhou, Ruilin, 117–118, 126, 143, 191
 Ziegler, Christina, 122–123

Subject Index

- Accelerationist Phillips curve, 8, 111, 361, 377, 396
- Adaptive expectations, 37, 40, 283
- Ad hoc mechanisms, 9–12
- “Aggregate Dynamics and Staggered Contracts” (Taylor), 12–13
- Aggregate supply/demand
- adaptive expectations and, 37, 40, 283
 - alternatives to, 36
 - DSGE models and, 20–22
(*see also* Dynamic stochastic general equilibrium (DSGE) models)
 - forced item substitution and, 297–303, 310, 312
 - Friedman and, 361
 - global trends and, 420–422
 - hysteresis and, 361–378
 - increased commodity prices and, 85–89 (*see also* Prices)
 - inflation forecasts and, 39–42
 - labor costs and, 450–451
 - microeconomic price data and, 291–325
 - OPEC oil shock and, xi (*see also* Shocks)
 - policy and, 361 (*see also* Policy)
 - Rotemberg-Woodford model and, 19–21
- Akaike Information Criterion (AIC), 105
- American Economic Association (AEA), 7, 73–77, 361
- American Economic Review*, 71, 76
- “Analytical Aspects of Anti-Inflation Policy” (Samuelson), 71
- Answers.com, 210
- ARMA structure, 21, 118–119, 127, 178n7
- Atkeson-Ohanian model
- accuracy of, 117
 - Bernanke and, 117, 122
 - equation for, 105
 - G7 countries and, 118–119
 - Greenbook forecasts and, 117, 119
 - inflation forecasts and, 105, 115–128, 143, 152, 161
 - performance of, 116–128, 143
 - resuscitating multivariate forecasts and, 121–123
 - vector autoregression (VAR) and, 118–119
- Australia, 191, 365, 374
- Autoregressive (AR) models, 189
- inflation forecasts and, 105, 107, 111–112, 115–128
 - performance of, 132
 - VAR, 118–119 (*see also* Vector autoregression (VAR))
- Autoregressive distributed lag (ADL) model
- comparative performance of, 144–145

- equation for, 129
 - inflation forecasts and, 107, 111–112, 114, 126, 129, 132–162, 166–167, 171–173, 176–177
 - using other predictors, 129
- Autoregressive integrated moving average (ARIMA)
 - models, 104, 178n4, 189
- Backward-looking Phillips curve
 - historical perspective on, 11, 17, 21, 24, 32–35, 39–40, 60n12, 18, 62n35, 81
 - inflation forecasts and, 101–102, 106, 111, 115
 - price determination and, 316, 319, 334, 339–342
- Bagging, 123–124
- Bank of England, xi–xii, 274–275, 410
- Bank of Israel, 55, 62n38, 62n40, 404, 406, 410
- Bayesian Model Averaging (BMA), 123–124
- Bayesian models, 342
 - Riksbank and, 434
 - likelihood, 21
 - shrinkage, 201
 - unobservable index, 313–314
- Bayes Information Criterion, 128–129
- Behavioral economics, 47, 259–260
- Belgium, 368
- Beveridge curve, 44–46
 - Blanchard comments and, 231–236
 - Bureau of Labor Statistics and, 210, 234
 - business cycle fluctuations and, 241–242
 - confidence intervals for, 223–224
 - Dickens model and, 212–227, 231–243
 - equilibrium and, 213–214, 238
 - help-wanted index and, 209–210
 - job vacancies and, 209–211, 221, 226, 240
 - Monte Carlo trials and, 224–225
 - non-accelerating inflation rate of unemployment (NAIRU) and, 208–227, 237–242
 - shocks and, 238–239
 - trade-off stability and, 210–211
 - UV-analysis and, 237
 - vacancy-unemployment relationship and, 209, 214–216, 232–233, 237–239
- Beveridge-Nelson model, 30
- Blue Chip surveys, 119
- Boston University, 94
- Bureau of Labor Statistics (BLS)
 - Beveridge curve and, 210, 234
 - CPI Research Database of, 295
 - metropolitan data sets and, 295
 - microeconomic price data and, 295–301, 307–308, 310, 320, 351, 354
 - Phillips curve and, 26–27, 45, 49–50
- Business cycles, 337
 - Beveridge curve and, 241–242
 - Great Moderation and, xii, 41–42, 200
 - hysteresis and, 383–384, 389
 - unemployment and, 241–242
- Calvo fairy, 15, 17
- Calvo model
 - Cogley-Sbordone model and, 30
 - contract price and, 60n12
 - dynamic indexation and, 342–344
 - dynamic stochastic general equilibrium (DSGE) models and, 50, 293
 - forward-looking price setting and, 338–339
 - magnitude of adjustment and, 345–347
 - micro-founded inflation dynamics and, 23–24, 26

- microeconomic price data and, 292–293, 311–312, 316–319, 325, 327n32, 328n24, 334–348, 354
- New Keynesian Phillips curve and, 31, 37, 77–79, 250–251
- persistent inflation and, 341–344
- policy and, 79
- price level dynamics and, 15, 17, 28–29, 46, 338, 340
- sticky prices and, 62n34, 334–341
- Careerbuilder.com, 210
- Central banks. *See also* Federal Reserve
 - Bank of England, xi–xii, 274–275, 410
 - Bank of Israel, 55, 62nn38,40, 404, 406, 410
 - Bernanke on, 447, 453–454, 457
 - Czech National Bank, 441
 - Deutsche Bank, 90
 - dual targeting role of, xi–xii
 - dynamic stochastic general equilibrium (DSGE) models and, 55–56, 338
 - European Central Bank, xi–xiii, 48, 55, 89, 109, 302, 306, 427–432
 - Federal Open Market Committee (FOMC) and, 416
 - forward-looking analysis and, 56
 - global trends and, 420–422
 - historical perspective on Phillips curve and, 16, 19, 29, 36, 41–42, 47, 52–57, 63n42, 87, 89
 - hysteresis and, 367–371, 378
 - inflation forecasts and, 41–42, 109, 115, 188
 - informal averaging by, 55
 - Lucas critique and, 403, 407, 412, 419, 423n6
 - Norges Bank, 438
 - Phillips curve framework for, 415–426
 - policy and, 53–57, 401–444 (*see also* Policy)
 - prices and, 56–57, 302, 333, 338, 415–423
 - rational expectations and, 249, 258, 267, 269, 272, 282–287
 - recessions and, xii, 367
 - Reserve Bank of New Zealand (RBNZ), 438
 - Riksbank, xiii, 55–56, 63n42, 433–444
 - Sedlabanki Islands, 441
 - shocks and, 417–420
 - structural models and, 419–420
 - time-varying NAIRU and, 207
 - welfare costs and, 400, 415–416
- Chicago Fed National Activity Index (CFNAI), 117, 121–122, 143
- Clothing, 298–301
- Combination forecast models, 123–124, 129–131
- Comfort zone, 3–4
- Communism, 262, 281
- Congressional Budget Office, 207
- Consensus Forecast, 59n1
- Consumer Price Index (CPI)
 - ARIMA models and, 104, 178n4, 189
 - Bureau of Labor Statistics and, 295
 - Dickens model and, 219, 222
 - forecasts and, 102, 104, 108–128, 132n3, 143–144, 178n7, 179n10
 - inflation and, 35, 102, 104, 108–128, 132n3, 143–144, 178n7, 179n10, 197, 219, 222, 274, 295, 326n14, 334, 342, 404, 406, 433
 - Harmonized Index of Consumer Prices and, 127, 316
 - Israel and, 404, 406
 - microeconomic price data and, 295, 334, 342, 345
 - multivariate models and, 121–128
 - performance of models and, 132, 143–144, 197–198
 - Research Database for, 295

