Economic Quarterly—Volume 94, Number 4—Fall 2008—Pages 301–309

Introduction to the New Keynesian Phillips Curve

Andreas Hornstein

In most industrialized economies inflation tends to be pro-cyclical; that is, inflation is high during times of high economic activity. When economic activity is measured by the unemployment rate this statistical relationship is known as the Phillips curve. The Phillips curve is sometimes viewed as a menu for monetary policymakers, that is, they can choose between high inflation and low unemployment or low inflation and high unemployment. But this interpretation of the Phillips curve assumes that the relationship between unemployment and inflation is structural and will not break down once a policymaker attempts to exploit the perceived tradeoff. After the high inflation episodes experienced by many economies in the 1970s, this structural interpretation of the Phillips curve was discredited. Yet, after a period of low inflation in the 1980s and early 1990s, economists have again worked on a structural interpretation of the Phillips curve. This New Keynesian Phillips curve (NKPC) assumes the presence of nominal price rigidities. In this special issue of the *Economic Quarterly*, we publish four surveys on the history of the Phillips curve, the structural estimation of the New Keynesian Phillips curve, and the policy implications of the nominal rigidities underlying the New Keynesian Phillips curve.

The Phillips Curve and U.S. Economic Policy

Robert King surveys the evolution of the Phillips curve itself and its usage in U.S. economic policymaking from the 1960s to the mid-1990s. He first describes how, in the 1960s, the Phillips curve became an integral part of U.S. macroeconomic policy in its pursuit of low unemployment rates. A stylized version of the Phillips curve that emerges from this period relates current

The views expressed do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

inflation, π, to the current unemployment rate, u, and lagged inflation,

$$\pi_t = \sum_{i \geq 1} \gamma_i \pi_{t-i} - \beta u_t.$$

Similar to other elements of the then-standard Keynesian IS-LM macromodel, economists would tell stories that motivated the Phillips curve but the Phillips curve was not derived from an explicit theory. Furthermore, the estimated parameters were taken as structural, in particular as invariant to policy interventions. In the late 1960s, Phelps (1968) and Friedman (1968) interpreted the Phillips curve as arising from search and information frictions in labor markets, and they argued that the relation between a real variable such as unemployment and nominal inflation was based on misperceptions about inflation on the part of the public. Phelps proposed an expectations-augmented Phillips curve,

$$\pi_t - \rho \pi_t^e = -\beta u_t,$$

where π^e denotes expected inflation. If, as Phelps and Friedman argued, $\rho = 1$, then a tradeoff between inflation and unemployment exists only to the extent that actual inflation deviates from expected inflation. At the time, inflation expectations were modeled as adaptive, that is, a geometric distributed lag of past actual inflation. In this case, for a constant actual inflation rate the expected inflation rate would eventually converge to the actual inflation rate and the unemployment rate would settle down at its natural rate. Thus, there is no long-run tradeoff between inflation and unemployment. Although Phelps and Friedman's argument originally represented a minority view in the profession, the argument became more widely accepted in the 1970s after periods of high inflation and unemployment.

Accounting for the instability of the Phillips curve in the 1970s had lasting effects on the way macroeconomic analysis was done and continues to be done today. First, since expectations play a crucial role in the expectations-augmented Phillips curve, it seemed necessary not to resort to some arbitrary assumption on the expectations mechanism. For this purpose, macroeconomists started to assume that expectations are rational. By this we mean that expectations are such that they do not lead to systematic mistakes given the available information. Sargent and Wallace (1975) used the idea of rational expectations in an otherwise standard IS-LM macromodel with an expectations-augmented Phillips curve to argue that systematic monetary policy actions do not systematically affect unemployment or output. Second, macroeconomists not only started to work with model-consistent expectations in otherwise ad hoc models, but they started to study the optimal choices of economic agents in explicitly specified environments agents; that is, they started to study macroeconomic questions using the tools of general equilibrium analysis. The seminal work was Lucas' (1972) formal analysis of the

Phelps-Friedman Phillips curve in an environment where agents had difficulty sorting out their own relative price shocks from aggregate price level shocks.

King describes how, at the end of the 1970s after years of persistently high inflation and high unemployment, monetary policymakers moved to lower the inflation rate. At that time, the debate centered on the perceived cost (in terms of elevated unemployment) associated with a reduction of the inflation rate. On the one hand, proponents of the more standard Phillips curve argued that these costs would be substantial. On the other hand, proponents of a rational expectations-augmented Phillips curve argued that the costs could be quite low, especially if the low inflation policy was credible to the public. In the end, the Federal Reserve under Paul Volcker reduced inflation over a relatively short time period at some cost, but not as high a cost as predicted by standard Phillips curves. For the remainder of the 1980s and the early 1990s, the Federal Reserve under Alan Greenspan further lowered average inflation and, in the process, strengthened its credibility for continued low inflation policies. King ends his survey in the mid-1990s when the Federal Reserve Board's monetary policy model incorporated an expectations-augmented Phillips curve with elements of rational expectations, and the Federal Open Market Committee debated the desirability of a target for low long-run inflation and what that target should be.

The New Keynesian Phillips Curve

At the time that U.S. inflation started to decline in the 1980s there was a resurgence of interest in business cycle analysis. Continuing the general equilibrium program in macroeconomics started with Lucas (1972), real business cycle analysis developed quantitative models of the aggregate economy based on the stochastic neoclassical growth model, e.g., Kydland and Prescott (1982) or Long and Plosser (1983). Using simulation studies, one could show that these models were able to mimic the U.S. business cycle in terms of the statistical properties of the time series of a limited number of aggregate variables (output, consumption, investment, and employment). As the name indicates, real business cycle theory addressed the behavior of quantities and relative prices over the business cycle, implicitly assuming that money is neutral. Working on the assumption that money is not neutral, economists in the mid-1990s then started to introduce nominal price rigidities into these models, now also known as Dynamic Stochastic General Equilibrium (DSGE) models. From this research program emerged the New Keynesian Phillips curve that relates actual and expected inflation not to the unemployment rate but to a measure of aggregate marginal cost. The second and third paper in this issue discuss the estimation of the structural parameters of the NKPC.

Once one assumes that nominal prices do not continuously adjust to clear markets, one has to decide how these prices are set in the first place.

Almost all of the work on nominal price rigidities has answered this question using the framework of monopolistic competition, which assumes that the product whose price has to be determined is produced by a profit-maximizing monopoly. There may be imperfect substitutes for the monopolist's product; that is, the demand for the product depends not only on its own price but also on the prices of the substitutes. When the monopolist decides on his own price he will, however, take these other prices as given, hence the term monopolistic competition. A monopolist that can continuously adjust his nominal price will set the price to equate contemporaneous marginal revenue and marginal cost and the price will be a markup over marginal cost. Compare this with flexible prices in perfectly competitive markets where the price and marginal cost are equated. If nominal prices cannot be continuously readjusted, then the monopolist will choose the current nominal price such that he equates the expected present value of marginal revenue and marginal cost over the time that the price remains fixed.

The model of an individual monopolistically competitive producer is then typically embedded into a general equilibrium model with a large number of these producers, e.g., Blanchard and Kiyotaki (1987). These producers are identical except for the time when they can adjust their nominal price. Various mechanisms for price adjustment have been proposed; most assume that the opportunity for price adjustment is exogenously given. One popular modeling technique is a Calvo-type price adjustment where, each period, a firm gets to adjust its price with some probability that is fixed over time. Using Calvo-type price adjustment, Woodford (2003) shows that the aggregation of the linearized optimal price adjustment rules for the individual firms yields an expression in current and expected future inflation and a measure of aggregate marginal cost, mc,

$$\pi_t = \gamma_f E_t \pi_{t+1} + \lambda mc_t + \xi_t.$$

This is the structural NKPC where γ_f and λ are functions of structural parameters, including the probability of price adjustment, α, and ξ_t is a random variable. The random disturbance is often interpreted as an exogenous shock to the firms' markup. Solving this difference equation forward, one can see that current and expected future marginal cost are driving today's inflation.

For most measures of inflation and what could be considered reasonable measures of marginal cost, inflation tends to be more persistent than marginal cost. Since marginal cost "drives" inflation in the basic NKPC, this makes it hard for the model to match the data. Economists have, therefore, modified the basic NKPC by introducing "rule of thumb" price adjusters or firms that simply index their price to the aggregate inflation rate, e.g., Galí and Gertler (1999). These assumptions lead to the inclusion of lagged inflation,

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \pi_{t+1} + \lambda mc_t + \xi_t,$$

and, therefore, make the NKPC a hybrid of the basic NKPC and more standard Phillips curves. The coefficients γ_b, γ_f, and λ are again functions of structural parameters. The ability of monetary policy to control inflation with a NKPC depends on the relative magnitudes of these coefficients. Loosely speaking, monetary policy affects inflation through its effects on marginal cost. Thus, the smaller the coefficient on marginal cost, the less impact monetary policy will have on inflation. In the extreme case when $\lambda = 0$, inflation evolves independently of monetary policy and whatever else happens in the rest of economy. How "costly" it is to reduce inflation depends on the relative magnitude of the coefficients on past and future inflation, γ_b and γ_f. If the coefficient on lagged inflation is large, then inflation is mostly driven by its own past and policy actions might affect inflation only with a long time lag. In order to evaluate the effectiveness of monetary policy actions we, therefore, need estimates of these parameters.

Single-Equation Estimation of the NKPC

In the second paper of this issue, James Nason and Gregor Smith survey the estimation of the parameters of the NKPC using only the NKPC itself. Single-equation estimation of the NKPC parameters is appealing because it does not require any assumption on how the rest of the economy should be specified. Yet standard ordinary least squares estimation of the NKPC is not applicable since expected inflation in the NKPC is an endogenous variable that is correlated with the error term of the estimation equation. Consistent parameter estimates can still be obtained through the use of the General Method of Moments (GMM) technique, which in turn requires instrumental variables that are correlated with expected inflation but uncorrelated with the other variables in the NKPC.

Nason and Smith report that, in general, estimated parameters for the hybrid NKPC are consistent with prior restrictions. For example, estimated price adjustment probabilities are between zero and one. They also find that the coefficient on expected future inflation tends to be larger than the coefficient on lagged inflation. This suggests that monetary policy can affect inflation in the short term. Nason and Smith also discuss the finding that the estimated coefficient on marginal cost tends to be small and barely significant. This is bad news for the NKPC as a model of inflation and for monetary policy.

The ambiguous evidence on the marginal cost coefficient may be related to weak identification through weak instrumental variables in the GMM estimation. Instrumental variables are essentially used to forecast expected inflation independent of the other variables in the NKPC. For an instrumental variable to serve its purpose it has to be correlated with expected future inflation and it should not be correlated with marginal cost and current and lagged inflation. But as Nason and Smith point out, past empirical work on inflation has shown that lagged inflation tends to be a good forecast of future inflation and it is difficult to improve on that forecast. This suggests that the instrumental variables in the GMM procedure are quite weak. Nason and Smith then show that after one takes into account that we have weak instruments, the evidence in favor of the NKPC is weakened or the NKPC is rejected outright.

System Estimation of the NKPC

In the third paper of this issue, Frank Schorfheide surveys system methods to estimate the parameters of the NKPC. For this approach one specifies a more or less complete model of the aggregate economy, a DSGE model, and then identifies the structural parameters from the restrictions that the equilibrium process imposes on the moments of a set of observable variables.

Using a simple example, Schorfheide interprets the various identification schemes used in the literature. He explains why it may not be possible to obtain consistent parameter estimates using single-equation methods. System methods on the other hand can obtain consistent parameter estimates through the imposition of prior constraints on elements of the DSGE model other than the NKPC. Essentially these prior restrictions allow one to identify exogenous shocks that may serve as instruments for the NKPC. As an example, Schorfheide points to the procedure of identifying monetary policy shocks from the restriction that the public cannot respond to contemporaneous monetary policy shocks. Schorfheide also suggests that it may not be possible to identify the coefficient on lagged inflation in the NKPC if one allows for serially correlated markup shocks. Indeed, single-equation estimates of the NKPC identify γ_b through the implicit prior restriction that the markup shock is i.i.d. This lack of identification affects the evaluation of policy effectiveness if it also implies that the coefficient on future inflation is not identified.

Schorfheide then surveys papers that estimate the NKPC as part of a more complete DSGE model. Most of this empirical work uses data on output, inflation, and a nominal interest rate. Marginal cost in the NKPC is then treated as a latent variable that is constructed from the observable variables and the equilibrium relationships implied by the DSGE model. But some empirical work also includes measures of marginal cost in the set of observable variables. Schorfheide observes that the range for the estimated coefficients on marginal cost in the NKPC is much larger when marginal cost is a latent variable. The range of estimated NKPC coefficients on marginal cost becomes much closer to that obtained from single-equation estimations once observations on marginal cost are included. Thus, with marginal cost as a latent variable, features of the DSGE model that are different from the NKPC can become much more important for the determination of the NKPC marginal cost coefficient. As is apparent from the work of Krause, López-Salido, and Lubik (2008), the implied process for the latent marginal cost variable is then very different from the process of various measures of marginal cost used in the literature.

In general, the literature review suggests that there is no consensus on the magnitude and role of nominal rigidities in the estimated price-setting process. Furthermore, introducing additional nominal rigidities in the wage-setting process affects the estimates for nominal rigidities in the price-setting process, that is, the NKPC. It also appears as if the relative role of nominal price and wage rigidities is not identified from the data.

Policy Implications of Nominal Price Rigidities

In the final paper of this issue, Stephanie Schmitt-Grohé and Martín Uribe discuss the implications of nominal price rigidities for optimal monetary policy. They first ask how the presence of nominal price rigidities affects the design of optimal policy when fiscal and monetary policy are jointly determined. They then go on to study if simple policy rules such as the Taylor rule can get the economy close to the optimal policy outcome. They find that with small amounts of nominal price rigidities, optimal policy involves price stability, i.e., it tightly stabilizes inflation at zero, and that simple rules that exclusively focus on deviations from price stability get the economy very close to the optimum.

These results provide a nice contrast between optimal monetary policy in environments with and without nominal rigidities. When nominal prices are flexible and there is a well-defined demand for real balances, a zero nominal interest and, hence, deflation minimize the welfare costs from holding money. Furthermore, if in a stochastic environment fiscal policy has to use distortionary taxes to finance given expenditures, mean zero unanticipated changes in the inflation rate represent lump-sum taxes and are an efficient way to raise revenues. Thus, optimal policy leads to low and volatile inflation. In contrast with nominal rigidities, deviations from price stability introduce relative price distortions among the monopolistically competitive producers and make production inefficient. Schmitt-Grohé and Uribe argue that in environments that contain both a well-defined demand for real money and nominal rigidities, even small amounts of nominal rigidities imply that price stability is optimal. This is a useful result since the surveys of Nason and Smith and Schorfheide provide some evidence for the presence of nominal rigidities, but also show that there is no agreement on how substantial nominal rigidities are.

Optimal policies that determine fiscal and monetary policies jointly can be quite complicated, yet Schmitt-Grohé and Uribe show that simple policy rules involve only minor welfare losses relative to the optimal policy. These simple rules are modeled on the Taylor rule that has the nominal interest responding to deviations of inflation and output from their targets with some dependence on past interest rates. It turns out that a simple rule that aggressively targets price stability involves only minimal welfare losses relative to the optimal policy, and that a response to deviations of output from trend significantly decreases welfare. An open question remains why most monetary policymakers prefer to target some positive inflation rate rather than price stability with a zero inflation rate.

Conclusion

The surveys in this special issue show that discussions of the Phillips curve have been at the core of monetary policymaking since the 1960s. Our understanding of what underlies the correlation between unemployment and the inflation rate and what that means for monetary policymaking has changed over the years. At first, many economists and policymakers took the statistical relationship as a fixed menu of choices between inflation and unemployment and targeted relatively low unemployment outcomes. From the period of high inflation and high unemployment in the 1970s, economists emerged believing that there is no inflation-unemployment tradeoff that remains invariant to policy interventions, and policymakers agreed that the objective of monetary policy should be low and stable inflation. Finally, in the 1990s, economists again started to study the inflation-output tradeoff using the new techniques developed in macroeconomics in the 1970s and 1980s, rational expectations and explicit quantitative general equilibrium models of the aggregate economy. This research program gave rise to the NKPC, which is based on the maintained assumption of nominal price rigidities. As is apparent from the surveys in this issue, there is some support for the NKPC in aggregate data, but there is no agreement on the extent of nominal price rigidities in the aggregate economy. Furthermore, one should be aware that not all macroeconomists agree that nominal rigidities are relevant for an understanding of the aggregate economy, e.g., see Williamson (2008) or Chari, Kehoe, and McGrattan (2009) for a skeptical view on this research program. To be sure, research on the relationship between unemployment and inflation will remain an active area in macroeconomics for anyone with an interest in applied monetary economics.

