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 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
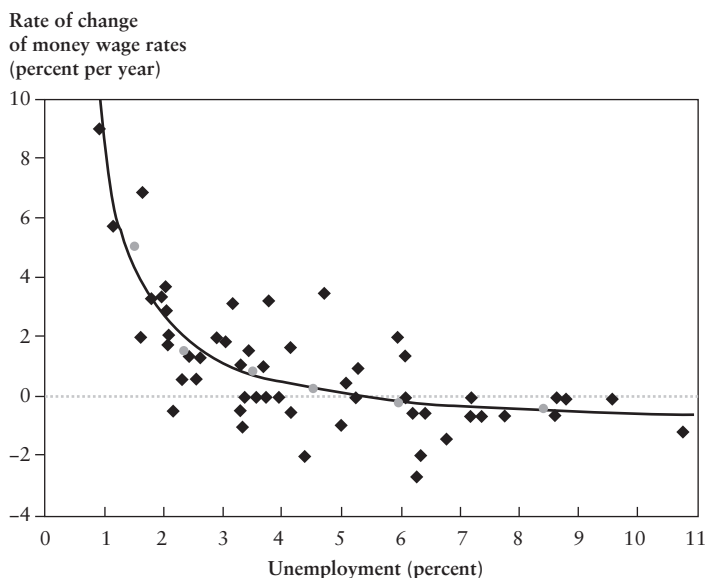
**Rate of change
of money wage rates
(percent per year)**

Unemployment (percent)

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 pol-
icy 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-unem-
ployment 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) $\pi_t = \pi_t^e - aU_t = \pi_{t-1} - aU_t$

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

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:

$$p_t = a p_t^e + b D_t.$$

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

$$p_t^e = c p_{t-1},$$

so that the resulting price equation becomes

$$p_t = ac p_{t-1} + b D_t.$$

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

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

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" model[8] (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)
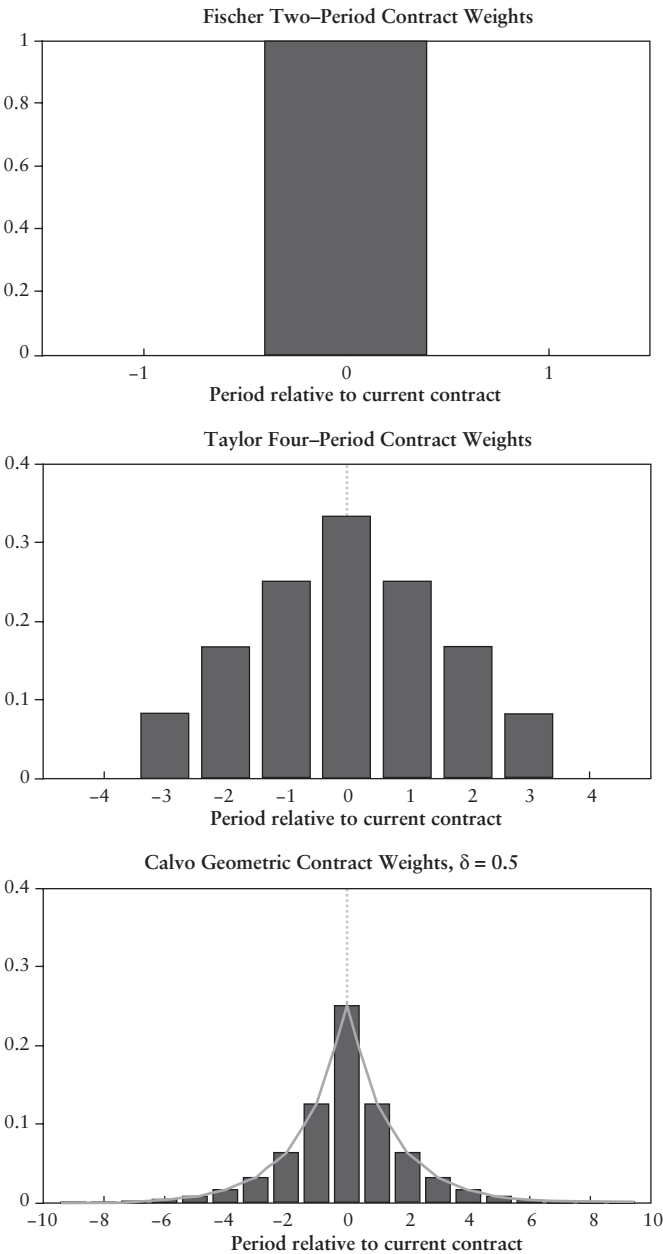
Figure 1.2
The Influence of Neighboring Wage Contracts on the Current Contract Wage
Source: Authors' calculations.

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 micro-economic foundations for such models began in earnest.

Three papers that arrived on the scene around the same time pro-vided partial answers to this question regarding the existence of rigid (or "sticky") wages and prices. The two authors, Calvo (1983) and Rotem-berg (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 tractabil-ity 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 $\frac{1}{\delta}$. Equiva-lently, this probability implies a geometric distribution of price contract lengths, with shorter durations most likely and longer durations increas-ingly unlikely. As in Taylor (1980), firms set their contract price (when tapped by Calvo's price-change signal) in reference to overlapping con-tracts, 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:

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

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

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

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 monopolisticly 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*,

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

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,

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

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

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. Theirs 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.

Microeconometric 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. Bils 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 Bils 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 infla-tion. Consequently, Phillips curve specifications that try to match a per-sistent 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 pre-dictions for aggregate shocks that do not square well with most of the extant empirical evidence. This conundrum has found a potential resolu-tion 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 idio-syncratic 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 mon-etary 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 dis-cipline 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 infla-tion. 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

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

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

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't[27]) 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:

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

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

$$\ln \sigma^2_{\varepsilon,t} = \ln \sigma^2_{\varepsilon,t-1} + \nu_{\varepsilon,t,}$$

$$\ln \sigma^2_{u,t} = \ln \sigma^2_{u,t-1} + \nu_{u,t,}$$

where $\zeta_t = (\zeta_{\varepsilon,t}, \zeta_{u,t})$ is independent and identically distributed (i.i.d.)N(0, I_2), $\nu_t = (\nu_{\varepsilon,t}, \nu_{u,t})$ is i.i.d., N(0, γI_2), ζ_t and ν_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 Phillps 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):

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

$$\beta_s = \int_0^\infty f_r f_{r+s} dr,$$

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

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 Eliasz (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 Eliasz. 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 Lünnemann, Fabio Rumler, 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.