- Riksbank and, 433
- unobserved components-stochastic volatility (UC-SV) model and, 153–162
- Consumption
 - dynamic stochastic general equilibrium (DSGE) models and, 20–22
 - personal consumption expenditures (PCEs) and, 35, 89, 102, 106–111, 114, 117–120, 123, 127–128, 132, 143–144
- Continuous-time labor market model, 212
- Contracts
 - Fuhrer-Moore model and, 17–18
 - staggered, 12–13, 24, 37, 325n1
 - utility maximization and, 15–16
- Cost-of-living effect, 5
- Council of Economic Advisors, 112
- Credit market, xii
- Creeping inflation, 73
- Customer costs, 308–310
- Czech National Bank, 441

- Data Resources, Inc., 119
- Deutsche Bank, 90
- Dickens model
 - Beveridge curve estimation and, 216–225
 - Blanchard comments on, 231–236
 - Blanchard-Diamond model and, 212
 - business cycle fluctuations and, 241–242
 - capital cost and, 213
 - confidence intervals for, 223–224
 - Consumer Price Index (CPI) and, 219, 222
 - equations of, 212–214, 219, 231–233, 238–239
 - equilibrium and, 213–216
 - Job Openings and Labor Turnover Survey (JOLTS) and, 221, 226
 - Kalman filter and, 219, 221
 - Monte Carlo trials and, 224–225
 - NAIRU estimation and, 212–227
 - natural rate hypothesis and, 233–234
 - Phillips curve estimation and, 216–225
 - Pissarides comments on, 237–244
 - policymakers and, 225
 - vacancy-unemployment relationship and, 214–216, 232–233
- Disinflation, 424n11
 - historical perspective on, 4, 24, 32–33, 52, 75, 80, 87
 - hysteresis and, 366–367, 371–377, 379n3
 - inflation forecasts and, 109–110, 143
 - rational expectations and, 251, 269–272
 - sacrifice ratio and, 33
- Dixit-Stiglitz aggregator, 318
- Dummy variables, 106, 208, 219–221, 227, 263n4
- Dynamic factor models, 121–123
- Dynamic indexation model, 338, 340, 342–344
- Dynamic stochastic general equilibrium (DSGE) models
 - business cycles and, 337
 - Calvo parameter and, 293
 - central banks and, 55–56, 338, 410–412
 - Dixit-Stiglitz aggregator and, 318
 - European Central Bank and, 420
 - historical perspective on, 20–23, 27, 49–50, 55–56
 - Israel and, 410–412
 - Kimball aggregator and, 318
 - microeconomic price data and, 293–294, 316–319, 321, 325, 334–341, 347
 - rigidities and, 321
 - Smets-Wouter model and, 293–294, 316–320

- sticky prices and, 50, 334–341
 - vector autoregression (VAR) and, 316
- Economica* journal, 72, 74
- Entry Level Items (ELIs), 295, 297, 301–302
- Equal opportunity laws, 210
- Equations
- accelerationist Phillips curve, 361, 377, 396
 - Atkeson-Ohanian model, 105
 - autoregressive distributed lag (ADL) model, 107
 - autoregressive forecast, 105
 - backward-looking Phillips curve, 32
 - Beveridge curve, 239
 - Calvo-Rotemberg model, 16
 - change in inflation rate, 8
 - combination forecasts, 130
 - Dickens model, 212–214, 219, 231–233, 238–239
 - exponentially weighted moving average (EWMA), 189
 - forward-looking price setting, 338–339
 - hybrid New Keynesian Phillips curve, 18
 - inflation as marginal cost function, 17
 - inflation persistence, 341
 - interest rates, 20
 - labor income share, 17
 - long-run target value, 20
 - New Keynesian Phillips curve, 16, 18
 - non-accelerating inflation rate of unemployment (NAIRU), 9, 43, 129, 381
 - optimization, 16, 20
 - Phillips curve forecast, 106–107
 - Phillips's log-log wage/unemployment, 5
 - price, 9–10, 338
 - relative predictability, 196
 - root mean squared error (RMSE), 104
 - second-generation Phillips curve, 389
 - Sims model, 252
 - sticky-information, 24
 - time-varying inflation targets, 30
 - triangle model, 106, 129
 - unobserved components-stochastic volatility (UC-SV) model, 40
 - vector autoregression (VAR), 202
- Equilibrium
- Beveridge curve and, 213, 238
 - Dickens model and, 213–216
 - dynamic stochastic general equilibrium (DSGE) models and, 20–23, 27, 49–50, 55
 - hysteresis and, 361–378
 - inflation forecasts and, 46–48
 - information-processing ability and, 25–26
 - integrating Phillips curve into new models and, 19–23
 - micro-founded Phillips curve and, 19–21
 - natural rate of unemployment and, 8–9
 - rational expectations and, 78, 268
- Euro area
- forced item substitution and, 303
 - frequency of price changes and, 292, 302–306, 310, 326n2
 - hysteresis and, 399
 - microeconomic price data and, 292–293, 302–306, 316–319
 - non-accelerating inflation rate of unemployment (NAIRU) and, 363, 365
- European Central Bank (ECB), xi–xiii, 55, 89, 115
- analytical framework of, 429–432
 - credit data and, 430–431

- dynamic stochastic general
 - equilibrium (DSGE) models and, 420
- Governing Council of, 429, 432
- Inflation Persistence Network and, 48, 109, 302, 306, 308, 322, 333, 351
- New Keynesian Phillips curve and, 428–429
- policy and, 282, 427–432
- Exchange Rate Mechanism, 368–369
- Exchange rates, 50, 413n2
 - hysteresis and, 368–369, 386
 - Israel and, 404
 - microeconomic price data and, 355
- Exogenous timing models, 320–321
- Exponentially weighted moving average (EWMA) formula, 189–191

- Federal Open Market Committee (FOMC), 416
- Federal Reserve, 267
 - Bernanke and, xii, 36, 57–58, 84, 117, 122, 282, 316, 447–454
 - Burns and, xii
 - Chicago Fed National Activity Index and, 117, 121–122, 143
 - comfort zone of, 3–4
 - credibility of, 4
 - employment maximization and, 415–416
 - Greenbook forecasts and, 117, 119, 128, 179n8, 195–199
 - Greenspan and, xii–xiii, 95, 97
 - Guerrieri and, 80
 - imperfect information and, 270
 - inflation forecasts and, 40
 - Livingston Survey and, 119
 - non-accelerating inflation rate of unemployment (NAIRU) and, 33
 - Philadelphia Fed's Livingston Survey and, 119
 - Phillips curve conference of, xi, xiii, 4, 36, 53
 - price stability and, 415–416
 - Quarterly Review* and, 81
 - recession policy of, 367, 378
 - sacrifice ratio and, 33
 - time-varying inflation targets and, 29–30
 - Volcker and, xii, 40, 95, 97, 252, 270, 361
- Feedback rule, 20
- Finland, 374
- Food prices, 3, 295, 298, 300–301, 418, 452
- Forced item substitutions
 - categories of, 298–299
 - clothing and, 298–299
 - criteria for recording, 297–298
 - discontinued samples and, 297–298
 - euro area and, 303
 - microeconomic price data and, 297–303, 310, 312, 326n11, 327n20
 - transportation and, 298–299
 - United States and, 297–303
- Forward-looking Phillips curve, 429
 - historical perspective on, 16, 18–24, 30, 33, 56, 60n18
 - inflation forecasts and, 198, 203
 - price determination and, 324, 334–335, 339–342
- France, 316, 361, 369–370
- Fuel prices, 295, 298–299
- Fuhrer-Moore model, 17–18, 79–80