REFERENCES

Blanchard, Olivier Jean, and Nobuhiro Kiyotaki. 1987. "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review* 77 (September): 647–66.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan. 2009. "New Keynesian Models: Not Yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics* 1 (January): 242–66.

Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review* 58 (March): 1–17.

Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (October): 195–222.

Krause, Michael, David López-Salido, and Thomas Lubik. 2008. "Inflation Dynamics with Search Frictions: A Structural Econometric Analysis." *Journal of Monetary Economics* 55 (July): 892–916.

Kydland, Finn E., and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (November): 1345–70.

Long, John B., Jr., and Charles I. Plosser. 1983. "Real Business Cycles." *Journal of Political Economy* 91 (February): 39–69.

Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (April): 103–24.

Phelps, Edmund S. 1968. "Phillips Curve, Expectations of Inflation, and Optimal Inflation over Time." *Economica* 34: 254–81.

Sargent, Thomas J., and Neil Wallace. 1975. "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *Journal of Political Economy* 83 (April): 241–54.

Williamson, Stephen D. 2008. "New Keynesian Economics: A Monetary Perspective." Federal Reserve Bank of Richmond *Economic Quarterly* 94 (Summer): 197–218.

Woodford, Michael. 2003. *Interest and Prices*. Princeton, N.J.: Princeton University Press.

Economic Quarterly—Volume 94, Number 4—Fall 2008—Pages 311–359

The Phillips Curve and U.S. Macroeconomic Policy: Snapshots, 1958–1996

Robert G. King

The curve first drawn by A.W. Phillips in 1958, highlighting a negative relationship between wage inflation and unemployment, figured prominently in the theory and practice of macroeconomic policy during 1958–1996.

Within a decade of Phillips' analysis, the idea of a relatively stable long-run tradeoff between price inflation and unemployment was firmly built into policy analysis in the United States and other countries. Such a long-run tradeoff was at the core of most prominent macroeconometric models as of 1969.

Over the ensuing decade, the United States and other countries experienced stagflation, a simultaneous rise of unemployment and inflation, which threw the consensus about the long-run Phillips curve into disarray. By the end of the 1970s, inflation was historically high—near 10 percent—and poised to rise further. Economists and policymakers stressed the role of shifting expectations of inflation and differed widely on the costliness of reducing inflation, in part based on alternative views of the manner in which expectations were formed. In the early 1980s, the Federal Reserve System undertook an unwinding of inflation, producing a multiyear interval in which inflation fell substantially and permanently while unemployment rose substantially but temporarily. Although costly, the disinflation process involved lower unemployment losses than predicted by consensus macroeconomists, as rational expectations analysts had suggested that it would.

The author is affiliated with Boston University, National Bureau of Economic Research, and the Federal Reserve Bank of Richmond. Thanks to Michael Dotsey, Andreas Hornstein, Mark Watson, and Alexander Wolman for valuable comments. Devin Reilly helped with the figures. The views expressed herein are the author's and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: rking@bu.edu.

By 1996, the central bank of the United States had constructed a large-scale rational expectations model, without any long-run tradeoff, which it began to use to evaluate alternative policy scenarios. Monetary policymakers at that time accepted the idea that there was no long-run tradeoff at, and above, the then-prevailing price inflation rate of 3 percent. Yet many felt that there were important tradeoffs over short-run horizons, diversely defined, and some saw long-run tradeoffs near zero price inflation.

This article reviews the evolving role of the Phillips curve as an element of macroeconomic policy during 1958–1996, as well as academic and central bank research on it, via a series of snapshots over this roughly 40-year period. In conducting the research summarized in this article, my motivation is to better understand the mindset about the tradeoff between inflation and unemployment over an important period of U.S. history with an eye toward ultimately better understanding the joint behavior of the Federal Reserve and the U.S. economy during that period. Diverse research in macroeconomics—notably Sargent (1999), Orphanides (2003), and Primiceri (2006)—has sought an explanation of inflation's role in the behavior of a central bank that has an imperfect understanding of the operation of the private economy. The perceived nature of the Phillips curve plays an important role in these analyses, so that my reading of U.S. history may provide input into future work along these lines. I draw upon two distinct and complementary sources of information, published articles and documents of the Federal Open Market Committee (FOMC), to trace the evolving interpretation of the tradeoff over this roughly 40-year period.

The discussion is divided into six sections that follow this introduction. Section 1 provides a quick overview of the U.S. experience with price inflation and unemployment during 1958–1996. As the objective of this article is to provide a description of how policymakers' visions of the Phillips curve may have evolved during this time, resulting from empirical and theoretical developments, it is useful to have these series in mind as we proceed. Section 2 describes the birth of the Phillips curve as a policy tool, highlighting three core contributions: Phillips' original analysis of U.K. data, Samuelson and Solow's (1961) estimates of the curve on U.S. data and their depiction of it as a menu for policy choice, and the econometric analysis by Klein et al. (1961) and Sargan (1964) of the interrelationship between wage inflation, price inflation, and unemployment, which formed the background for wage and price blocks of macroeconomic policy models. Section 3 depicts the battle against unemployment that the United States waged during the 1962–1968 period and its relationship to the Phillips curve in then-prominent macroeconomic policy models. Section 4 discusses the breakdown of the empirical Phillips curve during 1969–1979, a period including intervals of stagflation in which unemployment and inflation rose together, and theoretical criticisms of the Phillips curve as a structural macroeconomic relation. Section 5 indicates the role of the Phillips curve during the unwinding of inflation in the United States during 1980 through 1986. Section 6 concerns several aspects of policy modeling and policy targeting in 1996 as the United States returned to a sustained interval of relatively low inflation. Section 7 concludes.

1. INFLATION AND UNEMPLOYMENT, 1958–1996

Since my discussion focuses on studies of inflation and unemployment that were written during 1958 through 1996, it seems useful to start by providing information on U.S. inflation and unemployment over that historical period, augmented by a few initial years, as in Figure 1. As measured by the year-over-year percentage change in the gross domestic product deflator, inflation averaged just under 4 percent, starting and ending the 1955–1996 interval at about 2.5 percent. Inflation twice exceeded 10 percent, in 1974–75 and 1981. The unemployment rate averaged 6 percent, starting and ending the sample period near 5 percent. Recession intervals, as dated by the National Bureau of Economic Research (NBER), are highlighted by the shaded lines in Figure 1.

My snapshots of the Phillips curve and its role in macroeconomic policy are usefully divided into five periods.

- *The formative years* in which the initial studies were conducted, 1955–1961. During this interval, there were two recessions (August 1957 through April 1958 and April 1960 through February 1961), each of which was marked by declining inflation and rising unemployment.

- *The battle against unemployment* from 1962 through 1968 during which unemployment fell substantially, with inflation being at first quiescent and then rising substantially toward the end of the period.

- *The breakdown of the Phillips curve* empirically and intellectually came from 1969 through 1979. In this period, there were two recessions. During December 1969–November 1970, both inflation and unemployment rose but there was a brief decline in inflation within the recession. During November 1973 through March 1975, inflation and unemployment both rose dramatically. This period was a tumultuous one, marked by departure from gold standard, wage and price controls, energy shocks, as well as difficult political and social events.

- *The unwinding of inflation* took place during 1980 through 1985, with a substantial reduction in inflation accompanied by a sustained period of unemployment.

- In *the aftermath*, 1986–1996, the Phillips curve assumed a new form in monetary policy models and monetary policy discussions.

Figure 1 Inflation and Unemployment, 1955–1996

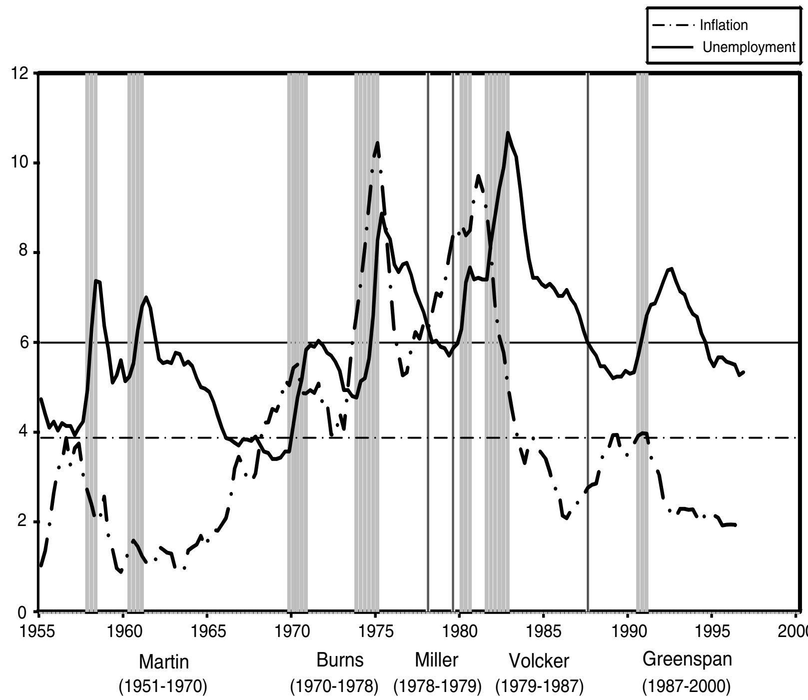

Notes: The inflation rate is the year-over-year change in the gross domestic product (GDP) deflator, the unemployment rate is the civilian unemployment rate, quarterly averages of monthly figures. All data from the Federal Reserve Economic Database (FRED) at the Federal Reserve Bank of St. Louis. The shaded grey areas are NBER recessions, while the vertical black lines separate the tenures of the various Federal Reserve chairman whose names appear below the horizontal axis.

2. THE FORMATIVE YEARS

Figure 2 is the dominant image from Phillips' initial article: a scatter plot of measures of wage inflation and unemployment in the United Kingdom over 1861–1913 supplemented by a convex curve estimated by a simple statistical procedure. During the 1960s, U.S. macroeconomic policy analysis and models were based on a central inference from this figure, which was that a permanent rise in inflation would be a necessary cost of permanently reducing unemployment. However, as background to that period, it is useful for us to understand how the Phillips curve was estimated initially, how it crossed the Atlantic, and how it was modified so that it could be imported into macroeconomic policy models.

Figure 2 Phillips' Curve

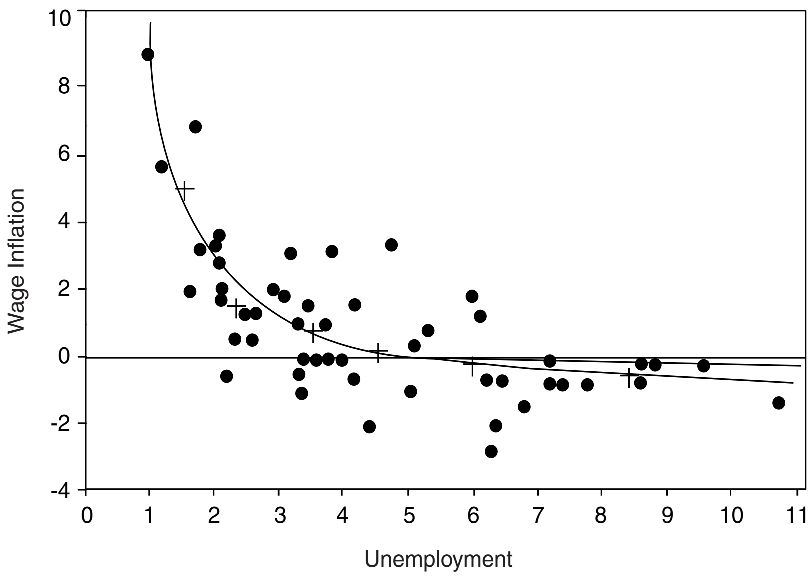

Notes: This is Figure 1 from Phillips (1958), displaying the relationship between unemployment and wage inflation over 1861–1913. The dots represent annual observations, while the crosses represent trade cycle averages.

The Original Study

Phillips (1958) described the objective of his study as follows: "to see whether statistical evidence supports the hypothesis that the rate of change of money wage rates in the United Kingdom can be explained by the level of unemployment and the rate of change of unemployment, except in or immediately after those years in which there was a very rapid rise in import prices, and if so to form some quantitative estimate of the relation between unemployment and the rate of change of money wage rates."[1] He began with the study of inflation and unemployment over multiyear periods, which he called trade cycles, and then he assembled these intervals into the overall curve that bears his name.

Trade cycles and the Phillips curve

The celebrated trade-off curve was derived by a complicated procedure. First, Phillips explored the behavior of a measure of wage change and unemployment

[1] Phillips (1958, 284).

over a series of historical United Kingdom "trade cycles," an alternative label for the sort of business cycles that Burns and Mitchell (1946) had identified for the United States.[2] The cycle for 1868–1879 is shown in Panel A of Figure 3. It begins with several years of falling unemployment and rising wage inflation, then an interval of rapidly declining wage inflation and modestly rising unemployment, then a number of years of wage declines accompanied by substantially increasing unemployment. Over the course of this cycle, there was an initial interval (1868–1872) during which inflation rose by about 10 percent, while unemployment dropped by about 5 percent. Then, from 1872–1875, there was a period of sharply declining inflation accompanied by modestly rising unemployment. Finally, from 1876–1879, there was a period of negative inflation (-1 to -3 percent per year) coupled with dramatically rising unemployment.

Phillips' identification of the tradeoff between inflation and unemployment did not rely on the shape of the cyclical pattern over the course of this and other individual trade cycles. Instead, the wage inflation and unemployment observations over the 1868–1879 trade cycle were averaged by Phillips to produce one of the "+" points in Figure 2, with the long-run curve adjusted so that it fit through these cycle averages.[3] The curve, fitted to six "+" points, contained three free parameters and implied that very low values of unemployment would lead to very high inflation, while very high values of unemployment would lead to very low inflation.[4]

Thus, the Phillips curve was based on average inflation and unemployment observations over the course of trade cycles of varying lengths. Although it was sometimes criticized as capturing short-run relations, Phillips' procedure contained significant lower frequency information. Yet, these cycle averages were drawn from the period when the United Kingdom was on the gold standard so that there were limits to the extent of price inflation or deflation.

Exploration of subsequent periods

After estimating the long-run curve on 1861–1913 data, Phillips then examined the extent to which the subsequent behavior of wage inflation and unemployment could be understood using the curve.

---

[2] Phillips' annual wage inflation observations are effectively a two-period average of the inflation rate in the future year and the current year.

[3] To explain the cyclical pattern around the long-run curve, Phillips developed a theory in which wage inflation was affected negatively both by the rate of change and level of unemployment. That part of his analysis was less broadly taken up by subsequent researchers, although there was a significant literature on "Phillips loops" during the 1970s.

[4] As in Phillips, letting y be the wage inflation and x be unemployment, the fitted curve took the form $y = -.9 + 9.638x^{-1.394}$, with the paremeters selected by a combination of least-squares and trial-and-error (Phillips 1958, 285).

Figure 3 Trade Cycles and the Phillips Curve

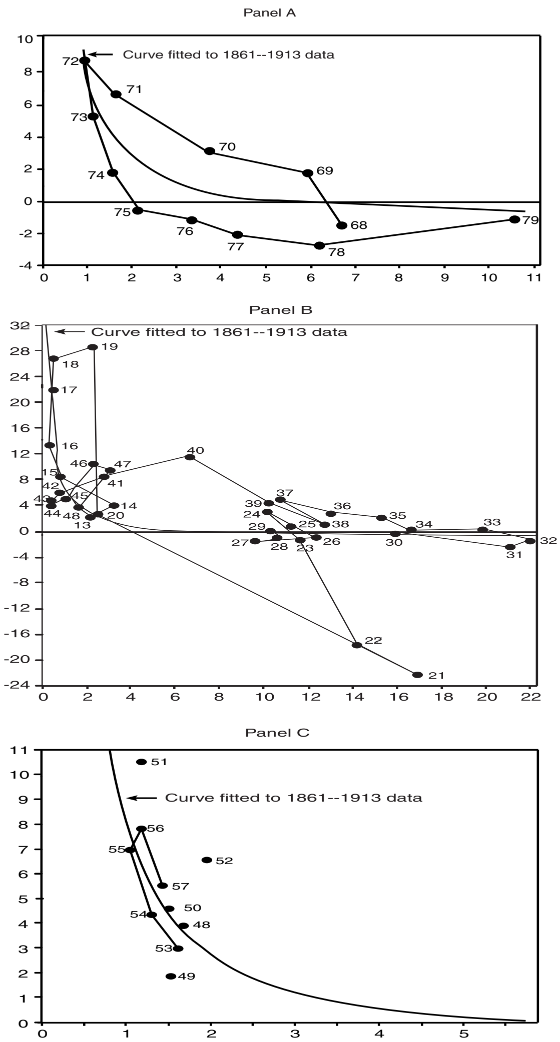

Notes: A key part of Phillips' analysis was to study the behavior of unemployment and wage inflation over various trade cycles, with three early cycles shown in this figure. For each, Phillips computed a cycle average, which was then used as one of the central data points through which he drew the long-run curve shown in Figure 1 (+). Panel A is Figure 3 from Phillips (1958); Panel B is Figure 9 from Phillips (1958); Panel C is Figure 10 from Phillips (1958). In all panels, the horizontal axis is unemployment and the vertical axis is wage inflation, as in Figure 2.