- Gap variables, 108
- Germany, 316, 327n20
- Great Depression, 4, 384, 386
- Great Moderation, xii, 41–42, 200
- Greenbook forecasts, 117, 119, 128, 179n8, 195–199
- Gross domestic product (GDP) data, 78
 - consumption deflator and, 222
 - counterfactual exercise for, 201–202
 - Great Depression and, 384

- hysteresis and, 384
 - inflation forecasts and, 102, 106, 109, 111–112, 120, 173–177, 196–202, 222
 - Israel and, 403–404
 - New Keynesian Phillips curve and, 252
 - non-accelerating inflation rate of unemployment (NAIRU) and, 207
 - rational expectations and, 260
 - Riksbank and, 441
 - root mean squared error (RMSE) model and, 173–177
 - stabilization of, 74–75
 - World War II and, 384
- Handbook of Macroeconomics* (Taylor), 291
- Harmonized Index of Consumer Prices, 127, 316
 - Hazard rates, 310–311
 - Help-wanted index, 209–210
 - Hodrick-Prescott filter, 62n40, 192, 364, 381, 397, 410
 - Household surveys, 90–91
 - inflation forecasts and, 119
 - microeconomic price data and, 308–310
 - rational expectations and, 272, 274
 - “How to Whip Inflation Now” (Ford), 75
 - H-period inflation, 103–104, 178n1
 - Hysteresis, 52–53
 - accelerationist Phillips curve and, 361, 377, 396
 - alternative approaches and, 398–399
 - business cycles and, 383–384, 389
 - Chari comments and, 383–393
 - disinflation and, 366–367, 371–377, 379n3
 - euro area and, 399
 - European Central Bank, 427–432
 - Exchange Rate Mechanism and, 368–369
 - explaining mechanism of, 376–377
 - further research for, 375–378
 - Galí comments on, 395–400
 - Great Depression and, 384, 386
 - Hodrick-Prescott (HP) filter and, 364, 381, 397
 - insider-outsider wage bargaining theory and, 364, 376
 - IS curve and, 368
 - Lawson boom and, 368–369
 - Lucas and, 383
 - measuring degree of, 377
 - natural rate of unemployment and, 361–364, 378, 383, 396–400
 - new evidence for, 369–375
 - non-accelerating inflation rate of unemployment (NAIRU) and, 362–378, 383, 395
 - nonlinearities and, 377
 - policy and, 367–368, 378, 383
 - previous evidence of, 366–369
 - responses to recessions and, 367–368
 - shocks and, 363, 365, 369, 381n3, 397, 400
 - state-dependence and, 377
 - United Kingdom and, 368–369, 373–374
 - United States and, 262, 363–364, 373, 399
 - Iceland, 441
 - Imperfect information models, 269–272
 - Impulse response, 292, 313–316
 - Inertia
 - built-in nature of, 79–81
 - central bank and, 267
 - persistent inflation and, 18–19
 - prices and, 10–12, 47, 60n18, 62n36, 79, 81, 240, 262, 351, 354, 356

- rational expectations and, 36–37, 259–261
- shocks and, 22–26
- sticky-information and, 24
- uncertainty modeling and, 259–261
- without Phillips curve, 257
- Inflation, xi–xiii
 - augmented expectations and, 79
 - central banks and, 29, 417–425 (*see also* Central banks)
 - comfort zone and, 3–4
 - compensation and, 274
 - consumer price index (CPI), 35 (*see also* Consumer price index (CPI))
 - core measures of, 3
 - cost-push, 73
 - creeping, 73
 - demand-pull, 73
 - different measures of, 40–41
 - disinflationary pressure and, 4 (*see also* Disinflation)
 - Fuhrer-Moore model and, 17–18
 - Great Moderation and, 42
 - historical perspective on, 386–393
 - household surveys and, 90–91
 - h-period, 103–104, 178n1
 - hysteresis and, 361–400
 - information costs and, 23–25
 - Israel and, 404–412
 - lagged, 9, 18–24, 30, 34 (*see also* Lagged inflation)
 - microeconomic price data and, 291–325
 - micro-founded alternatives and, 23–26
 - monetary expansion and, 8
 - NAIRU and, 207 (*see also* Non-accelerating inflation rate of unemployment (NAIRU))
 - natural rate of unemployment and, 51–53
 - New Keynesian Phillips curve and, 16–19, 250–257
 - output gap and, 19–20, 24 (*see also* Output gap)
 - persistent, 46–48, 79–80, 109, 302, 306, 308, 322, 333, 341–344, 351
 - personal consumption expenditures (PCEs) and, 35, 89, 102, 106–111, 114, 117–120, 123, 127–128, 132, 143–144
 - Phillips paper and, 5–6
 - rational expectations and, 267–269 (*see also* Rational expectations)
 - real-time policymaking and, 451–453
 - recession and, 3, 24, 41, 51, 61n29, 73, 80, 102, 114, 143, 186–188, 192, 232, 263n6, 367–368, 378, 403
 - sacrifice ratio and, 33
 - specification approaches to, 17
 - stagflation and, xii–xiii, 3–4, 284, 286
 - steady-state, 30
 - time-varying targets and, 29–31
 - unanswered questions on, 57–58
 - welfare costs and, 48, 51, 263, 325, 400, 415–416
- Inflation determination. *See also*
 - Rational expectations
 - historical perspective of Phillips curve and, 247–249
 - inertia models and, 259–261
 - New Keynesian Phillips curve and, 250–257
 - with Phillips curve and, 247–249
 - Primiceri model and, 248
 - without Phillips curve, 257–259
- Inflation forecasts, 39–42, 79
 - Akaike Information Criterion and, 105
 - ARIMA models and, 104, 178n4, 189
 - Atkeson-Ohanian model and, 101–102, 105, 115, 116–128, 143, 152, 1161
 - autoregressive (AR) models and, 105, 107, 111–112, 115–116, 118–128, 189