Looking at 1913–1948 as shown in Panel B of Figure 3, Phillips concluded that the general approach worked well. In particular, the trade cycle of 1929–1937 fit his general pattern, but he puzzled somewhat over the relatively high rates of inflation in 1935–1937. Other parts of the 1913–1948 interval fit in less well. As potential explanations of the behavior of wage inflation during both the First World War and the subsequent deflation to return the pound to its pre-war parity, Phillips discussed the potential importance of cost-of-living changes (effects of price inflation on wage-setting) as contributing to wage inflation in the wartime period and to wage deflation during the post-WWI interval). The sharp declines in nominal wage rates in 1921 and 1922, as Britain returned to pre-war parity with gold, stand out dramatically in Panel B of Figure 3. Although these points lie far from the curve based on 1861–1913 data, they are dramatic outliers that nevertheless show a negative comovement of inflation and unemployment in line with Phillips' general ideas.

Phillips also explored the consistency of the period following the Second World War, 1948–1957, with his long-run curve. During this period, U.K. unemployment was at a remarkably low level (between 1 and 2 percent) and inflation varied widely, as shown in Panel C of Figure 3. Phillips commented on several aspects of this interval. First, he noted that a governmental policy of wage restraint was in place in 1948 and apparently temporarily retarded wage adjustments. Second, he noted that the direction of the trade-cycle "loop" had reversed from the earlier period, which he suggested might be due to a lag in the wage-setting process. Third, he used this period to show how his curve could be used to partition wage inflation into a "demand-pull" component, associated with variation in unemployment along the curve and other factors, which induced departures from the curve. After looking at retail price inflation during this period, Phillips suggested that some of the wage inflation observations, such as the 1948 value that lies well above the curve, could have arisen from "cost-push" considerations in which workers bargained aggressively for higher nominal wages. However, Phillips concluded that the post-WWII period was broadly consistent with the curve fit to the 1861–1913 data.

U.S. Background to a Vast Experiment

In 1960, Paul Samuelson and Robert Solow examined U.S. data on the rate of change in average hourly earnings in manufacturing and the annual average data on unemployment over an unspecified sample period, which is most likely 1890 through the late 1950s.[5] Their empirical analysis most closely resembles

[5] Samuelson and Solow (1960) indicate that their study is based on the data of Rees, which is most likely his 1961 monograph on real wages in manufacturing, where the earliest data is 1890.

Figure 4 Phillips Curves of Samuelson and Solow

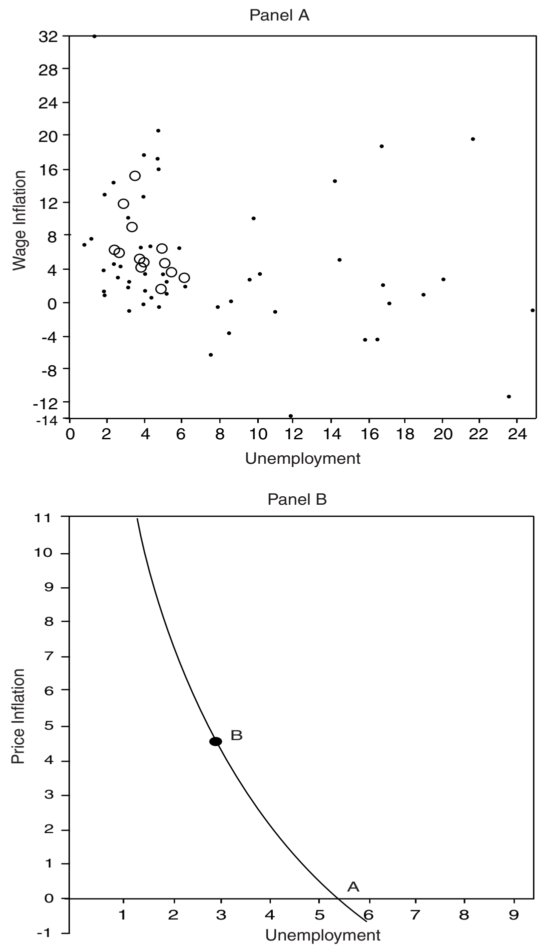

Notes: Panel A shows annual unemployment and wage inflation in U.S. data (this is Figure 1 from Samuelson and Solow [1961]), with their figure notes indicating that circled points are for "recent years." Panel B shows the trade-off curve that Samuelson and Solow discussed for the United States (this is Figure 2 in their article).

the 1861–1913 Phillips analysis that we have just looked at, but the overall association was looser, as dramatically displayed in Panel A of Figure 4.

Like Phillips, Samuelson and Solow looked at sub-samples, noting that money wages rose or failed to fall during the high unemployment era of 1933 to 1941, which they suggested might be due to the workings of the New

Deal. Further, they noted that the World War I period also failed to fit into the expected pattern, in line with Phillips' findings discussed above.

Overall, though, Samuelson and Solow argued that "the bulk of the observations—the period between the turn of the century and the first war, the decade between the end of that war and the Great Depression, and the most recent ten or twelve years—all show a rather consistent pattern. Wage rates do tend to rise when the labor market is tight, and the tighter the faster." They noted with interest that "the relation, such as it is, has shifted upward slightly but noticeably in the forties and the fifties." In the early years, before and after the first war, "manufacturing wages seem to stabilize absolutely when 4 or 5 percent of the labor force is unemployed; and wage increases equal to the productivity increase of 2 to 3 percent per year is the normal pattern at about 3 percent unemployment" and described this finding as "not so terribly different" from Phillips' results. In the later years, 1946–1959, Samuelson and Solow judged that it "would take more like 8 percent unemployment to keep money wages from rising and that they would rise at 2 to 3 percent per year with 5 or 6 percent of the labor force unemployed."[6] It is these later years that Samuelson and Solow circled in Panel B of Figure 4.

To describe the policy implications of their findings, Samuelson and Solow (1961) drew a version of the Phillips curve as representing tradeoffs between *price* inflation and unemployment. Essentially, this involved using the idea that price inflation and wage inflation were different mainly by the growth of labor productivity, suggesting an implicit model of relatively quick pass-through from wages to prices.

To obtain price stability under the assumption that real wages would grow at 2.5 percent per year, they suggested that the American economy would have to experience a 5 to 6 percent rate of unemployment (this option is marked as point A in Panel B of Figure 4). By contrast, they suggested that "in order to achieve the nonperfectionist's goal of 3 percent unemployment, the price index might have to rise by 4 to 5 percent per year" (this option is marked as point B).[7]

Seeking to understand whether inflation originated from cost-push or demand-pull factors, Samuelson and Solow described a "vast experiment" in which "by deliberate policy one engineered a sizeable reduction in demand" so as to explore the effects on unemployment and inflation. Although they were not explicit about the mechanism, they likely shared the prevailing Keynesian view of the time that fiscal and other policies that cut aggregate demand would first increase unemployment, with higher unemployment then reducing wage and price inflation. One interpretation of the subsequent 30

---

[6] All quotations in this paragraph are from Samuelson and Solow (1961, 189).

[7] Samuelson and Solow (1961, 192).

years of U.S. history is that versions of such experiments, with both increases and decreases in demand, were repeatedly undertaken.[8]

The Samuelson and Solow analysis led to a detailed research program of estimating the long-run tradeoff between inflation and unemployment in the United States. Given that interpretation and the subsequent development of macroeconometric models, it is interesting to note that Samuelson and Solow (1961) included a foreshadowing of future critiques of the long-run tradeoff: "aside from the usual warning that these are simply our best guesses, we must give another caution. All of our discussion has been phrased in short-run terms, dealing with what might happen in the next few years. It would be wrong, though, to think that our menu (Figure 4B) that relates obtainable price and unemployment during the next few years will maintain its shape in the longer run." They pointed to two reasons for potential instability—one was that "wage and other expectations" might shift the position of the Phillips curve and the other was that "institutional reforms" including product and labor market regulations or direct wage and price controls might shift the American Phillips curve downward and to the left.[9] Both expectations and wage-price controls were to play an important role in the subsequent history of the Phillips curve in the United States and other countries.

Wages, Prices, and Lags

Macroeconomic models along Keynesian lines first aimed at capturing the dynamics of aggregate demand. Thus, for example, the Duesenberry, Eckstein, and Fromm (1960) simulation study of the U.S. economy in recession used a quarterly econometric model with 14 equations governing aggregate demand: it contained neither a monetary sector nor a wage-price block. That is, the interaction between shocks and the components of aggregate demand was viewed as first order for understanding the behavior of the U.S. economy in a recession, with implications for wages and prices or their influences taken as less important. Fiscal policy measures rather than monetary policy measures were introduced in many studies of the time, reflecting a professional focus on fiscal rather than monetary policy tools.

Yet, after these first stages, U.K. and U.S. modelbuilders introduced a block of equations for wages and prices, stimulated in part by the work of Phillips (1958). The monograph by Klein et al. (1961) reports on a multiyear project to construct quarterly U.K. data and to estimate an econometric model with a

[8] All quotations in this paragraph are from Samuelson and Solow (1961, 191).

[9] All quotations in this paragraph are from Samuelson and Solow (1961) page 193 except for the final one, which is from page 194.

wage-price sector. These authors found that distributed lags were important in the wage and price equations.[10]

In contrast to the specifications of Klein et al. (1961), one notable element of Sargan's (1964) investigation was that he required that his equations display homogeneity, so that the absolute levels of wages and prices were not important for model properties. Sargan, therefore, studied a wage equation of the form

$$W_t - W_{t-1} = \lambda(W_{t-1} - P_{t-1}) + \beta u_{t-1} + \gamma t + \xi f_t + \phi(P_{t-1} - P_{t-4}), \quad (1)$$

where W_t is the log nominal wage rate in quarter t, P_t is the log nominal price level, u_t is the unemployment rate, and f_t is a measure of the political party in power. He divided the analysis of this equation into two components. First, an equation that described wage changes as deriving from deviations from an equilibrium real wage,

$$W_t - W_{t-1} = \lambda[W_{t-1} - P_{t-1} - \overline{w}_{t-1}], \quad (2)$$

and a specification for the equilibrium real wage

$$\overline{w}_{t-1} = [\frac{\beta}{\lambda}u_{t-1} + \frac{\gamma}{\lambda}t + \frac{\xi}{\lambda}f_t + \frac{\phi}{\lambda}(P_{t-1} - P_{t-4})]. \quad (3)$$

This is simply an algebraic decomposition, but Sargan (1964) was insistent that the elements of the equilibrium wage process made internal sense, for example requiring that the coefficient $\frac{\gamma}{\lambda}$ is interpretable as the effect of productivity growth on the real wage. He also interpreted the parameter λ as a speed of adjustment toward the equilibrium.[11]

Following the work of Klein et al., Sargan also estimated a price equation that linked prices to wages. Sargan explored measures of productivity, demand, and relative input costs as additional determinants of prices. Combining the wage and price equations, Sargan was able to trace out dynamic consequences of changes in the unemployment rate on wages and prices. These were influenced by the strength of the equilibrating tendencies (λ) and the influence of the price terms (ϕ) from the wage equation. For example, even if there were no lags of wages in the wage equation, there still could be indirect effects coming from the presence of price lags.

[10] These results echoed the earlier findings of Fisher (1926) who had, in fact, *invented* the concept of a distributed lag for the purpose of empirical analysis of inflation and interest rates. More generally, the estimation of wage-price blocks has provided the basis for many advances in time series econometrics. In particular, Sargan (1964) used the wage-price block of Klein et al. (1961) as the basis for an investigation that was the starting point for the so-called London School of Economics (LSE) approach to econometric dynamics. For a recent study of the UK Phillips curve, using Sargan's work as its starting point, see Castle and Hendry (Forthcoming).

[11] Sargan also investigated generalization of the first specification to allow for additional lags of wage changes, $W_t - W_{t-1} = \lambda[W_{t-1} - P_{t-1} - w^e_{t-1}] + \sum_{j=1}^{J} \delta_j (W_{t-j} - W_{t-j-1})$ to enrich this dynamic adjustment process toward equilibrium.

Sargan (1964) concluded that there was a long-run tradeoff between wage inflation and unemployment, but that there were also lengthy average lags so that changes in unemployment and other variables would take several years to be fully reflected in wage inflation. These broad properties were widely built into Keynesian macroeconometric models, as wage and price sectors were added to the initial aggregate demand constructions.

3. THE BATTLE AGAINST UNEMPLOYMENT

The 1962 *Economic Report of the President* was the first prepared by the Kennedy Council of Economic Advisors (CEA), which was eager to implement "The New Economics" originating in the work of Keynes.[12] The 1962 *Report* discussed the origins of unemployment in labor market frictions and in aggregate demand conditions, concluding that the "objective of maximum employment" would have to use policies aimed principally at labor market conditions. The 1962 *Report* argued that "in the existing economic circumstances, an unemployment rate of about 4 percent is a reasonable and prudent full employment target for stabilization policy," further stressing that additional policy interventions to reduce structural unemployment would make it possible to further reduce that target. At the same time, the *Report* built the case that the macroeconomic conditions of the late 1950s and early 1960s had led to an "output gap" of between 4 and 10 percent. As shown in the top panel of Figure 5, the output gap was the difference between actual output and a smooth trend line, based on an assumed level and rate of growth of capacity. Based on the work of a young economist at the CEA (Okun 1962), unemployment was linked to the output gap, so that a 2-percentage-point higher unemployment rate was related to an output gap of 5 percent. That is, the *Report* built in an Okun's Law coefficient of 2.5 to produce the second panel of Figure 5. While the "Phillips curve tradeoff" is now frequently discussed in terms of inflation and the output gap using some version of Okun's Law, this article will maintain the original linkage between inflation and unemployment as its focus.

In fact, the Kennedy-Johnson administration did deliver a substantial decline in unemployment, as a look back at Figure 1 confirms. In keeping with the tenor of the times, in which a package of fiscal, structural, and monetary policies was viewed as necessary for and capable of producing this decline, the present article will not seek to separately identify the contributions of different types of policies. Histories of the period, such as Hetzel (2008, chapters 6 and 7), stress the coordination of fiscal and monetary decisionmaking, so that such an identification could be quite subtle.

[12] Walter Heller was the chairman of the Council of Economic Advisors from 1961–64, with the other members being Kermit Gordon and James Tobin. The terminology "new economics" was widely used at the time and apparently dates back to a 1947 volume by Seymour Harris.