- autoregressive distributed lag (ADL)
 - model and, 107, 111–112, 114, 126, 129, 132–162, 166–167, 171–173, 176–177
- backward-looking Phillips curve
 - and, 101–102, 106, 111, 115
- based on forecasts of others, 107–108
- based on other predictors, 108
- based on past inflation, 104–106
- Bayes Information Criterion and, 128–129
- Bernanke and, 447–450
- changes in univariate performance
 - and, 128–144
- combination forecast models and, 123–124, 129–131
- commodity price shocks and, 417–418
- counterfactual exercise for, 201–202
- data sets and, 108
- decreasing relative predictability
 - and, 196–201
- difficulty of, 101
- direct, 103–107, 117–118, 122, 127–129, 179n8
- dynamic factor models and, 121–123
- exponentially weighted moving average (EWMA) formula and, 189–191
- gap variables and, 108
- GDP and, 102, 106, 109, 111–112, 120, 173–177, 196–202, 222
- good in-sample fit and, 178n1
- Greenbook, 117, 119, 128, 179n8, 195–199
- household surveys and, 119
- h-period inflation, 103–104, 178n1
- Israel and, 407, 409
- iterated, 103, 105–108, 112, 128
- literature survey on, 109–128
- mean squared error (MSE) and, 112, 130
- Mishkin specification and, 107
- M1 growth and, 108
- multivariate models and, 102–103, 109, 118–119, 121–128, 152
- Naïve, 143, 195–200
- non-accelerating inflation rate of unemployment (NAIRU) and, 110, 115, 125
- oil prices and, 3, 38, 85–87, 111, 295, 303, 417, 418, 422, 454n1
- Pagan comments and, 187–193
- personal consumption expenditures (PCEs) and, 102, 106–111, 114, 117–120, 123, 127–128, 132, 143–144
- Phillips curve forecasts and, 106–107
- Producer Price Index (PPI) and, 112
- prototypical models for, 104–108
- pseudo out-of-sample, 101–104, 108–115, 118, 120, 124–131, 178n1
- random walk benchmark and, 101, 105, 116–118, 121, 129, 178n5
- Reichlin comments and, 195–203
- Riksbank and, 434, 441
- root mean squared error (RMSE), 104, 111–112, 115–117, 120–128, 131–151, 178nn4,7, 179n10
- step function lag restriction and, 128–129
- successful, 144–152
- Taylor rules and, 201–202
- terminology for, 103–104
- time-variation and, 109–110
- Treasury Inflation-Protected Securities (TIPS) and, 94, 107

- triangle model and, 102, 106–107, 111–113, 129–152, 156–157, 161–162, 166–167, 171–172, 176–177, 191
- univariate models and, 101–102, 104, 106, 109, 116–121, 124–129, 132–141, 152–153, 158, 163, 168, 170nn8,10, 173
- unobserved components-stochastic volatility (UC-SV) model and, 101, 105, 116, 124, 131–141, 145–153, 157–158, 163, 168, 173, 187–192
- vector autoregression (VAR) and, 112, 124, 201–202
- vector error correction model (VECM) and, 112
- warning signs of 1990s and, 110–116
- Inflation Persistence Network (IPN), 48, 109, 302, 306, 308, 322, 333, 351
- Information costs, 23–25
- Interest rates, 80, 449
 - central banks and, 407–412, 417, 423, 424n14, 441
 - consumption reaction and, 47
 - Exchange Rate Mechanism and, 368–369
 - feedback rule and, 20
 - island economies and, 249
 - Israel and, 410
 - Lawson boom and, 368–369
 - long/short spread and, 203
 - microeconomic price data and, 367–368
 - Mishkin specification and, 107
 - rational expectations and, 252–254, 258–259, 262, 274
 - real, 38, 82
 - Rotemberg-Woodford model and, 19–20
 - as spending stimulus, 8
 - term structure and, 126, 198
 - vector autoregression (VAR) and, 201
- Intermittent Observation theory, 260
- International Monetary Fund (IMF), 409
- Internet, 210, 260
- Inventory-theoretic models, 355
- Ireland, 369
- IS curve, 10–11, 20, 38, 368
- Island economies, 249
- Israel
 - Consumer Price Index (CPI) and, 404, 406
 - dynamic stochastic general equilibrium (DSGE) models and, 410–412
 - economic demographics of, 403–404
 - exchange rate and, 404
 - inflation and, 404–412
 - interest rates and, 410
 - Lebanese War and, 404
 - Lucas critique and, 403, 407, 412
 - Taylor rule and, 409
- Italy, 316, 327n20, 374
- Iterated forecasts, 103, 105–108, 112, 128
- Japan, 365, 374
- Job Openings and Labor Turnover Survey (JOLTS), 45, 210, 221, 226, 240
- Kalman filter, 209, 219, 221
- Kalman predictor, 189
- Keynesianism
 - capacity point of, 247–248
 - hysteresis and, 395
 - inflation forecasts and, 39–42
 - microeconomic price data and, 292
 - New Keynesian Phillips curve and, 16–19, 24, 26, 31, 37–38 (*see also* New Keynesian Phillips curve)

- null hypothesis and, 310–312
 - policy making and, 248
 - rational expectations and, 11, 247–249
 - triangle model and, 11
 - UV-analysis and, 237
 - wages and, 247
- Kimball aggregator, 318
- Labor
 - Beveridge curve and, 44–46, 208–211
 - Bureau of Labor Statistics (BLS) and, 26–27
 - Careerbuilder.com and, 210
 - continuous-time labor market model and, 212
 - Dickens model and, 212–227, 231–243
 - dynamic stochastic general equilibrium (DSGE) models and, 20–22
 - equal opportunity laws and, 210
 - firing restrictions and, 366
 - help-wanted index and, 209–210
 - hysteresis and, 361–400 (*see also* Hysteresis)
 - inflation specification and, 17
 - Job Openings and Labor Turnover Survey (JOLTS) and, 45, 210, 221, 226, 240
 - Monster.com and, 210, 232
 - NAIRU and, 42–46 (*see also* Non-accelerating inflation rate of unemployment (NAIRU))
 - New Keynesian Phillips curve and, 16–19
 - output per hour, 386
 - reallocation issues and, 235
 - role of labor costs and, 57–58
 - skilled, 52
 - staggered contracts and, 12–16, 24, 37, 325n1
 - unemployment issues and, 3 (*see also* Unemployment)
 - unions and, 366
 - vacancy rates and, 45, 209–211, 214–216, 221, 226, 232–233, 240
- Lagged inflation
 - historical perspective on Phillips curve and, 9, 18–24, 30, 34, 40, 52, 54, 59n8, 61n24, 90–91
 - household surveys and, 90–91
 - hysteresis and, 361
 - inflation forecasts and, 106, 115, 178n3, 191
 - policy and, 283
 - price determination and, 324–343
 - time-varying NAIRU and, 227n1, 222
- LASSO algorithms, 123–124, 197–198
- Lawson boom, 368–369
- LM curve, 10–11, 81
- London School of Economics (LSE), 72
- Lowpass filters, 108, 144
- Lucas critique
 - central banks and, 403, 407, 412, 419, 423n6
 - Israel and, 403, 407, 412
 - rational expectations and, 248–251
- Lucas supply curve, 249
- Lucas surprise model, 79
- Macroeconomics, xi. *See also*
 - Keynesianism
 - dynamic stochastic general equilibrium (DSGE) models and, 20–23
 - future research for, 87–88
 - general equilibrium representation and, 19–21
 - inflation forecasts and, 39–42 (*see also* Inflation forecasts)

- integrating Phillips curve into new models and, 19–23
- learning from Phillips curve, 71–98
- microeconomic price data and, 48–51, 291–325
- natural rate of unemployment and, 51–53 (*see also* Unemployment)
- price inertia and, 12
- price rigidities and, 12–13, 15
- rational expectations and, 9–12, 272–277
- setting heterogeneity and, 291–292
- stabilization issues and, 281–282
- Mean squared error (MSE), 112, 130
- Median consumer price, 61n32, 292, 310
- Menu cost model
 - Golosov and, 49
 - Lucas and, 49
 - microeconomic price data and, 297, 300, 307–308, 311–313, 316, 320, 322, 324, 326n12, 327nn28,33, 351–357
 - Phillips curve and, 27, 49–50
 - rigidities and, 322, 324
 - sticky prices and, 354
- Mexico, 308
- Michigan Survey of Consumer Sentiment, 119
- Michigan Survey of Households, 272, 274
- Microeconomic price data, 48–51
 - alternative approaches and, 23–26
 - Bayesian unobservable index model and, 313–314
 - Bureau of Labor Statistics and, 295, 297–299, 301, 307–308, 310, 320, 351, 354
 - Calvo model and, 292–293, 311–312, 316–319, 325, 327n32, 328n24, 334–348, 354
 - causality and, 291
 - Consumer Price Index (CPI) and, 295, 342, 345
 - cross-country evidence and, 306–308
 - customer costs and, 308–310
 - data sources and, 291
 - dynamic indexation and, 338, 340, 342–344
 - dynamic stochastic general equilibrium (DSGE) models and, 293–294, 316–319, 321, 325, 334–341, 347
 - elasticity of substitution and, 311
 - Entry Level Items (ELIs) and, 295, 301–302
 - euro area and, 292–293, 302–306, 316–319
 - exogenous timing models and, 320–321
 - forced item substitutions and, 297–303, 310, 312, 326n11, 327n20
 - frequency of price changes and, 291–310, 335–337, 347
 - heterogeneity in, 291–303, 310, 312, 319, 322, 344, 351–352
 - imperfect information and, 313
 - impulse responses and, 292, 313–316
 - King comments on, 333–348
 - magnitude of adjustment and, 345–347
 - mapping into macroeconomic models, 312–319
 - markup deviation and, 312
 - menu cost and, 297, 300, 307–308, 311–313, 316, 320, 322, 324, 326n12, 327nn28,33, 351–357
 - Midrigan comments on, 351–357
 - New Keynesian Phillips curve and, 293
 - null hypothesis and, 310–312
 - rational inattention and, 323–324
 - real exchange rates and, 355
 - regularities in, 292

- representative-firm model and, 297
- rigidities and, 321–324, 351, 355–357
- rule-of-thumb behavior and, 319
- sales data and, 292, 297–306, 335–337, 352–353
- shocks and, 292–294, 297, 299–302, 311–324, 326nn2,11,12, 327nn28,32,33, 337, 353–356
- Smets-Wouter model and, 293–294, 316–320
- state-dependent models and, 320–321, 344–347
- sticky information and, 323–324
- survey evidence and, 308–310
- United States and, 292–302, 316–319
- vector autoregression (VAR) and, 316
- V-shape data and, 299–300
- Models. *See also* Specific model
 - ad hoc mechanisms and, 9–12
 - dummy variables and, 106, 208, 219–221, 227, 263n4
 - empirical evidence and, 26–29
 - endogenous persistence and, 59n10
 - microeconomic, 13–21, 26–29
 - need for structural, 53–54
 - newer macroeconomic, 19–23
 - price-setting behavior and, 26–29
 - rational expectations and, 9–12
 - stagflation and, 4
 - staggered contracts and, 12–13
- M1 growth, 108
- Monetary Policy Report* (Riksbank), 434
- Monetary Policy Update* (Riksbank), 434
- Money illusion, 250–251
- Monster.com, 210, 232
- Monte Carlo trials, 224–225
- Multivariate models, 40, 252
 - combination forecast, 123–124
 - dynamic factor, 121–123
 - inflation forecasts and, 102, 109, 118–128, 152
 - nonlinear, 124–125
 - price cross-section and, 127–128
 - structural term, 125–127
 - time-varying NAIRU and, 200, 203
- Naïve forecasts, 143, 195–200
- National Bureau of Economic Research, 196–197, 453
- National Income and Product Accounts, 386
- National Science Foundation, 152
- Natural rate hypothesis
 - Dickens model and, 233–234
 - historical perspective on Phillips curve and, 8–9, 37–38, 42–45, 51–55, 77–78, 81–85, 92
 - hysteresis and, 361–364, 378, 383, 396–400
 - NAIRU and, 77, 84, 231–236
 - New Keynesian Phillips curve and, 81
 - Taylor rule and, 85, 92
- Neoclassical Synthesis, 49
- Netherlands, 368–369
- New Keynesian Phillips curve, 16–19, 24, 26, 31
 - Calvo model and, 77–78
 - development of, 77
 - distance from capacity and, 250
 - European Central Bank and, 428–429
 - improved NAIRU measurement and, 52–56
 - inflation forecasts and, 106–107, 110, 119 (*see also* Inflation forecasts)
 - instability of, 250–251
 - Lucas critique and, 250–251
 - microeconomic price data and, 293, 317, 327n30
 - money illusion and, 250–251
 - natural rate of unemployment and, 81

- non-neutrality and, 46
- output gap and, 251
- policy and, 250–257, 427–432
- prices and, 250
- rational expectations and, 283
- real tightness and, 251
- shocks and, 254–255
- Sims model and, 252–257
- Solow and, 37–38
- theory of, 250
- unemployment/inflation
 - relationship and, 250
 - usefulness of, 250–257, 281
- New Neoclassical synthesis, 310–312
- New Zealand, 5, 72–73, 365, 371, 373–374, 438
- Non-accelerating inflation rate of unemployment (NAIRU), 36, 152
 - Beveridge curve and, 44–46, 208–227, 237–242
 - Blanchard comments and, 231–236
 - business cycle fluctuations and, 241–242
 - coefficient models and, 208
 - confidence intervals and, 43, 209
 - determining mechanisms of, 51–53
 - Dickens model and, 212–227
 - disinflation and, 366–367, 371–377, 379n3
 - dummy variables and, 208
 - estimation of, 42–46, 92, 207–228, 364–365, 381
 - euro area and, 363, 365
 - Federal Reserve and, 33
 - firing restrictions and, 366
 - help-wanted advertising and, 44–45
 - historical levels of, 207–208
 - hysteresis and, 52–53, 361–378, 381, 383, 395
 - improved measurement of, 42–46, 92
 - inflation forecasts and, 41, 115, 125
 - Kalman filter and, 209, 221
 - Monte Carlo trials and, 224–225
 - natural rate hypothesis and, 77, 84
 - policy and, 32–33
 - potential GDP estimation and, 207
 - role of, 9
 - sacrifice ratio and, 33
 - structural break models and, 208
 - time varying, 32, 110, 129–130, 207–228, 231–243
 - United States and, 363
 - unobserved components-stochastic volatility (UC-SV) model and, 188, 191
 - vacancy rates and, 44
- Non-synchronization phenomenon, 79
- Norges Bank, 438
- Norway, 373
- Null hypothesis, 310–312
- Oil prices, 3, 38, 85–87, 454n1
 - central banks and, 417–418, 422
 - inflation forecasts and, 111
 - OPEC shock and, xi
 - price determination and, 296, 298–303
- Optimization, 16–17
 - IS curve and, 20
 - micro-founded Phillips curve and, 19–21
- Organisation of Economic Co-operation and Development (OECD), 365, 368, 396
- Organization of the Petroleum Exporting Countries (OPEC), xi
- Output gap
 - ADL Phillips curve models and, 129
 - Bernanke on, 451–453
 - central banks and, 403, 412, 433, 441, 451–452
 - historical perspective on Phillips curve and, 19–20, 24, 29, 47, 54, 58, 60n14, 61n27, 74, 81
 - inflation forecasts and, 106, 112, 115, 118, 120, 129, 178n7, 192