Figure 5 Unemployment and Output Gap in the 1961 CEA Report

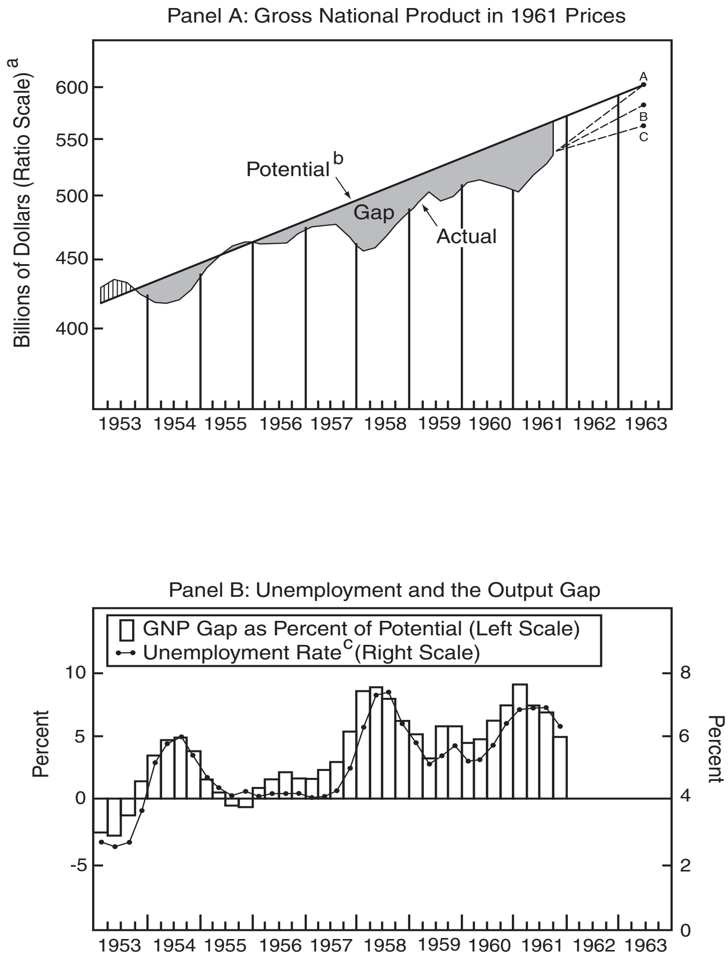

Notes: Panel A is the CEA's potential, actual, and output gap decomposition, while Panel B shows the link between the output gap and the unemployment rate. The CEA economists attached notes as follows: [a]Seasonally adjusted annual rates; [b]potential output based on a $3\frac{1}{2}$ percent trend line through the middle of 1955; [c]unemployment as a percent of civilian labor force, seasonally adjusted; A, B, and C represent GNP in the middle of 1963, assuming an unemployment rate of 4 percent, 5 percent, and 6 percent, respectively. They listed their sources as: Department of Commerce, Department of Labor, and Council of Economic Advisors.

The 1962 *Report* did note that "the economy last experienced 4 percent unemployment in the period May 1955–August 1957. . . . During this period, wages and prices rose at rates which impaired the competitiveness of some U.S. goods in world markets. However, there is good reason to conclude that upward pressures of this magnitude are not a permanent and systematic feature of our economy when it is operating in the neighborhood of 4 percent unemployment." Looking back at Figure 1, the reader will notice that the inflation rate rose by several percentage points during the 1955–1957 period alluded to in the CEA report, while unemployment averaged about 4 percent.

By the late 1960s, some version of the long-run Phillips curve tradeoff had become a cornerstone of economic policy. It entered centrally in macroeconomic models and more ephemerally in macroeconomic reports.[13]

Macroeconomic Models

In fall 1970, the Federal Reserve System sponsored a major conference on "The Econometrics of Price Determination," which contained a wide range of studies and later appeared in 1972 as a volume edited by Otto Eckstein. Drawn from one of these studies, Figure 6 displays the long-run relationship between price inflation and unemployment as of 1969 within three prominent macroeconomic models of the sort used by the U.S. private sector for forecasting purposes, by the executive branch of the U.S. government, and by the U.S. central bank. This figure is reproduced from the Hymans (1972) survey of the price dynamics within the Office of Business Economics (OBE) model used by the executive branch, the Federal Reserve-MIT-Penn (FMP) model used by the central bank, and the DHL-III model developed by Hymans and Shapiro at the University of Michigan.

As Hymans explains (1972, 313), this figure was produced by taking the wage-price block of the various models and evaluating these equations at alternative unemployment rates. One first finds the long-run inflation rate when the unemployment rate is constant at, say 5 percent, and then one finds the long-run inflation rate at 4 percent and so on.

By and large, these estimates of the long-run relationship accord well with that portrayed by Samuelson and Solow (1961) and reproduced as Panel B of Figure 4 in this article. Further, the increase in the inflation rate from about 1 percent in 1960–61 to about 4.5 percent in 1968–69 is particularly well captured by the FMP and DHL models. Although each long-run Phillips curve is nonlinear and there are differences in models, the "average" tradeoff

---

[13] As, for example, in the "Phillips plot without a Phillips curve" of the 1969 *Report*, discussed further below. Presumably, the economists at the CEA were not too interested in taking "credit" for the effect of low unemployment on inflation, while the economists at the Federal Reserve Bank (FRB) had a model that featured the tradeoff and could not escape the connection.

Figure 6 The Long-Run Tradeoff in Major Macroeconomic Models

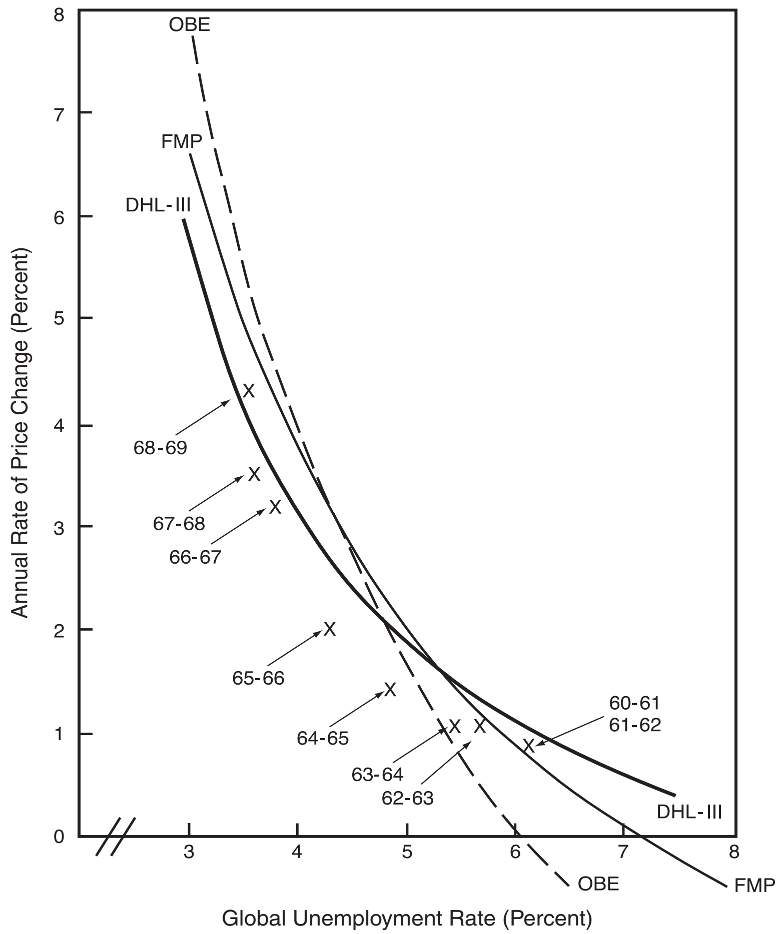

Notes: The long-run inflation and unemployment tradeoff in three prominent macroeconometric models as of 1970. The figure also shows actual unemployment and inflation (annual averages) during 1960–1969.

Source: Hymans (1972).

over this range is that lowering unemployment by 2.5 percent (from 6 to 3.5) costs about 3.5 percent in terms of inflation (from 1 percent to 4.5 percent). Smaller changes in inflation and unemployment feature a roughly one-for-one tradeoff.

The dynamics of wages and prices were studied by many authors under a variety of assumptions within the FMP and other large models. For example, de Menil and Enzler (1972) considered the effect of changing the unemployment rate from 4 percent to 5 percent in the FMP model: The results of their investigations are shown in Figure 7. The economy is assumed to be in an initial steady state with price inflation of 3.4 percent per year (wage inflation is just over 6 percent) and unemployment of 4 percent. Then, unemployment increased to 5 percent at date 1 and at all future dates. The inflation rate declines to 2.7 percent after four quarters, to 2.2 percent after three years and to 1.9 percent after five years. Compensation per man-hour (the wage measure) drops from 6 percent to 4.9 percent after a year, to 4.6 percent after three years, and to 4.4 percent after five years. The more rapid response of wage inflation is related to the fact that unemployment affects wages immediately, with effects of wages on prices occurring only with a distributed lag.[14]

Overall, the short-run Phillips curve reported by de Menil and Enzler (1972) is flatter than the long-run one that they report (essentially that for the FMP in Figure 6): a 1 percent increase in unemployment brings about a .7 percent change after a year's time, but a 1.5 percent decline in inflation after five year's time. From the standpoint of the econometric modelers of the time, this was a natural result of the lags in the wage-price components of their model, built in along Klein-Sargan lines. Looking at Figure 6, economists such as de Menil and Enzler likely saw a consistency with the dynamic specification of the macroeconomic policy model: the historical inflation rate initially lies below the long-run Phillips curve in the early 1960s during the start of the transition to lower unemployment.

Macroeconomic Reports

The 1969 *Economic Report of the President* was the last report of the Kennedy-Johnson era and was prepared under the leadership of Arthur Okun. With inflation rising, early 1968 saw a new cabinet-level Committee on Price Stability charged to recommend actions to contain inflation. President Johnson's introductory remarks in the Report distinguished between "roads to avoid" and "roads to reducing inflation." The roads to avoid were an "overdose of fiscal and monetary restraint" or "mandatory wage and price controls." The "roads to follow" included a combined fiscal and monetary program—including a continuation of the 1968 tax surcharge—as a "first line of defense," but also

[14] Note that de Menil and Enzler's (1972) experiment assumes an immediate and permanent change in unemployment induced by macroeconomic policies. Given the nature of the wage-price block in the FMP model, it was conceptually feasible to simply change unemployment and trace out the implications for wage and price inflation. However, since changes in fiscal and monetary instruments had only a gradual effect on aggregate demand and unemployment within the FMP and similar models, responses to more standard policy changes were more complicated and had gradual effects on both inflation and unemployment.

Figure 7 The Dynamic Tradeoff in the FMP Model circa 1970

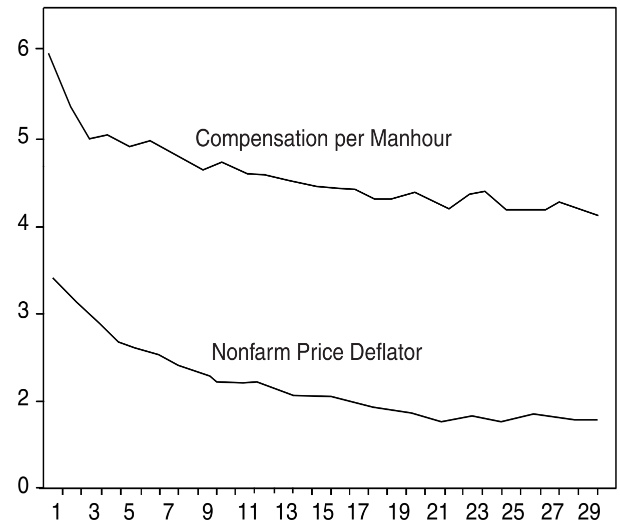

Notes: The vertical axis measures annualized percentage rates of change of the wage and price series. The horizontal axis is in quarters of a year.

Source: de Menil and Enzler (1972).

voluntary cooperation in wage and price setting to aid the process of reducing inflation.

The 1969 *Report* portrayed the U.S. economy as running at or slightly above potential output, with associated unemployment in the neighborhood of 4 percent. As portrayed in Figure 8, the potential output series resembles that presented in the 1962 *Report* (Figure 5), but the detailed notes make clear that potential output grew at 3.5 percent over 1955–1962; at 3.75 percent over 1963–1965; and at 4 percent over 1966–1968. Thus, an acceleration of potential output growth was necessary to fit together the 1962 *Report*'s view of 4 percent as the unemployment target, which finally was hit in the latter years of the Kennedy-Johnson era, with the behavior of output during those years. With more modest potential output, there would have been a very negative output gap during the final years of the Kennedy-Johnson era, whereas it is only slightly negative in Figure 8.

The 1969 *Report* also featured a Phillips scatter plot, from 1954–1968, highlighting the historical relationship between "price performance and

Figure 8 Unemployment and Output Gap in the 1961 CEA Report

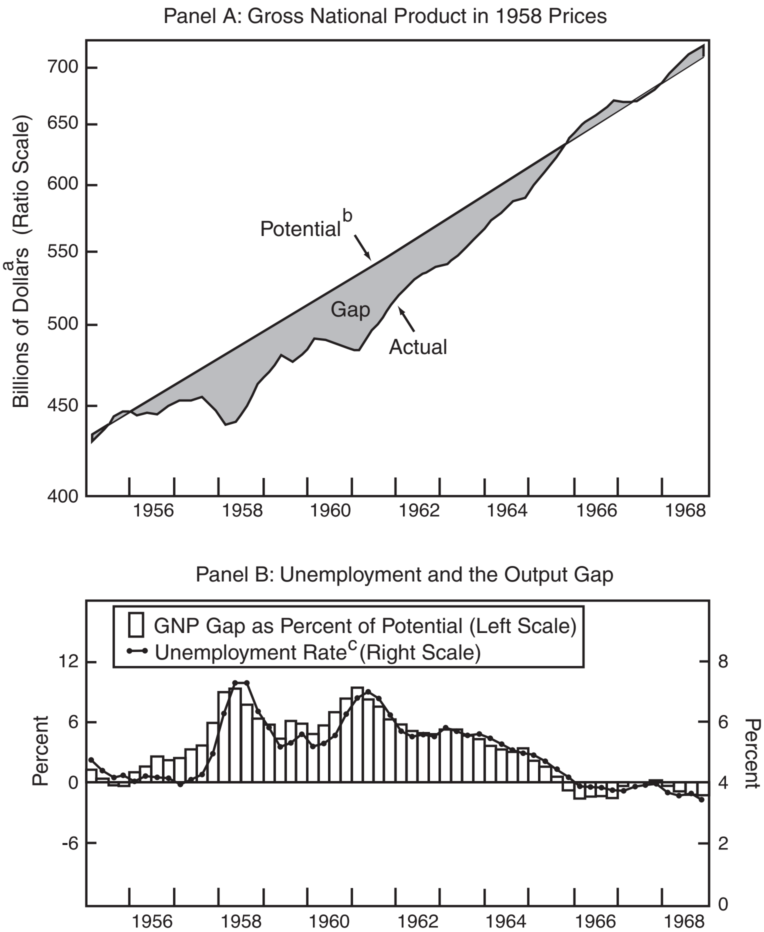

Notes: Panel A shows the CEA's potential, actual, and output gap decomposition, while Panel B shows the link between the output gap and the unemployment rate. The CEA economists attached notes as follows: [a]Seasonally adjusted annual rates;[b]potential output is a trend line of $3\frac{1}{2}$ percent through the middle of 1955 to 1962 IV, $3\frac{1}{4}$ percent from 1962 IV to 1965, and 4 percent from 1965 IV to 1968 IV; [c]unemployment as a percent of civilian labor force, seasonally adjusted. They list as sources: Department of Commerce, Department of Labor, and Council of Economic Advisors.

unemployment" during the Kennedy-Johnson years, although no trade-off curve was displayed. The decline in unemployment from the 5.5 percent level of 1963 to the 3.5 percent level of 1968 was associated with a rise in inflation from 1.5 percent to close to 4 percent. The *Report*'s accompanying discussion of the historical record stresses the importance of a wage-price spiral arising from demand growth in excess of capacity growth. It argued that "once such a spiral starts, it becomes increasingly difficult to arrest, even after productive capacity has caught up with demand and the initial pressures have largely subsided."

Thus, the *Report*'s explicit stand is that the U.S. economy in 1969 was operating close to capacity, with a rate of unemployment that was not necessarily inflationary, a viewpoint that echoes the appraisal in the 1962 *Report*. The rise in inflation over the eight years of the Kennedy-Johnson administration is, therefore, implicitly portrayed as arising from the effects of a wage and price spiral, not a purposeful movement along a long-run Phillips curve.

4. BREAKDOWN, 1969–1979

The breakdown of the consensus concerning the long-run Phillips curve involved a major revision of macroeconomic theory along with an unusual pattern of inflation and unemployment, with these intertwined developments reinforcing each other.

Macroeconomic Theory

As the Phillips curve played an increasing role in macroeconomic models, and as inflation rose during the mid-1960s, economists began to take a harder look at its theoretical underpinnings. The implications of new models were then compared to unemployment and inflation data, with results that sparked a major empirical controversy and a revolution in macroeconomic modeling.

The natural rate hypothesis

In the late 1960s, Milton Friedman and Edmund Phelps made separate arguments about why the long-run Phillips curve should be vertical. Friedman (1968) began from a vision of the labor market in which real wages and employment (or unemployment) were jointly determined in response to local and aggregate conditions of supply and demand. His natural rate of unemployment was "the level that would be ground out by the Walrasian system of general equilibrium equations, provided there is imbedded in them the actual structural characteristics of the labor and commodity markets, including market imperfections, stochastic variability in demands and supplies, the cost of gathering information about job vacancies and labor availabilities, the costs of mobility, and so on." He then imagined a situation in which firms offered workers nominal rather than real wages, with workers evaluating labor supply opportunities based on their best estimate of the purchasing power of those wages. With a known path for the price level, this calculation is easy for workers, but it becomes harder when the price level is changing. Friedman (1968) imagined the central bank increasing the growth of the money supply and stimulating the demand for the final product. To hire additional workers, firms would offer higher nominal wages. Faced with higher nominal wages that were interpreted as higher real wages, workers would supply more hours and potential workers would accept more jobs. So, it was possible for Friedman's model to reproduce a Phillips curve of sorts. However, if workers correctly understood that the general level of prices was increasing as a result of a monetary expansion, there would be no real effects: The rate of inflation and the rate of wage growth would jointly neutralize the effect of a higher rate of monetary growth, leaving real activity unaffected.

Phelps (1967) analyzed the problem in more Keynesian terms, based on a specification of a price equation of the following type:[15]

$$P_t - P_{t-1} = \pi_t = \beta(u_t - u^*) + \pi_t^e, \tag{4}$$

where u^* is the "natural rate of unemployment," in Friedman's terminology, and π^e is the expected rate of inflation. Phelps argued for this sort of "expectations-augmented Phillips curve" specification on grounds similar to those of Friedman that we have already discussed: labor suppliers should make their decisions on real, not nominal grounds.

Further, Phelps studied this inflation equation under the assumption of adaptive expectations,

$$\pi_t^e = \theta\pi_{t-1} + (1-\theta)\pi_{t-1}^e, \tag{5}$$

where $0 < \theta < 1$ governs the weight placed on recent information in forming expectations. This specification implies that if inflation were maintained at any constant level, $\overline{\pi}$, then expected inflation would ultimately catch up to it since the sum of coefficients in the distributed lag representation of expected inflation,

$$\pi_t^e = \theta \sum_{j=0}^{\infty} (1-\theta)^j \pi_{t-j-1},$$

is equal to one.

The "expectations-augmented Phillips curve" means that unemployment would be low ($u < u^*$) only if agents are surprised. Thus, a policy of maintaining low unemployment requires consistent underforecasting of inflation, $\pi < \pi^e$. But low unemployment could only be brought about by raising

---

[15] There are several cosmetic differences with Phelps' (1967, equation 3) specification. First, Phelps worked in continuous time while the text equation is in discrete time. Second, Phelps' specification is in terms of a general utilization variable rather than unemployment. Third, Phelps worked with the expected return on money, which is the negative of the expected inflation rate. Fourth, Phelps employed a nonlinear (convex) specification of the link from utilization to inflation rather than a linear one as in the text.

the inflation rate to a higher and higher level, with expectations always lagging behind because of the adaptive mechanism (5). Hence, the view of Friedman and Phelps became known as the "accelerationist hypothesis" in some quarters.

Phelps also stressed that the accelerationist model meant that a temporary period of low unemployment would bring about a permanently higher rate of inflation. He recognized that this led to an important new dynamic element in policy design, relative to traditional work that had stressed the Phillips curve as a stable "menu of policy choice" for the long run. Generally, a period of temporarily high unemployment would be necessary to permanently reduce inflation and some basic economic mechanisms—such as a momentary social planner objective that attributes increasing cost to high unemployment—making it desirable to smooth the adjustment process.

Overall, the natural rate/accelerationist hypothesis moved the worries of Samuelson and Solow (1961) about the effects of expectations on the Phillips curve from second-order to first-order status.

Tests of Solow and Gordon

From the perspective of wage and price adjustment equations of the form prominent in Keynesian theoretical and empirical models, the arguments of Friedman and Phelps suggested that there were important omitted expectational terms. Solow (1969) and Gordon (1970) devised tests to determine whether there was a long-run tradeoff between inflation and unemployment. To capture the spirit of these tests, consider a wage equation along the Klein-Sargan lines, such as (1) above. As above, the nominal wage at a given date will be W_t and the real wage will be $w_t = W_t - P_t$. However, the inflation terms in (1) are replaced by expected inflation, π^e_t, with a coefficient α attached,

$$\Delta W_t = \lambda(w_{t-1} - w^*_{t-1}) + \beta(u_t - u^*_t) + \alpha\pi^e_t; \tag{6}$$

other time-varying terms will be omitted for simplicity.[16]

The tests of Solow and Gordon made diverse use of price and wage equations, but the essential features can be simply described using this expression. First, in the wage equation above, the Friedman-Phelps conclusion obtains if $\alpha = 1$ since this is simply a restriction on expected real wages and unemployment,

$$w_t - w_{t-1} = \lambda(w_{t-1} - w^*_{t-1}) + \beta(u_t - u^*_t) - (\pi_t - \pi^e_t). \tag{7}$$

There is, thus, no influence of inflation if expectations are correct and no tradeoff between real and nominal variables. Solow and Gordon proposed

[16] In this expression, w^* would be the "natural" real wage, similar to the natural rate of unemployment, u^*.

to directly estimate the parameter α and to evaluate the accelerationist view by testing whether α differed significantly from unity. Second, this test is challenging to implement because expectations are unobservable. However, if expectations are formed adaptively, as in (5), then it is possible to conduct the test. Solow (1969) estimated parameters such as α for a range of different values of θ, while Gordon (1970) used a more general distributed lag but maintained the requirement that the coefficients summed to unity. This sum of the coefficients restriction was rationalized by the Phelpsian thought experiment, comparing a zero inflation steady state to a positive inflation steady state at rate π: If the sum of coefficients is one, then expected inflation is $\pi^e = 0$ in the first case and $\pi^e = \pi$ in the second case.

All of their diverse estimates suggested values of α that were positive, but significantly less than one. Thus, an aggregate demand policy that lowered unemployment in a sustained manner would create rising inflation over time, as expectations increased, but it would not ultimately be unsustainable.

Rational expectations critique

Sargent (1971) and Lucas (1972a) criticized the tests of Solow and Gordon by invoking two arguments. First, Sargent and Lucas insisted that expectations formation should be rational along the lines of Muth (1961). Second, they constructed example economies in which the accelerationist position was exactly correct, but in which an econometrician using the methods of Solow and Gordon would reach an incorrect conclusion.

A valuable example arises when inflation (π) has a persistent (x) and temporary (η) component, so that it is generated according to

$$\begin{aligned} \pi_t &= x_t + \eta_t \\ x_t &= \rho x_{t-1} + e_t, \end{aligned}$$

where η_t and e_t are serially uncorrelated, zero mean random variables and $|\rho| < 1$.

As analysis along the lines of Muth (1960) determines, rational expectations then are formed according to

$$E_{t-1}\pi_t = \rho[\theta\pi_{t-1} + (1-\theta)E_{t-2}\pi_{t-1}] = \rho\theta \sum_{j=0}^{\infty}(1-\theta)^j \pi_{t-j-1}, \qquad (8)$$

with $0 < \theta < 1$. This is broadly the same form as the adaptive expectations formula above, except that the distributed lag now is multiplied by ρ, which captures the degree of persistence of inflation.[17]

[17] The inflation process implies $\pi_t = \rho\pi_{t-1} + e_t + \eta_t - \rho\eta_{t-1}$. A "Wold representation" is $\pi_t = \rho\pi_{t-1} + a_t - \rho(1-\theta)a_{t-1}$, with a_t a forecast error for π_t relative to its own past history. These two processes are observationally equivalent if they imply the same restrictions on

In particular, suppose that inflation is somewhat persistent but stationary so that $0 < \rho < 1$, and that there is no effect of inflation on the real wage, i.e., $\alpha = 1$. Then, an econometrician constructing a measure of expectations $\widehat{\pi}_{t-1} = \theta \sum_{j=0}^{\infty}(1-\theta)^j \pi_{t-j-1}$ for any $\theta < 1 = \rho\pi_{t-1}$ would face data generated according to

$$W_t - W_{t-1} = \gamma(w_{t-1} - w^*_{t-1}) + \beta(u_t - u^*_t) + \rho\widehat{\pi}_{t-1}, \tag{9}$$

and would conclude that $\alpha = \rho$. That is, the econometrician would estimate that there was a long-run tradeoff, $\alpha < 1$, even though none, in fact, existed.

The critique played an important role in the evolution of macroeconomic modeling. Lucas (1972b) built a small-scale general equilibrium macroeconomic model with a short-run Phillips curve arising without any long-run tradeoff, providing an analytical interpretation of Friedman's (1968) suggestion about the nature of the link between inflation and real activity. Lucas (1976) expanded the critique of the policy invariance of parameters within 1970s macroeconometric models into a general challenge, noting that similar difficulties were contained in the consumption function and the investment function, as well as the wage-price block. He stressed that rational expectations and dynamic optimization, which seemed useful as basic postulates for model construction, inevitably led to such problems. A major revolution in econometric model construction ensued, leading central banks around the world to develop new models for forecasting and policy analysis during the 1990s, as we discuss in Section 6.

Shifting unemployment-inflation tradeoffs

What is wrong with the sum of coefficients restriction employed by Solow and Gordon? Sargent (1971) notes that inflation through the mid-1960s did not display much serial correlation, so it is not well forecasted by a moving average with weights that sum to unity. Intuitively, in an economy in which there are never any permanent changes in inflation, rational expectations are not designed to guard against this possibility.

Yet, as inflation rose and stayed high through the 1970s, empirical model-builders found estimates of parameter values drifting toward one and were

---

the variance and first-order autocorrelation of $\pi_t - \rho\pi_{t-1}$. (All other autocovariances are zero.) Using the Wold representation, the forecast is

$$\begin{aligned} E_{t-1}\pi_t &= \rho\pi_{t-1} - \rho(1-\theta)a_{t-1} = \rho\pi_{t-1} - \rho(1-\theta)(\pi_{t-1} - E_{t-2}\pi_t) \\ &= \rho[\theta\pi_{t-1} + (1-\theta)E_{t-2}\pi_{t-1}] \end{aligned}$$

as reported in the text. The covariance restrictions imply $var(a_t) = var(\eta_t) + \frac{1}{1+\rho^2} var(e_t)$ and $\theta = 1 - \frac{var(\eta_t)}{var(a_t)}$.

more open to accelerationist models of the Phillips curve. I have searched for, but have not found, a 1970s study that succinctly shows this drift in coefficients. McCallum (1994) calculates the sum of coefficients implied by a fifth-order autoregression for inflation estimated over various periods: it is about one-third in 1966–67, two-thirds in 1968–70, and between .88 and 1.02 for 1973–1980. McCallum's exercise indicates why econometric modelers found increasing evidence for the accelerationist hypothesis as the evidence from the 1970s was added.

Up a derivative: NAIRU models

As inflation rose during the mid-1970s, the wage and price blocks of standard macroeconometric models were augmented in various ways. One route was to include expectations terms explicitly but to make these expectations respond sluggishly to macroeconomic conditions. Another, arguably initially more popular strategy originates in the work of Modigliani and Papademos (1975). These authors argued that empirical models of price and wage dynamics would have better success if they related changes in inflation measures to levels of real variables. Modigliani and Papademos viewed inflation as likely to accelerate if unemployment was low relative to a benchmark value and as likely to decelerate if unemployment was high relative to a benchmark. They argued that this feature, as indicated by points A and C in Figure 9, was far more important for policy analysis than the question of whether the long-run Phillips curve had a negative slope (as in P P´) or was vertical (as in F F´).[18] Further, they stressed that "the shading of an area on either side of NAIRU indicates both uncertainty about the exact location of NAIRU and the implausibility that any single unemployment rate separates accelerating and slowing inflation."

The NAIRU model has been commonly captured by a simple model of the form,

$$\pi_t = \pi_{t-1} + \beta(u_t - u^n), \tag{10}$$

a specification closely related to that of Phelps (equation [4]), but with past inflation replacing expected inflation. Moving up a derivative allowed for a continuation of empirical research on price and wage inflation, which investigated the consequences adding lags of inflation and unemployment as well as adding shift variables to the basic NAIRU model. Thus, such empirical investigations built in a particular assumption on the "accelerationist hypothesis" that differed from the earlier research following Phillips, but continued to study many similar questions. However, uncertainty about the location of

[18] In the original unnumbered figure early in the Modigliani and Papademos (1975) article, the non-accelerating inflation rate of unemployment is marked "NIRU," but Figure 9 follows the now-standard acronym. It also corrects a typo in the labelling of the vertical axis.

Figure 9 The NAIRU Substitute for the Original Phillips Curve

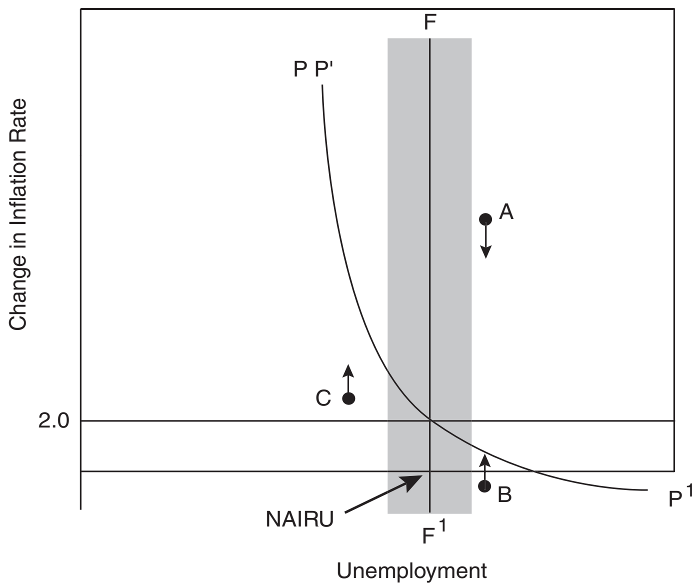

Notes: Modigliani and Papademos (1975) posited a relationship between unemployment and the *change* in the inflation rate. They stressed that there was a range of unemployment rates (gray area) over which the effect on inflation was quite uncertain. In the longer run, they did not take a firm stand on whether there was a tradeoff between the level of unemployment and the acceleration of inflation (the curve P P´) or not (the curve F F´).

the NAIRU, as Modigliani and Papademos had suggested, led to challenges in the application of this approach in forecasting and policy analysis.[19]

Macroeconomic Policy

The Nixon administration came into office in 1969 with the aim of reducing inflation via a combination of orthodox fiscal and monetary methods, but sought to do so without prolonging the recession that the country was then experiencing.

[19] Staiger, Stock, and Watson (1997) document the considerable uncertainty surrounding estimates of the NAIRU.

Gradualism

The Nixon administration initially embarked on a course of "policy gradualism" in an attempt to reduce inflation while maintaining real activity and unemployment at relatively constant levels. In a mid-course appraisal, Poole (1970) provides a useful definition of such policies: "The prescription of gradualism involves the maintenance of firm but mild restraint until the objectives of anti-inflationary policy are realized. Real output is to be maintained somewhat below potential until the rate of inflation declines to an acceptable level." The policies outlined in the 1970 *Economic Report of the President*, produced under the leadership of CEA chairman Paul McCracken, contained both fiscal and monetary components designed to generate a modest reduction in output for the purpose of reducing inflation.[20]

Rising understanding of the importance of expectations

Arthur Burns became head of the Federal Reserve System in early 1970, replacing William McChesney Martin, who had served since 1951. During 1969, Martin had undertaken restrictive monetary policy to reduce inflation, indicating that "expectations of inflation are deeply imbedded. . . . A slowing in expansion that is widely expected to be temporary is not likely to be enough to eradicate such expectations" and "a credibility gap has developed over our capacity and willingness to maintain restraint."[21] The phrase "credibility gap" had a particularly harsh ring to it, even as part of a self-criticism, as it had been widely used to describe the Vietnam policies of the Johnson administration: The Tet offensive of September 1968 had convinced many that there was no credibility to the administration's previous upbeat forecasts for military success or its description of the offensive as a disastrous defeat for the Viet Cong. Martin believed that unemployment was unsustainably low and that economic growth would need to slow to reduce inflation, which was at about 5 percent during late 1968 when Richard Nixon was elected president. But the restrictive monetary policy of 1969–1970 started by Martin, envisioned as part of a gradualist strategy by McCracken, and continued by Burns, resulted only in a slowing of inflation, but not a major decline, during the recession of December 1969–November 1970.