- natural rate of unemployment and, 81
- New Keynesian Phillips curve and, 251
- policy issues and, 251, 274, 281–282
- rational expectations and, 274
- real-time policymaking and, 451–453
- Oxford Economic Papers*, 237
- Pass band filters, 108
- Personal consumption expenditures (PCE) models, 35
 - increased commodity prices and, 89
 - inflation forecasts and, 102, 106–111, 114, 117–120, 123, 127–128, 132, 143–144
 - performance of, 143
 - unobserved components-stochastic volatility (UC-SV) model and, 163–172
- Philadelphia Fed's Livingston Survey, 119
- Phillips curve
 - accelerationist, 8, 111, 361, 377, 396
 - ad hoc mechanisms and, 9–12
 - assessment of, 36–39
 - augmented expectations and, 79
 - backward-looking, 32, 115 (*see also* Backward-looking Phillips Curve)
 - Calvo and, 13–16, 23–24, 30–31
 - causality and, 37, 76–77
 - central banks and, 52–57, 415–426
 - coining of term, 73
 - decreasing relative predictability and, 196–201
 - deeper micro-foundations and, 17–19
 - delayed policy effect of, 75
 - disinflationary pressure and, 4
 - dynamic stochastic general equilibrium (DSGE) models and, 20–23, 27
 - empirical challenges of, 31–35
 - expectations-augmented, 74
 - Federal Reserve conference on, xi, xiii, 4, 36, 53
 - first-generation, 389
 - first publication of, 71–72
 - Fischer and, 12
 - forecast performance of, 144–152
 - forward-looking, 16, 18–24 (*see also* Forward-looking Phillips curve)
 - Friedman and, 6–11
 - Gray and, 12
 - hybrid, 18, 24
 - implementation issues of, 31–35
 - inflation forecasts and, 39–42, 101–152 (*see also* Inflation forecasts)
 - integration into newer macroeconomic models, 19–23
 - intellectual history of, 5–35, 247–249
 - IS curve and, 10–11, 20
 - labor costs and, 450–451
 - learning from, 71–98
 - LM curve and, 10–11
 - Lucas and, 9–11
 - macroeconomic model implications and, 48–51
 - microeconomic price data and, 355 (*see also* Microeconomic price data)
 - micro-founded alternative explanations and, 23–26
 - NAIRU and, 207–228 (*see also* Non-accelerating inflation rate of unemployment (NAIRU))
 - natural rate of unemployment and, 51–53
 - New Keynesian, 16 (*see also* New Keynesian Phillips curve)
 - Phelps and, 6–11
 - policy and, 6–9, 46–48
 - price-setting behavior and, 26–29 (*see also* Prices)

- rational expectations and, 9–12, 267–269 (*see also* Rational expectations)
- real time data and, 57
- relative stability of, 34
- reversed-direction, 40, 46, 249
- Riksbank and, 433–444
- role of labor costs and, 57–58
- Rotemberg and, 13–16
- Sargent-Wallace and, 9–11
- second-generation, 389
- sticky-information context and, 24
- Taylor and, 12–13
- time-varying inflation targets and, 29–31
- unemployment and, 6 (*see also* Unemployment)
- Poland, 308
- Policy, 23–26
 - Bernanke and, 447–454
 - Calvo model and, 79
 - central banks and, 53–57, 403–413, 417–432
 - Dickens model and, 225
 - Exchange Rate Mechanism and, 368–369
 - forward-looking analysis and, 56
 - hysteresis and, 52–53, 367–368, 378, 383
 - inflation forecasts and, 46–48, 102 (*see also* Inflation forecasts)
 - Keynesianism and, 248
 - labor costs and, 450–451
 - lower unemployment, 7–9
 - microeconomic price data and, 355 (*see also* Microeconomic price data)
 - monetary expansion and, 8
 - New Keynesian Phillips curve and, 250–257
 - non-accelerating inflation rate of unemployment (NAIRU) and, 32–33
 - output gap and, 19–20 (*see also* Output gap)
 - rational expectations and, 9–12, 247–249 (*see also* Rational expectations)
 - rational inattention and, 261–263
 - real-time, 57, 451–453
 - recessions and, 367–368, 378
 - Riksbank and, 433–444
 - Sims model and, 252–257
 - stagflation and, xii–xiii, 4
 - staggered contracts and, 12–13
 - Taylor rule and, 38, 81, 85, 92–93, 201, 272, 274, 409
 - trade-offs and, 6–9, 210–211
 - triangle model and, 11 (*see also* Triangle policy)
 - welfare costs and, 48, 51, 263, 325, 400, 415–416
- Portugal, 369, 373
- Prices. *See also* Consumer Price Index (CPI)
 - ad hoc mechanisms and, 9–12
 - asymmetry issues and, 92–93
 - Bayesian unobservable index model and, 313–314
 - Bernanke on, 447–451
 - Calvo model and, 15, 17, 28–29, 46, 60n12, 62n34, 338, 340
 - central bank influence and, 56–57, 415–423
 - commodity shocks and, 417–423 (*see also* Shocks)
 - cross-section forecasts and, 127–128
 - customer costs and, 308–310
 - data quality and, 452–453
 - demand conditions and, 9
 - Dickens model and, 212–227
 - disinflationary pressure and, 4
 - duration of median, 292
 - dynamic indexation and, 338, 340, 342–344
 - dynamic stochastic general equilibrium (DSGE) models and, 20–23, 27
 - Federal Reserve and, 415–416