As Hetzel (1998) stresses, Burns had a long-standing belief that expectations were important for inflation, writing in the late 1950s that: "One of the main factors in the inflation that we have had since the end of World War II is

[20] Gradualist policies were advocated by a range of economists and policymakers. On this dimension, the monetarist economists of the Shadow Open Market Committee (SOMC) shared some of the reservations of their Keynesian counterparts. See Meltzer (1980) for a discussion of the SOMC perspective on gradualism.

[21] Martin's February 26, 1969, testimony to Congress, quoted in Hetzel (2008, 75).

that many consumers, businessmen, and trade union leaders expected prices to rise and therefore acted in ways that helped to bring about this result."[22] After taking over the Fed, he made the case that it was possible to side-step the Phillips curve through the imposition of wage and price controls, which he believed would exert a substantial effect on expectations.[23]

Incomes policies

By the middle of 1971, the Nixon administration lost patience with gradualism, imposing a 90-day temporary freeze on prices on August 15, 1971, and ending the convertibility of the dollar into gold. The temporary freeze evolved into a multiyear incomes policy with various phases differing in intensity and coverage. The 1973 *Economic Report of the President* noted that "1972 was the first full year in American history that comprehensive wage and price controls were in effect when the economy was not dominated by war or its immediate aftermath" (page 51). The *Report* indicated that there had been three purposes for the controls (page 53). First, the controls were intended to directly affect the rate of inflation, lower the probability of its increase, and raise the probability of its decline. Second, the controls were aimed at "reducing the fear that the rate of inflation would rise or not decline further." Third, the controls were designed to "strengthen the forces for expansion in the private economy and to free the Government to use a more expansive policy." In describing the conditions in the spring and summer of 1971 that led to the imposition of the controls, the *Report* specifically discussed "anxiety" about increasing inflation as holding back consumer spending and rising long-term interest rates that "may have signalled rising inflationary expectations." Overall, a key motivation for the controls was to affect expectations of inflation and their incorporation into price and wage setting.

Many analysts see the incomes policy period as involving expansionary monetary policy—as in the third *Report* point above—with fiscal policy under the Nixon team and monetary policy under Burns producing an economic expansion. Unemployment hovered in the 6 percent range through 1971, dropping to 5 percent by the end of 1972 as Nixon won a landslide re-election. Inflation, according to the gross domestic product deflator shown in Figure 1, had fallen from 5 percent in 1971 to 4 percent in 1972, but it then rose to a 7 percent annual rate by the end of 1973.

---

[22] Burns 1957, p. 71, quoted in Hetzel (1998).

[23] Burns had played a leading role in the Council of Economic Advisors during the Eisenhower administration and knew Nixon well. As chairman of the Fed, he played an important role in administration policy more broadly, including the mid-August 1971 meetings at Camp David that formulated major changes in Nixon administration economic policies.

Skepticism about government goals

By the time of the publication of the 1973 *Economic Report of the President*, many economists were becoming more skeptical about both the long-run Phillips curve and, more specifically, about whether government plans were consistent with the available information on the historical linkage between inflation and unemployment. The review of the 1973 *Report* by the mainstream economist Carl Christ provided one clear presentation of this skepticism. He noted that every economic report contains an overall statement of objectives by the President followed by a more detailed and nuanced report by his economic advisors.

An initial quote from Christ's review summarized the condition of the previous year for us: "Mr. Nixon's report begins with a review of the good things about 1972: a $7\frac{1}{2}$ percent rise in real output, a reduction of the inflation rate (measured by the consumer price index) to about 3 percent from about 6 percent in 1969, and a reduction of the unemployment rate to 5.1 percent in December 1972 from 6 percent in December 1971."

The 1973 goals of the Nixon administration were summarized by Christ using a series of quotations from Nixon's message: "Output and incomes should expand. Both the unemployment rate and the rate of inflation should be reduced further, and realistic confidence must be created that neither need rise again. The prospects for achieving these goals in 1973 are bright—if we behave with reasonable prudence and foresight" (p. 4); "We must prepare for the end of wage and price controls..." (p. 6). Christ notes that this buoyant optimism on inflation and unemployment is "reminiscent of Mr. Nixon's statement in January 1969, shortly after taking office, that he would reduce inflation without increasing unemployment and without imposing wage and price controls."

Christ then proceeded to argue that the optimism was unwarranted: "The evidence strongly suggests it is not possible for the American economy, structured as it has been since World War II, to achieve simultaneously unemployment rates that remain at 4.75 percent or less, and consumer price increases that remain at 2.4 percent a year or less, without wage or price controls. In the 25 years since consumer prices leveled off at the end of World War II, this has been achieved in only 4 years: 1952, 1953, 1955, and 1965.... In those same 25 years, the average unemployment rate was 4.8 percent, and the average increase in consumer prices was 2.4 percent a year."

Figure 10 was produced by Christ using data from the 1972 *Report*. Notice that Christ's argument is *not* that the long-run Phillips curve is vertical, although that view is not inconsistent with his figure. Instead, it is that public policy goals ought to be consistent with the available evidence and that unemployment far below the 1972 level of 5 percent and inflation below the 1972 level of 3 percent did not seem consistent with U.S. experience.

Figure 10 Christ's Summary of U.S. Inflation and Unemployment Experience

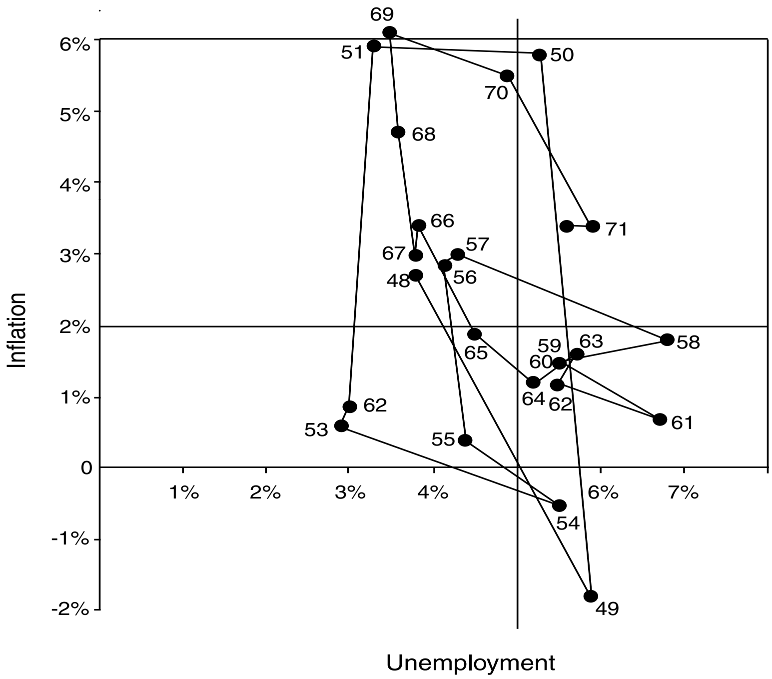

Notes: Using data from the 1973 Annual Report, Christ showed that the announced administration objectives of inflation below 3 percent and unemployment below 5 percent did not seem consistent with prior U.S. experience.

Christ's warning was a timely one: The remainder of the Nixon-Ford administration saw the onset of stagflation, as a look back at Figure 1 reminds us. The United States was not to see 5 percent unemployment and 3 percent inflation at any time in the next two decades.

Humphrey-Hawkins

While Christ may have been skeptical about the internal consistency of the *Economic Report of the President* in 1973, the debates over the legislation put forward by Representative Augustus Hawkins and Senator Hubert H. Humphrey five years later illustrated that other elements of government and society continued to seek very low unemployment.

As initially passed by the House in March 1978, the Full Employment and Balanced Growth Act specified the goal of lowering the unemployment rate to 4 percent for all working age individuals and 3 percent for all individuals over the age of 20. Early drafts of the House bill had mandated that the government be an "employer of last resort" for the long-term unemployed, but this provision was dropped, while the focus on unemployment was maintained. The national unemployment goal was to be reached within five years and the bill called for cooperation between the executive branch, Congress, and the Federal Reserve Board in working toward the specified target. It also specified that the President should submit an annual economic report to Congress including numerical goals for employment, unemployment, and inflation, as well as some other macroeconomic indicators. Amendments to add budget balance as a goal at the five-year horizon and to include an inflation goal of 3 percent at that time were defeated. The bill evolved substantially in order to gain Senate approval. In the process, inflation objectives were reinstated. It was signed into law by President Jimmy Carter on October 27, 1978.

The Full Employment and Balanced Growth Act in final form established national goals of full employment, growth in production, price stability, and balance of trade and public sector budgets. More specifically, it specified that by 1983, unemployment rates should be no more than 3 percent for persons aged 20 or over and no more than 4 percent for persons aged 16 or over. Inflation should be no more than 4 percent by 1983 and 0 by 1988. Thus, in its nonbinding goals, it displays the same tendencies that Christ identified in the *Economic Report of the President*.

While these goals were nonbinding, the Humphrey-Hawkins Act did require that the Federal Reserve Board of Governors transmit a report to Congress twice a year outlining its monetary policy.

5. UNWINDING INFLATION

By the late 1970s, a wide range of economists and politicians were becoming concerned about high inflation and recommending disinflation. However, economists and politicians differed widely on the costs of reducing inflation.

Forecasting the Costs of Disinflation

Surveying six estimates of "macroeconomic Phillips curves," Okun (1978) found that the experience of the 1970s had led to the abandonment of the long-run Phillips curve. Yet, he also stressed that "while they are all essentially accelerationist, implying no long-run tradeoff between inflation and unemployment, they all point to a very costly short-run tradeoff." Thinking about the Phelpsian question of how fast unemployment should be raised from a situation of initially low capacity utilization, in terms of the consequences for long-run inflation, he calculated "for an extra percentage point of unemployment maintained for a year, the estimated reduction in the ultimate inflation rate." Comparing the various studies, he found that this disinflation gain for a given amount of unemployment ranged between $\frac{1}{6}$ and $\frac{1}{2}$ of a percent, with an average estimate of 0.3 percent.

To put Okun's numbers in a specific context, consider the unemployment cost attached to a 5 percent reduction in long-run inflation. The estimates reported by Okun meant that this cost would be between $10 = [5/(1/2)]$ and $30 = [5/(1/6)]$ "point years of unemployment," with an average estimate of $16.7[5/.3]$. That is, if the cost of eliminating a 5 percent inflation was spread evenly over four years, then each year would see an unemployment rate that was between 2.5 and 7.5 percent above the natural rate with a mean estimate of over 4 percent.

It is now more standard to discuss disinflation costs as a ratio of point years of unemployment arising from a one percent change in inflation, which is called the "sacrifice ratio."[24] Okun's estimates were that the sacrifice ratio was in the range of 2 to 6, with a mean of 3.3 in that each percentage point reduction in inflation would involve very major economic costs.

Put another way, by the late 1970s, policymakers may have abandoned the long-run Phillips curve in the face of evidence and theory. But most major econometric models continued to maintain a tradeoff over horizons of four or more years, as originally described by Samuelson and Solow. Just as there had been a protracted period of low unemployment as inflation had risen, so too did Okun envision a protracted period of high unemployment as inflation was reduced.

The perceived severity of a potential reduction in inflation is perhaps best illustrated in an excerpt from James Tobin's (1980) review of stabilization policies at the close of the first decade of the Brooking's Panel. To put the excerpt in context, Tobin's review described the accelerationist hypothesis as having been a core part of macroeconomics for the better part of the previous decade. Tobin wrote that it was broadly recognized that "inflation accelerates at high employment rates because tight markets systematically and repeatedly generate wage and price increases in addition to those already incorporated in expectations and historical patterns. At low utilization rates, inflation decelerates, but probably at an asymmetrically slow pace. At the Phelps-Friedman "natural rate of unemployment," the degrees of resource utilization and market tightness generate no net wage and price pressures up or down and are consistent with accustomed and expected paths, whether stable prices or any other inflation rate. The consensus view accepted the notion of a nonaccelerating inflation rate of unemployment (NAIRU) as a practical constraint on

---

[24] The sacrifice ratio is now perhaps more commonly described as the output gap.

policy, even though some of its adherents would not identify NAIRU as full, equilibrium, or optimum employment."[25]

To put the potential costs of a disinflation in front of his audience, Tobin used a very simple inflation model with a NAIRU of 6 percent unemployment and assumed that the economy originated from an inflation rate of 10 percent. As displayed in Figure 11, he studied a gradual disinflation in which the central bank reduced the growth rate of nominal aggregate demand smoothly so that it falls by 1 percent each year for 10 years. Thus, after a decade, the conditions for price stability are met from the aggregate demand side. Tobin also assumed that the expectations term was an eight-quarter, backward-looking average of recent inflation rates. Tobin stressed that the result was "not a prediction!. . . but a cautionary tale. The simulation is a reference path, against which policymakers must weigh their hunches that the assumed policy, applied resolutely and irrevocably, would bring speedier and less costly results." The cautionary tale of Figure 11 involves an initial decade in which unemployment looks to average about 8.5 percent, 2.5 percent higher than its equilibrium value, so that the sacrifice ratio during this period is about 2.5 since inflation is being reduced by 10 percent. So, while the tale was cautionary, the message was consistent with the range of Okun's sacrifice ratio estimates and, hence, meant to depict some potential consequences of disinflation.

There were some skeptics. William Fellner (1976) viewed the government policies of the 1970s as sharply inconsistent with the objective of bringing about low inflation, echoing Christ's 1973 concerns. Fellner argued that households and firms would be similarly skeptical and that the disinflation process was costly in part because of the imperfect credibility of policies, so he endorsed a policy of gradualism like that which Tobin explored in his simulation coupled with strong announcements about future policy intentions. However, economists like Tobin were quite skeptical about the practical importance of this line of argument, while accepting the basic logical point that expectations effects could mitigate some of the output losses associated with his gradualist simulation. Considering the benefits of preannounced stabilization plan credibility, Tobin (1980) wrote: "The question is how much. One obvious problem is that a long-run policy commitment can never be irrevocable, especially in a democracy. Important economic groups will not find it wholly credible, and some will use political power to relax or reverse the policy. Even assuming credibility and understanding by private agents,

[25] Thus, Tobin uses the "natural rate" and the "NAIRU" interchangeably when it comes to the analysis of inflation. However, many Keynesian economists did not want to assume that the level of real activity consistent with constant inflation was an efficient level. As discussed above, Friedman had written that the natural rate was to include "the actual structural characteristics of the labor and commodity markets, including market imperfections, stochastic variability in demands and supplies, the cost of gathering information about job vacancies and labor availabilities, the costs of mobility, and so on," but these conditions had sometimes been ignored in the debate over the efficiency and inevitability of the natural rate.

Figure 11 Tobin's Cautionary Tale

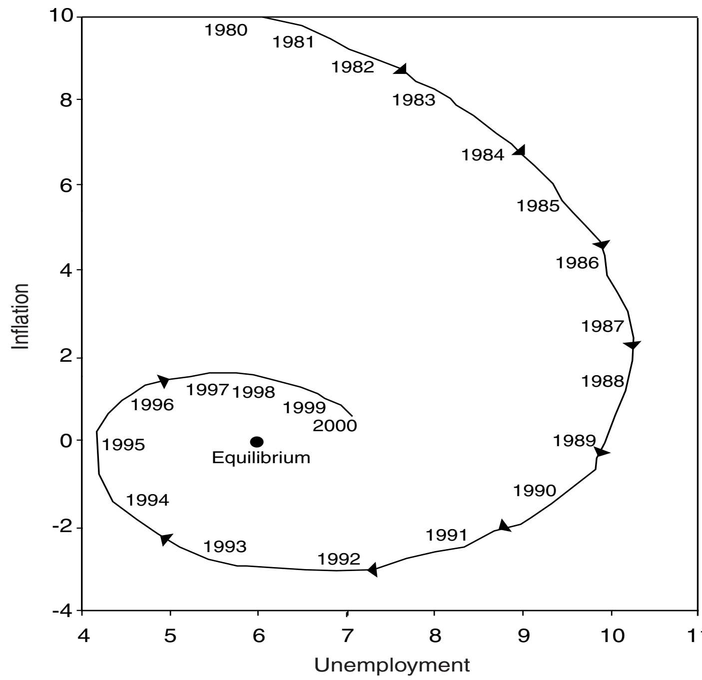

Notes: A simulation study of the unemployment and inflation implications of a gradual reduction in the rate of growth of the money supply, as described in more detail in Tobin (1980).

their responses are problematic. In the decentralized but imperfectly competitive U.S. economy, wage and price decisions are not synchronized but staggered. It is hard to predict how individual firms, employees, and unions will translate a threatening macroeconomic scenario into their own demand curves. If each group worries a lot about its relative status, each group will decide that the best strategy is to disinflate very little." Thus, Tobin (1980) argued that it would be "recklessly imprudent to lock the economy into a monetary disinflation without auxiliary incomes policies. The purpose of these policies would be to engineer directly a deceleration of wages and prices consistent with the gradual slowdown of dollar spending." In contrast to Fellner's case for gradualism, rational expectations theorists like Sargent began to explore actual disinflation experiences, using the lessons of the models that they had developed in the mid-1970s. In particular, in his "Ends of Four Big Inflations," circulated no later than spring 1981, Sargent argued that dramatic, sustained anti-inflation policies could bring about reductions in inflation with relatively low unemployment costs, as long as such policy changes were credible and that their dramatic nature enhanced their credibility.

The Volcker Disinflation

Paul Volcker assumed the chairmanship of the Federal Reserve System (FRS) in August 1979. Looking back at Figure 1, we can see that inflation was substantially reduced, while there was a lengthy period of high unemployment. As cataloged in many discussions, the Federal Reserve made a high-profile announcement of a shift to monetary targeting in October 1979 in the face of rapidly rising inflation; there were two recessions during the period, one short and relatively mild, one lengthy and severe, and the inflation rate had declined dramatically by 1984.

There are many questions about the Phillips curve during this important historical period, but our focus in this section will be limited to two. First, how did the unemployment cost of the actual disinflation line up with the suggestions of Okun and others? Second, how did the Federal Reserve perceive the menu of policy choice during this period?

The unemployment cost

Mankiw (2002, 369–701) calculates the unemployment cost of the Volcker disinflation under the assumption that there was a 6 percent natural rate of unemployment during 1982–1985, with the inflation rate falling from 9.7 percent in 1981 to 3 percent in 1985. His annual average unemployment numbers were 9.5 percent in 1982 and 1983, 7.4 percent in 1984, and 7.1 percent in 1985 so that there was a total cyclical unemployment cost of 9.5 percent of unemployment. One can argue about details of this calculation, for example with whether the disinflation should be viewed as starting in 1980 or in 1981, about the natural rate of unemployment, and so on. But it is a reference textbook calculation familiar to many: The sacrifice ratio during the Volcker disinflation is estimated by Mankiw to be about $9.5/6.7 = 1.5$.

Mankiw's sacrifice ratio is about one-half of that which Okun suggested on the basis of his mean estimate and lies below the low end of the range in the studies that he reviewed. The 6.7 percent decline in the inflation rate should have had the effect of raising the unemployment rate by a total of 22 percent over the period according to the average estimate, 13.4 percent according to the low estimate, and over 40 percent for the high estimate. Put another way, the mean estimate implies that unemployment should have been higher by more than 5 percent over each year of a four-year disinflation period.

Some have suggested that this lower cost was due to increased credibility of the Fed and its disinflationary policies under Volcker; others have suggested that the cost was largely due to the central bank's imperfect credibility.[26] However, as McCallum (1984) points out, the testing of hypotheses about credibility is subtle because the measures of private expectations about policy that must be constructed are more involved than in standard rational expectations models.

The Fed's Perceived Unemployment Cost

One key question about this disinflation period is "How did the Federal Reserve System view the tradeoff between inflation and unemployment?" The developments reviewed above, in which the nature of the tradeoff was subject to substantial controversy and evaluated in an evolving model, makes this a particularly interesting question. It is also not an easy question to answer, as any central bank is a large organization with many differing viewpoints and its policymakers do not file survey answers about their perceived tradeoffs.

However, the question can be answered in part because the Humphrey-Hawkins legislation requires testimony by the FRS chairman twice a year, in late January or early February and again in July. For the six FOMC meetings each year, the research staff under Volcker prepared a basic forecast of the economy's developments within the "Green Book" under a particular benchmark set of policy assumptions. For the FOMC meetings that precede the chairman's testimony, the staff also prepared a set of alternative policy options within the "Blue Book" of which Table 1 provides examples at two FOMC meetings. In both cases, then, the alternative strategies were framed in terms of growth rates for the M1 concept of money and were based on the Board's quarterly macroeconometric (MPS) model along with staff judgemental adjustments, so that they reflected the effects of pre-existing economic conditions as well as alternative paths of policy variables.[27] Projections

[26] To my mind, the role of imperfect credibility in the Phillips curve in U.S. history remains an open, essential area of research at present. Goodfriend and King (2005) argue that Volcker's actions were based significantly on his perception that his policy actions were not perfectly credible and that the nature of the disinflation dynamics was also substantially influenced by imperfect credibility of policy.

[27] The Federal Reserve Bank of New York (1998, 123) describes the preparation of "blue book" material, from which the Table 1 entries are taken, as follows: "The blue book provides the Board staff's view of recent and prospective developments related to the behavior of interest rates, bank reserves, and money. The blue books written for the February and July meetings contain two extra sections to assist the Committee in its preparation for the Humphrey-Hawkins testimony. The first of these sections provides longer-term simulations, covering the next five or six years. One of these simulations represents a judgmental baseline, while two or three alternative forecasts use a Board staff econometric model to derive the deviations from the judgmental baseline under different policy approaches. Typically, at least two scenarios are explored: one incorporates a policy path that is designed to bring economic activity and employment close to their perceived long-run potential paths fairly quickly, and another is intended to achieve a more rapid approach to stable prices. The section also offers estimates of how different assumptions about such factors as fiscal policy, the equilibrium unemployment rate, or the speed of adjustment to changed inflationary expectations would affect the predicted outcome."

Table 1 FRB Economic Projections Associated with Alternative Monetary Growth Strategies

January 1980

	Strategy	**1980**	**1981**	**1982**
Money Growth	1	6.0	6.0	6.0
	2	4.5	4.5	4.5
	$\Delta = 1 - 2$	−1.5	−1.5	−1.5
Inflation	1	9.5	8.7	7.7
	2	9.1	8.2	6.8
	$\Delta = 1 - 2$	−0.4	−0.5	−0.9
Unemployment	1	8.1	8.9	9.3
	2	8.4	10.1	11.6
	$\Delta = 1 - 2$	0.7	1.2	2.3

January 1982

	Strategy	**1982**	**1983**	**1984**
Money Growth	1	4.0	3.5	3.0
	2	5.5	5.0	4.5
	$\Delta = 2 - 1$	−1.5	−1.5	−1.5
Inflation	1	6.4	5.1	4.2
	2	6.5	5.4	5.3
	$\Delta = 2 - 1$	−0.1	−0.3	−0.9
Unemployment	1	9.3	9.1	8.9
	2	9.0	8.2	6.9
	$\Delta = 2 - 1$	0.3	0.9	2.0
T-Bill	1	13.0	12.5	11.0
	2	9.7	8.3	8.4
	$\Delta = 2 - 1$	3.3	4.2	2.6

Notes: These table items are drawn from the Federal Reserve Blue Books "Monetary Aggregates and Money Market Conditions," January 4, 1980, page 11, and "Monetary Policy Alternatives," January 29, 1982, page 7.

under alternative strategies during the Volcker deflation also took into account forecasted developments in fiscal policy, which were being revamped by the Reagan administration at the time.

The first meeting is the January 1980 FOMC session, at which the benchmark strategy (called strategy 1 in this meeting) was for 6 percent money growth over each of three years: 1980, 1981, and 1982. Under this benchmark policy, as can be seen by reading across the relevant row of the table, the forecast was that U.S. inflation would gradually decline from 9.5 percent in 1980 to 7.7 percent in 1982, but that there would be high and rising unemployment in each year (8.1 percent in 1980, 8.8 percent in 1981, and 9.3 percent in 1982). In this general sense, the forecast incorporated a Phillips curve but one that depended in a complex manner on initial conditions and shocks. However, the staff also forecasted unemployment and inflation under an alternative policy, 4.5 percent money growth. Hence, it is possible to use the difference in forecasts to gain a sharper sense of the tradeoff under alternative monetary policies. The forecast differences are listed as the row Δ in the table. The January 1980 FOMC meeting corresponded with the onset of a recession, as later dated by NBER researchers.

The second meeting is the January 1982 FOMC session, at which the benchmark strategy (again called strategy 1 in this meeting) was for a gradually declining path of money growth: 4.0 percent in 1982, 3.5 percent in 1983, and 3.0 percent in 1984. Again, the FRS staff forecasted declining inflation (from 6.4 in 1982 to 4.2 in 1984) but this time with declining unemployment (from 9.3 percent in 1982 to 8.9 percent in 1984). The perceived Phillips curve was less evident in these forecasts, but an opportunity to appraise its nature is afforded by the fact that the staff also prepared forecasts under the assumption of higher money growth (strategy 2).

There are a number of aspects of the benchmark policy projections that are notable. First, in each case, strategy 1 is the assumption under which the Federal Reserve staff made its "Green Book" forecast for inflation and real activity for the coming years. In both 1980 and 1982, inflation was expected to decline by about two percentage points under the benchmark forecast. Second, in both 1980 and 1982, the projections implied that a policy change (lowering money growth by 1.5 percent for three years) would have no effect on inflation within the first year. Third, looking out two years after such a policy change, the alternative Blue Book policy scenarios suggested substantial effects on both unemployment and inflation of changing monetary policy. A shift from strategy 1 to strategy 2 in 1980 was projected to produce a .9 percent decline in inflation in 1982 and a 2.3 percent increase in unemployment in 1982. Seen in terms of a "menu of policy choice," the unemployment cost in two years of a 1 percent reduction in inflation was 2.3 percent. A 1982 shift from strategy 1 to strategy 2 was predicted to have the same effect on inflation at a two-year horizon, at an unemployment of 2.0 percent.

Proceeding further, we can set a *lower bound* on the perceived unemployment cost of disinflation by cumulating the "deltas" over the three-year period and viewing the result as the first part of a transition to a 1.5 percent lower inflation rate.[28] From that standpoint, in 1980 and 1982, the Fed's perception was that the first three years of restrictive monetary policy would cost about 4 point years of unemployment to lower the inflation rate by 1.5 percent. Thus, a 6.7 percent decline in the inflation rate was perceived to cost no less than 18

[28] It is a lower bound because the staff projection would certainly have viewed unemployment as remaining high beyond the three-year projection in the Blue Book; the unemployment Δ is highest in the third year.

point years of unemployment and it could have been a good bit higher once fourth-year costs were included.

Overall, in 1980 and 1982, it seems that the Fed's perception was that the disinflation would be at least twice as costly as the cyclical unemployment that was actually experienced. The perceived cost was not too much different across these years, although it was modestly smaller in 1982, and it was in line with the consensus estimates of Okun (1978). Accordingly, it does not seem that the Fed undertook the disinflation because its research staff, at least, believed that the costs would be small.

The Consolidation of Disinflation Gains

The reduction in inflation during the early 1980s had to be followed up by a lengthy period of inflation fighting, as discussed in Goodfriend (1993). In particular, upon taking over as chairman in 1987, Alan Greenspan had to fight a series of inflation scares. Yet, by 1994–1995 it seemed that the United States had settled into a period of low inflation (about 3 percent) and low unemployment (about 5 percent), essentially returning to conditions that resembled those in the mid-1950s, where we started our Phillips curve documentary.

6. IN THE AFTERMATH

The year 1996 saw two novel developments on the Phillips curve front, which are closing snapshots: the completion of version 1 of a new quarterly Federal Reserve Board macroeconometric model of the United States and an explicit discussion of Phillips curve tradeoffs by the FOMC.

The Model

The structure and results of large-scale models are notoriously difficult to convey in a compact and coherent manner. In that regard, the 1996 "Guide to FRB/US: A Macroeconomic Model of the United States," edited by Brayton and Tinsley, is a remarkable document. It provides the reader with a clear model-building vision and a set of clean experiments that can be used to learn about the model.

The wage-price block of the new model combines the sort of forward-looking price-setting and wage-setting specifications that are standard in modern macroeconomic analysis, with a set of gradual adjustment specifications of the variety that applied econometricians have found useful for fitting data since the days of Fisher, Klein, and Sargan. The specific modeling is in the tradition of the approach to time series econometrics initiated by Sargan and refined by Hendry and others.

For example, the price-level specification of the model contains a long-run relationship that makes

$$P_t^* = .98 * (W_t - a_t) + .02 * P_t^f - .003u_t,$$

so that the "equilibrium" price level strongly depends on the gap between the log nominal wage rate (W) and productivity (a) with weaker effects from the nominal energy fuel price (P_t^f), consistent with factor share data. Further, the unemployment rate has a small effect, via an effect on the desired markup (the units of measurement imply that a 1 percent increase in unemployment lowers the desired markup by .3 percent). The adjustment dynamics indicate that inflation is high as the price level adjusts gradually toward this target level, via

$$\begin{aligned} P_t - P_{t-1} &= .10(P*_{t-1} - P_{t-1}) + .57\text{lags}_2[P_{t-1} - P_{t-2}] \\ &\quad +.43\text{leads}_\infty[P_{t+1}^* - P_t^*]^e. \end{aligned}$$

That is, there is a gradual elimination of .1 of the gap ($P_{t-1}^* - P_{t-1}$) each quarter, some additional backward-looking adjustment terms with substantial weight, and variations in the expected target. The requirement that the structural lead and lag coefficients sum to one, along with similar restrictions in the companion wage equation, means that the FRB/US model features no long-run tradeoff between inflation and unemployment.

Thus, the new model represented a blend of the Klein-Sargan approach, with a new macroeconomic theory that stresses expectational elements of pricing and other behavior. The new FRB/US model also had common elements with a set of small, fully articulated dynamic models then being developed in academia (King and Wolman [1996] and Yun [1996]), which were early examples of the types of new macroeconomic models explored elsewhere in this *Economic Quarterly* issue.

The frictions in the model are substantial, as Brayton and Tinsley (1996) make clear, in that they apply to changes as well as levels. The associated distributed lags and leads are lengthy, averaging 3.3 quarters for unanticipated shocks. Hence, there is a short-run Phillips curve in the model that involves dynamics over many quarters. Figure 12 shows the response to a permanent decline in the inflation rate within the FRB/US model, essentially obtained by shifting down the constant term in an interest rate rule along Taylor (1993) lines.

The FRB/US model can be solved under alternative assumptions about expectation formation, with rational expectations being one specification and a modern version of adaptive expectations being the other. More precisely, the second specification is expectations based on a vector-autoregression estimated from a model's data for a subset of just three of that model's variables. The model implies a gradual disinflation process as a result of the lags

Figure 12 Credible Disinflation Dynamics in the FRB/US Model

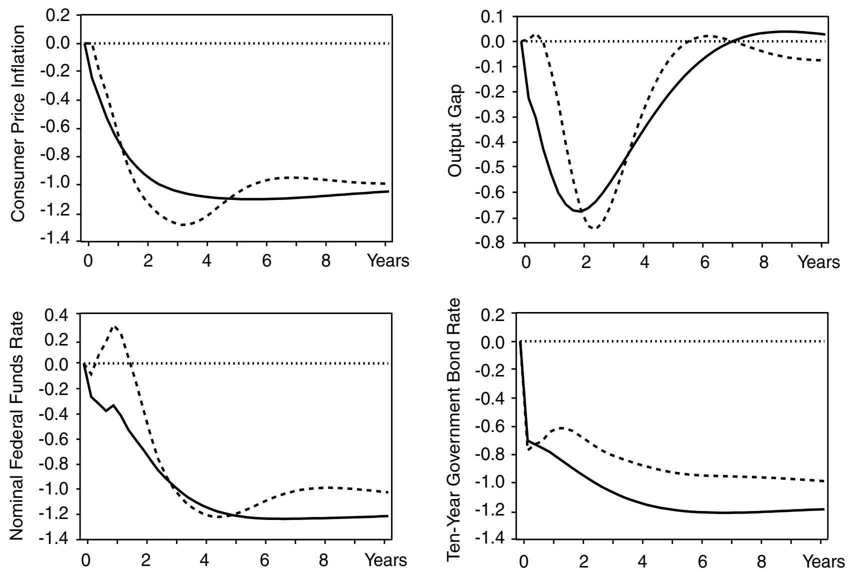

Notes: The FRB/US model can be used to calculate the implications of a permanent shift down in the target inflation rate, with results that are reported in Brayton and Tinsley (1996) and displayed above. The model can be solved under full rational expectations (solid line) or under a simpler procedure of vector-autoregression (VAR) expectations.

in the price and wage specification, but the transition to the new lower inflation rate is completed within about two years, although the real consequences are present for several years. Overall, since the reported simulations are in terms of an output gap, they cannot be directly compared to those considered above. Yet a back-of-the-envelope calculation suggests that they are quite major but smaller than those experienced in the Volcker period or in the MPS model discussed above. To see why, suppose that we summarize the figure as indicating that an upper bound on the output loss is a .5 percent output gap on average for 6 years; this is a 3 percent cumulative output loss. Suppose also that we use the same Okun's Law coefficient of 2.5 that links unemployment to output gaps as in the 1962 *Economic Report*. Then, the unemployment cost of a permanent disinflation is about 1.2 point years of unemployment for each point of inflation, a number that is in line with ratios reported by Brayton, Tinsley, and collaborators. The FRB/US team also reported that imperfect credibility of monetary policy actions can more than double the unemployment cost of a disinflation.