- firm optimization problem and, 16
- food, 3, 295, 298, 300–301, 418, 452
- frequency of price changes and, 291, 347
- Fuhrer-Moore model and, 17–18
- global trends and, 420–422
- Harmonized Index of Consumer Prices and, 127, 316
- hazard rates and, 310–311
- imperfect information and, 269–272
- impulse response and, 292, 313–316
- inertia and, 10–12, 47, 60n18, 62n36, 79, 81, 240, 262, 351, 354, 356
- increased commodity, 85–89
- labor costs and, 450–451
- macroeconomic model implications and, 48–51
- magnitude of adjustment and, 345–347
- markup deviation and, 312
- menu cost model and, 27, 49–50, 297, 300, 307–308, 311–313, 316, 320, 322, 324, 351–357
- microeconomic price data and, 26–29, 291–325 (*see also* Microeconomic price data)
- narrow measurement focus and, 93–94
- natural rates and, 8
- nature of price dynamics and, 335–337
- New Keynesian Phillips curve and, 16–19, 250
- null hypothesis and, 310–312
- oil, xi, 3, 38, 85–87, 111, 295, 298–303, 417–418, 422, 454n1
- personal consumption expenditures (PCEs) and, 35, 89, 102, 106–111, 114, 117–120, 123, 127–128, 132, 143–144
- Phillips paper and, 5–6
- Producer Price Index (PPI) and, 112, 301, 326n14
- rational expectations and, 9–11, 260–261 (*see also* Rational expectations)
- rational inattention and, 314 (*see also* Rational inattention)
- real time data and, 57–58
- rigidities and, 12–15, 40, 62n36, 268, 271–272, 321–324, 351, 355–357
- role of labor costs and, 57–58
- sector response and, 48–49, 61n33, 313–316
- setting heterogeneity and, 291–292
- stagflation and, xii–xiii, 3–4, 284, 286
- state-dependent models and, 320–321, 344–347
- sticky, 18, 78, 292, 315, 323–324, 334–341, 348, 353–356
- timing of adjustment in, 338–339, 344–347
- triangle model and, 11 (*see also* Triangle model)
- utility maximization and, 15–16
- vector autoregression (VAR) and, 316
- V-shape data and, 299–300
- Principal component (PC) regression, 197–198
- Producer Price Index (PPI), 112, 301, 326n14
- Pseudo out-of-sample forecasts
 - inflation and, 101–104, 108–115, 118, 120, 124–131, 178n1
 - rolling/recursive estimation and, 103–104
 - strict interpretation of, 178n2
- Quarterly Review* (Federal Reserve Bank of Richmond), 81

- Ramses model, 436–438
- Random walk benchmark, 178n5
 Atkeson-Ohanian model and, 101, 105 (*see also* Atkeson-Ohanian model)
 four quarter, 101, 105, 116–117, 118
 inflation forecasts and, 101, 105, 116–118, 121, 129, 178n5
 minor role of, 178n5
 performance of, 101, 121
- Rational expectations, xiii, 16, 95–96, 263n8
 adaptive expectations and, 283
 alternative methods and, 284, 286–287
 behavioral analysis and, 259–261
 Bernanke and, 282, 453–454
 commodity price shocks and, 417–420
 communism of, 262, 281
 cross-equation restrictions and, 248
 deviation from, 267–269, 281, 283–287
 disinflation costs and, 269
 equilibrium and, 78, 268
 Friedman model and, 89–90
 GDP data and, 260
 household surveys and, 90–91
 imperfect information and, 269–272
 inertia models and, 259–261
 inflation forecasts and, 39–40
 Intermittent Observation theory and, 260
 Kiley comments and, 267–270
 learning and, 58, 453
 Lucas and, 9–12, 78, 248–249, 267–269
 monetary neutrality and, 428
 money illusion and, 250–251
 Muth on, 277n1
 New Keynesian Phillips curve and, 250–257
 Orphanides comments and, 281–288
 perfect information and, 61n29, 283
 policy and, 247–249, 259–263, 272–277, 281–287
 prices and, 260–261
 psychology and, 259–260
 research effects of, 268–269
 reversed-direction, 46, 248–249
 rigidities and, 40, 267–268
 shocks and, 282
 sticky-price models and, 18, 257
 Taylor rule and, 272–274
 uncertainty models and, 259–261
 white noise deviations and, 81
 without Phillips curve, 257–259
- Rational inattention, 25–26, 47, 49, 56, 316
 foundation of, 270–271
 free information and, 260
 information use and, 260
 marginal value of information and, 309, 325
 microeconomic prices and, 260–261, 323–324
 optimal pricing and, 314
 policy implications and, 261–263
 research support and, 275
 sticky information and, 323–325
 uncertainty and, 299
- Real estate, xi
- Recessions, 232, 263n6, 403
 central banks and, xii, 367
 creeping inflation and, 73
 Federal Reserve policy and, 367, 378
 historical perspective on Phillips curve and, 3, 24, 41, 51, 61n29, 73, 80
 hysteresis and, 367–368
 inflation forecasts and, 102, 114, 143, 186–188, 192
 OPEC and, xi
 policy responses to, 367–368

- real estate and, xi
- Wall Street and, xi
- “Relations Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957” (Phillips), 5
- Relative predictability, 196–201
- Reserve Bank of New Zealand (RBNZ), 438
- Resource slack, 4
- Reuters survey, 272, 274
- Ridge regression, 197–200
- Rigidities, 12–13, 15
 - aggregate inertia and, 62n36
 - imperfect information and, 271–272
 - menu cost and, 324
 - microeconomic price data and, 321–324, 351, 355–357
 - rational expectations and, 40, 268
- Riksbank, xiii, 55–56, 63n42
 - Bayesian models and, 434
 - Consumer Price Index (CPI) and, 433
 - decision-making process of, 433–434
 - Executive Board and, 433–434
 - inflation forecasts and, 434, 441
 - instrument-rate paths and, 438, 441
 - policy and, 433–444
 - Ramses model and, 436–438
 - repo rate and, 433
 - vector autoregression (VAR) and, 434
- Root mean squared error (RMSE) forecasts
 - consumer price index (CPI) model and, 153–162
 - equations for, 104
 - GDP deflator and, 173–177
 - inflation and, 104, 111–112, 115–117, 120–128, 131–151, 153, 178n4,7, 179n10
 - performance of, 132, 143–144, 146
 - personal consumption expenditures (PCE) model and, 163–172
 - rolling, 9, 103–104, 109, 111, 114–118, 125–126, 129–142, 145–151, 157, 179n10
- Rule-of-thumb behavior, 319
- Shocks, 423n2
 - aggregate, 25, 27–28, 49–51, 292, 314–315, 320, 322–323
 - ARMA structure and, 21
 - Beveridge curve and, 238–239
 - central banks and, 417–420
 - DSGE models and, 22–23
 - historical volatility and, 199–201
 - hysteresis and, 363, 365, 369, 381n3, 397, 400
 - idiosyncratic, 25, 28, 31, 51, 292–293, 311, 320, 353, 356
 - imperfect information and, 313
 - inertia and, 22–26
 - inflation forecasts and, 40
 - microeconomic price data and, 292–302, 311–324, 326nn2,11,12, 327nn28,32,33, 337, 353–356
 - NAIRU and, 43, 129, 208 (*see also* Non-accelerating inflation rate of unemployment (NAIRU))
 - New Keynesian Phillips curve and, 19, 254–255
 - nominal, 50–51, 238–242, 297, 300, 311–313, 320–324, 326n12, 354–355
 - OPEC oil and, xi
 - Paciello model and, 25–26
 - price-setting behavior and, 26–29
 - propagation and, 201–202
 - rational expectations and, 10, 282
 - sector-specific, 292, 313–316, 322, 327n28
 - Sims model and, 254
 - sticky prices and, 292

- structural models for, 54
- total factor productivity suppression and, xii
- UC-SV model and, 188 (*see also* Unobserved components-stochastic volatility (UC-SV) model)
- vector autoregression (VAR) and, 21
- Simple state-dependent pricing (SDP) model, 344–347
- Spain, 327n20, 368, 374
- Speed limit effect, 5
- Stagflation, xii–xiii, 3–4, 284, 286
- Staggered contracts, 12–13, 24, 37, 325n1
- “Staggered Prices in a Utility-Maximizing Framework” (Calvo), 15–16
- State-dependent models, 27–28, 320–321, 344–347, 377
- Step function lag restriction, 128–129
- Sticky prices models, 15, 62n34, 271, 292, 314, 348
 - dynamic stochastic general equilibrium (DSGE) models and, 50, 334–340
 - exchange rates and, 50
 - information costs and, 23–24, 47, 51
 - lack of persistence in, 78
 - marginal costs and, 356
 - menu cost and, 354
 - non-neutralities and, 334
 - rational expectations and, 18
 - rational inattention and, 323–325
 - real exchange rate and, 355
 - sector indices and, 48–49
 - transmission effects and, 351–352
- Structural models
 - central banks and, 419–420
 - inflation forecasts and, 125–127
 - non-accelerating inflation rate of unemployment (NAIRU) and, 208
- Subprime mortgages, xi
- Survey of Professional Forecasters, 119, 272
- Sweden, 55–56, 63n42, 373–374. *See also* Riksbank
- Switzerland, 374
- Taxes, 207
 - imperfect indexation of, 416
 - personal income, 3
 - rebates and, 3
 - sales, 379n5
- Taylor rule, xii
 - consumer price inflation and, 274
 - conventional policy and, 92, 409
 - deviations from, 272
 - inflation forecasts and, 201–202
 - IS curve and, 38
 - Israel and, 409
 - LM curve and, 81
 - microeconomic price data and, 311
 - natural rate hypothesis and, 85, 92
 - rational expectations and, 272–274
 - real activity and, 201
- Technology, 52, 81
 - Careerbuilders.com and, 210
 - computer prices and, 424n15
 - menu costs and, 354, 356
 - Monster.com and, 210
 - shocks and, 25–26
 - worker placement and, 210, 232
- Three-equation model, 81–82
- Transportation costs. *See* Oil prices
- Treasury Inflation-Protected Securities (TIPS), 94, 107, 128, 272
- Triangle model, 11, 59n8, 61n24
 - equation of, 106
 - inflation forecasts and, 102, 106–107, 111–113, 129–152, 156–157, 161–162, 166–167, 171–172, 176–177, 191
 - time-varying NAIRU models and, 129–130

- unobserved components-stochastic volatility (UC-SV) model and, 191
- Uncertainty models, 259–261
- “Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective” (Federal Reserve Bank of Boston conference), 4, 36, 53
- Unemployment, xi–xiii, 3
 - Beveridge curve and, 44–46, 208–227, 231–243
 - business cycle fluctuations and, 241–242
 - determining rate of, 51–53
 - Dickens model and, 212–227
 - firing restrictions and, 366
 - hysteresis and, 52–53, 361–400 (*see also* Hysteresis)
 - inflation forecasts and, 102
 - monetary expansion and, 8
 - NAIRU and, 207–228 (*see also* Non-accelerating inflation rate of unemployment (NAIRU))
 - natural rate of, 84–85 (*see also* Natural rate hypothesis)
 - New Keynesian Phillips curve and, 16–19, 250
 - Phillips paper and, 5–6
 - policy trade-offs and, 6–9
 - rational expectations and, 9–11
 - speed limit effect and, 5
 - staggered contracts and, 12–13
 - Taylor rule and, xii (*see also* Taylor rule)
 - UV-analysis and, 237
- Unemployment insurance, 368
- United Kingdom, 5–6, 240, 262
 - Bank of England and, 274–275, 410
 - Chancellor of the Exchequer and, 369
 - Exchange Rate Mechanism and, 368–369
 - hysteresis and, 368–369, 373–374
 - Lawson boom and, 368–369
 - London School of Economics (LSE), 72
 - policy trade-offs and, 7
 - Thatcher and, 249
- United States
 - Consumer Price Index (CPI) and, 345 (*see also* Consumer Price Index (CPI))
 - forced item substitution and, 297–302
 - frequency of price changes and, 292, 294–302, 310, 326n2
 - global trends and, 421
 - Great Depression and, 384, 386
 - hysteresis and, 262, 363–364, 373, 399
 - inflation forecast data and, 108
 - median consumer price duration in, 292
 - microeconomic price data and, 292–302, 316–319
 - non-accelerating inflation rate of unemployment (NAIRU) (*see also* NAIRU) and, 363
 - persistent inflation and, 29–30
 - policy and, 6–9, 258–259 (*see also* Policy)
 - primary surpluses of, 259
 - reallocation issues and, 235
 - recession and, 3–4
 - World War II and, 384
- Univariate models
 - ARIMA, 104, 178n4, 189
 - Bayes Information Criterion and, 128–129
 - inflation forecasts and, 101–102, 104, 106, 109, 116–121, 124–129, 132–141, 152–153, 158, 163, 168, 170nn8,10, 173
 - NAIRU estimation and, 365
 - performance of, 132
 - step function lag restriction and, 128–129
 - success rate of, 102–103

- University of Minnesota, 92
- Unobserved components-stochastic volatility (UC-SV) model, 40–41
 - arbitrary persistence value and, 188
 - CPI models and, 153–162
 - exponentially weighted moving average (EWMA) formula and, 189–191
 - inflation forecasts and, 101, 105, 116, 124, 131–141, 145–153, 157–158, 163, 168, 173, 187–192
 - Kalman predictor and, 189
 - non-accelerating inflation rate of unemployment (NAIRU) and, 188, 191
 - one parameter of, 105–106
 - Pagan comments and, 187–192
 - performance of, 132, 143–144, 146, 189
 - personal consumption expenditures (PCE) model and, 163–172
 - smoothness control and, 105–106
 - time varying, 130
 - triangle model and, 191
 - unit root imposition of, 178n6
- UV-analysis, 237
- Variable selection algorithms, 197–198
- Vector autoregression (VAR), 20–21
 - Atkeson-Ohanian model and, 118–119
 - bivariate, 118–119
 - direct forecasts and, 103–107, 117–118, 122, 127–129, 179n8
 - dynamic stochastic general equilibrium (DSGE) models and, 316
 - inflation forecasts and, 42, 46, 112, 124, 201
 - interest rates and, 201
 - microeconomic price data and, 316
 - Riksbank and, 434
 - time-varying inflation targets and, 29–30
 - trivariate, 118–119
- Vector error correction model (VECM), 112
- V-shape price data, 299–300
- Wages, xii–xiii
 - cost-of-living effect and, 5
 - Dickens model and, 212–227, 231–243
 - dynamic stochastic general equilibrium (DSGE) models and, 20–22
 - Friedman on, 7–8
 - hysteresis and, 361–378
 - insider-outsider bargaining theory and, 364, 376
 - monetary expansion and, 8
 - New Keynesian Phillips curve and, 16–19, 247
 - Phillips paper and, 5–6
 - policy trade-offs and, 6–9
 - price inertia and, 12
 - rational expectations and, 9–11
 - relative to price inflation, 7–8
 - speed limit effect and, 5
 - staggered contracts and, 12–13
 - stickiness and, 22
 - utility maximization and, 15–16
- Wall Street, xi
- Welfare costs, 48, 51, 263, 325, 400, 415–416
- “Whip Inflation Now” (WIN) campaign, 75
- World War II, 384