A notable feature of the disinflation simulation displayed in Figure 12 is that monetary policy initially works heavily through expectational channels. On impact, at date 0, the long-term bond rate drops dramatically (say, 80 basis points), while the federal funds rate moves by much less and only averages about a 40-basis-point decline over the first year. By contrast, the MPS model simulations of a lower money growth rate displayed nominal interest rate increases by an average of over 300 basis points for the first three years as shown in Table 1. Thus, the FRB/US model differs importantly from its MPS predecessor in terms of other areas, notably the term structure of interest rates, in ways that are important for monetary policy.

The Background to the Meeting

In 1996, the FOMC conducted a remarkable discussion of its long-run policy goals, stimulated by earlier calls for an increased emphasis on price stability by some of its members, as well as the adoption of inflation-targeting systems by other countries around the world. Notably, at a January 1995 meeting, Al Broaddus, the then-president of the Richmond Fed, had called within the FOMC for a system of inflation reports each year to accompany the Fed chairman's Humphrey-Hawkins testimony.[29] Broaddus' suggestion was opposed by FRB Governor Janet Yellen in January 1995, but the FOMC had agreed to continue the discussion in the context of future meetings that preceded the Humphrey-Hawkins testimony.

When the FOMC met in January 1996, the U.S. economy had been experiencing low inflation and strong macroeconomic activity for some time. In the first quarter of 1996, inflation was running at about 2 percent per year, with unemployment in the neighborhood of 5.5 percent. Since 1980, the United States had experienced the major decline in inflation described in the last section, during which unemployment had ranged over 10 percent in the last quarter of 1982 and the last two quarters of 1983. In the last year of Volcker's chairmanship and during the first few years of Greenspan's, a rise in inflation had taken place—from the 2 percent range in 1986 to about 4 percent in 1990—which had been accompanied by a decline in unemployment.[30] Subsequently, during 1991 and 1992, there had been a rise in unemployment while inflation fell back in the 2 percent range. Most recently, from mid-1992 through the end

[29] Broaddus (2004) used his proposals at this meeting as one of three examples of his use of macroeconomic principles in practical monetary policy discussion.

[30] At the time of Greenspan's appointment in August of 1987, inflation was at 2.8 percent, while unemployment was 5.8 percent.

of 1995, there had been about 2 percent inflation, while unemployment was between 5.5 and 6 percent.

These developments are shown in Figure 13, which is a Phillips-style plot of unemployment and inflation during 1980 through 1996. Observations during the Volcker period are marked with a circle (o) and those during the Greenspan period are marked with a diamond ($\diamond$). This figure captures the background to the FOMC's 1996 discussion. The major disinflation is the first half-loop: an interval of declining inflation and rising unemployment between 1980 and mid-1983, followed by an interval of declining inflation and declining unemployment with inflation reaching the 4 percent range by the second quarter of 1984. Subsequently, there was a year in which inflation fell in the 2 percent range, with little accompanying change in unemployment. The increase in inflation between mid-1985 through 1989 was followed by a decline in inflation to the 2 percent range during late 1991 and early 1993, accompanied by increases in unemployment. In late 1993 through early 1995, unemployment fell sharply, with little change in inflation. Thus, the late-Volcker and early-Greenspan years trace out a full clockwise loop, after the disinflation of 1980–1984. The negative association between inflation and unemployment during the first stages of each of these three episodes (one of increasing inflation and one of decreasing inflation) corresponds to periods that FOMC members would have viewed as reflecting the phenomena isolated by Phillips.

The Meeting

At the time of the January 1996 meeting, there were two important economic conditions that occupied the FOMC's attention. First, there was a sense that key aspects of the U.S. economy were changing, with the possibility of a "New Economy" based on computer and communications advances. Second, and most important for the meeting, the inflation rate for personal consumption expenditures was running at about a 3 percent rate and its "core" component—that stripped of food, energy, and other volatile price components—was running at about 2.5 percent, but staff forecasts suggested that it was poised to rise to the 3 percent range as well. The strong real growth in the economy, coupled with a decline in unemployment to the range of 5.5 percent, had led some FOMC members to express concerns about inflation.

In detailed prepared remarks, Governor Janet Yellen discussed a cost-benefit approach to determining the optimal long-run rate of inflation and the transition path. She noted that the Board's new model indicated a cost of 2.5 point years of unemployment for every 1 percent decline in the long-run inflation rate, under imperfect credibility. To warrant a reduction in inflation, she argued that such a cost of permanently lower inflation had to be less than the discounted value of a stream of future benefits. However, Yellen also

Figure 13 Inflation and Unemployment, 1978Q4–1996Q4

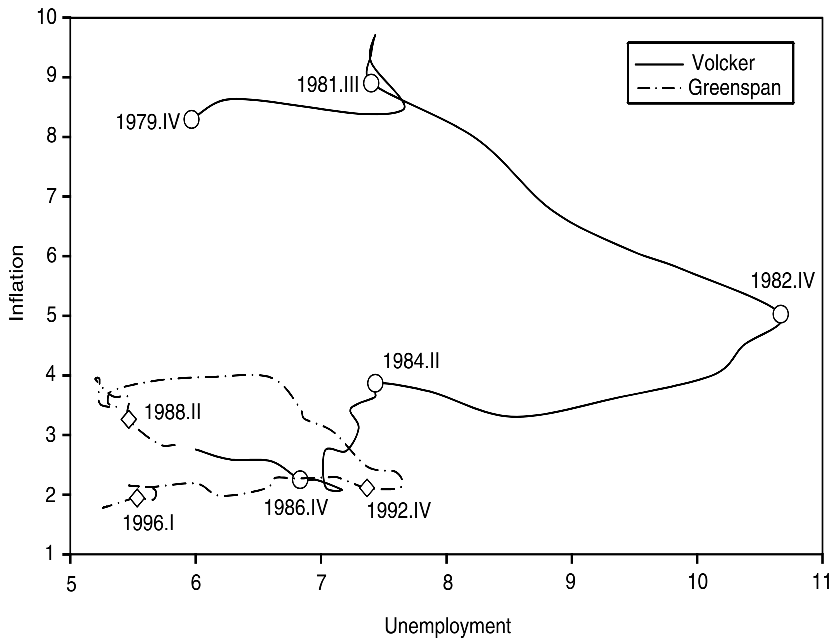

Notes: The inflation rate is the year-over-year change in the gross domestic product (GDP) deflator, the unemployment rate is the civilian unemployment rate, quarterly averages of monthly figures. All data are from the Federal Reserve Economic Database (FRED) at the Federal Reserve Bank of St. Louis.

returned to a theme suggested by Phillips' original research, which was that there could be particular costs to low rates of inflation. Citing research by Akerlof, Dickens, and Perry (1996) which argued that worker-resistance to nominal pay cuts produced a long-run Phillips curve with a negative slope at low rates of inflation, Yellen also argued for a positive rate of long-run inflation "to grease the wheels of the labor market."

As they considered the appropriate long-run rate of inflation, the FOMC decisionmakers took into account their perceived transition costs, their sense of the benefits from permanently low inflation, and their sense of the costs of permanently low inflation. There was diversity in the views reflected in the statements of various members on each of these topics. But, as Broaddus noted, there was a consensus that the long-run inflation rate should not be higher than the current level of 3 percent. Broaddus and then-Cleveland Fed President Jerry Jordan stressed the importance of explicit public discussion of inflation objectives as a means of enhancing Fed credibility and thus lowering the cost of further reductions in inflation.

The FOMC discussed how to define price stability as an objective of monetary policy. Greenspan suggested that "price stability is that state in which expected changes in the general price level do not effectively alter business or household decisions," but Yellen challenged him to translate that general statement into a specific numerical value. He responded that "the number is zero, if inflation is properly measured." Yellen said that she preferred 2 percent "imperfectly measured."

The FOMC settled on 2 percent inflation as an interim goal, with a policy of deliberately moving toward that lower level. Presumably, some members viewed it as the natural first step toward a lower ultimate inflation objective, while others thought of it as an end point. On the second day of the two-day meeting, Greenspan cautioned the committee that the 2 percent objective was included within "the highly confidential nature of what we talk about at an FOMC meeting." He noted that "the discussion we had yesterday was exceptionally interesting and important" but warned that "if the 2 percent inflation figure gets out of this room, it is going to create more problems for us than I think any of you might anticipate." He did not elaborate on whether he was concerned about market or political reactions to the inflation goal.

7. SUMMARY AND CONCLUSIONS

With a series of snapshots over a nearly 40-year period, this article has reviewed the evolution of the Phillips curve in macroeconomic policy analysis in the United States. During this period, U.S. inflation rose dramatically, initially during a decade of glittering economic performance and then further during an interval of stagflation. The reversal of inflation beginning in the early 1980s was associated with a major recession, although perhaps not as large a one as policymakers and economists had feared.

The rise and fall of inflation brought about a major change in the style of macroeconometric models that were used to evaluate policy choices. The earliest versions of these models featured a substantial long-run tradeoff consistent with the findings of Phillips over a near-century of U.K. data. The subsequent evolution of models first involved altering their wage-price block so that there was no long-run tradeoff and then, later, a more comprehensive rational expectations revision that included forward-looking wage and price-setting structured so that there was no long-run tradeoff.

More generally, the rise and fall of inflation led monetary policymakers to place greater weight on the role of expectations in governing macroeconomic activity, with central banks working to extract information in long-term interest rates about market expectations of inflation. Toward the end of the historical period examined here, the Federal Reserve System had decided to maintain a goal of a low, but positive rate of inflation. Yet, it also chose not to communicate that long-run target directly to the public. The decision to choose a positive rate of inflation was traced, in part, to a concern about the transitory unemployment costs of moving to a zero rate of inflation and in part to a concern about high long-run costs of low inflation, in the spirit of Phillips' analysis.

REFERENCES

Akerlof, George A., William R. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economics Activity* 1: 1–76.

Brayton, Flint, and Peter C. Tinsley. 1996. "A Guide to FRB/US: A Macroeconomic Model of the United States." FEDS Working Paper 96-42.

Broaddus, J. Alfred, Jr. 2004. "Macroeconomic Principles and Monetary Policy." Federal Reserve Bank of Richmond *Economic Quarterly* 90 (Winter): 1–9.

Burns, Arthur F., and Wesley C. Mitchell. 1946. Measuring Business Cycles. New York: National Bureau of Economic Research.

Castle, Jennifer, and David F. Hendry. Forthcoming."The Long-Run Determinants of UK Wages, 1860–2004." *Journal of Macroeconomics: Advances in Historical Analyses*.

Christ, Carl F. 1973. "The 1973 Report of the President's Council of Economic Advisers: A Review." *American Economic Review* 63 (September): 515–26.

de Menil, G., and Enzler, J. 1972. "Prices and Wages in the FR-MIT-PENN Econometric Model." In *The Econometrics of Price Determination*, edited by Otto Eckstein. Washington, D.C.: Federal Reserve Board of Governors.

Duesenberry, James S., Otto Eckstein, and Gary Fromm. 1960. "A Simulation of the United States Economy in Recession." *Econometrica* 28 (October): 749–809.

Federal Reserve Bank of New York. 1998. "The FOMC Meeting: Developing a Policy Directive." In *U.S. Monetary Policy and Financial Markets*. Available at www.newyorkfed.org/education/addpub/monpol/.

Fellner, William. 1976. *Towards a Reconstruction of Macroeconomics*. Washington, D.C.: American Enterprise Institute.

Fisher, Irving. 1926. "A Statistical Relationship between Unemployment and Price Changes." *International Labor Review* 47 (June): 785–92.

Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review* 58 (January): 1–17.

Goodfriend, Marvin. 1993. "Interest Rate Policy and the Inflation Scare Problem: 1979–1992." Federal Reserve Bank of Richmond *Economic Quarterly* 91 (Winter): 1–23.

Goodfriend, Marvin, and Robert G. King. 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics* 52 (July): 981–1015.

Gordon, Robert J. 1970. "The Recent Acceleration of Inflation and Its Lessons for the Future." *Brookings Papers on Economic Activity* 1: 8–47.

Hetzel, Robert. 1998. "Arthur Burns and Inflation." Federal Reserve Bank of Richmond *Economic Quarterly* 84 (Winter): 21–44.

Hetzel, Robert. 2008. *The Monetary Policy of the Federal Reserve: A History*. Cambridge, U.K.: Cambridge University Press.

Hymans, Saul H. 1972. "Prices and Price Behavior in Three U.S. Econometric Models." In *The Econometrics of Price Determination*, edited by Otto Eckstein. Washington, D.C.: Federal Reserve Board of Governors, 309–24.

King, Robert G., and Alexander L. Wolman. 1996. "Inflation Targeting in a St. Louis Model of the 21st Century." Federal Reserve Bank of St. Louis *Review* 78 (Fall): 83–107.

Klein, L. R., R. J. Ball, A. Hazelwood, and P. Vandome. 1961. *An Econometric Model of the United Kingdom*. Oxford: Basil Blackwell.

Lucas, Robert E., Jr. 1972a. "Econometric Testing of the Natural Rate Hypothesis." In *The Econometrics of Price Determination*, edited by Otto Eckstein. Washington, D.C.: Federal Reserve Board of Governors, 50–9.

Lucas, Robert E., Jr. 1972b. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (April): 103–24.

Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1 (January): 19–46.

Mankiw, N. Gregory. 2002. *Macroeconomics*, fifth edition. New York: Worth Publishers.

McCallum, Bennett T. 1984. "Credibility and Monetary Policy." In *Proceedings*. Federal Reserve Bank of Kansas City, 105–35.

McCallum, Bennett T. 1994. “Identification of Inflation-Unemployment Tradeoffs in the 1970s a Comment.” *Carnegie-Rochester Conference Series on Public Policy* 41 (December): 231–41.

Meltzer, Allan M. 1980. “The Case for Gradualism in Policies to Reduce Inflation.” Federal Reserve Bank of St. Louis *Review* April: 127–54.

Modigliani, Franco, and Lucas Papademos. 1975. “Targets for Monetary Policy in the Coming Year.” *Brookings Papers on Economic Activity* 6: 141–65.

Muth, John F. 1960. “Optimal Properties of Exponentially Weighted Forecasts.” *Journal of the American Statistical Association* 55 (June): 299–306.

Muth, John F. 1961. “Rational Expectations and the Theory of Price Movements.” *Econometrica* 29 (July): 315–35.

Okun, Arthur M. 1962. “Potential GNP: Its Measurement and Significance.” In *Journal of the American Statistical Association*; reprinted with slight changes in *The Political Economy of Prosperity* by Arthur M. Okun. Washington, D.C.: Brookings Institution.

Okun, Arthur M. 1978. “Efficient Disinflationary Policies.” *American Economic Review* 68 (May): 348–52.

Orphanides, Athanasios. 2003. “The Quest for Prosperity without Inflation.” *Journal of Monetary Economics* 50 (April): 633–63.

Phelps, Edmund S. 1967. “Phillips Curves, Expectations of Inflation, and Optimal Inflation over Time.” *Economica* 34 (August): 254–81.

Phillips, A. W. 1958. “The Relationship between Unemployment and the Rate of Change of Money Wages in the United Kingdom 1861–1957.” *Economica* 25 (November): 283–99.

Poole, William. 1970. “Gradualism: A Mid-Course View.” *Brookings Papers on Economic Activity* 1: 271–301.

Primiceri, Giorgio. 2006. “Why Inflation Rose and Fell: Policymakers’ Beliefs and U.S. Postwar Stabilization Policy.” *Quarterly Journal of Economics* 121 (August): 867–901.

Samuelson, Paul A., and Robert M. Solow. 1961. “Analytical Foundations of Anti-Inflation Policy.” *American Economic Review* 50 (May): 177–94.

Sargan, John D. 1964. “Wages and Prices in the United Kingdom: A Study in Econometric Methodology.” In *Econometric Analysis for National Economic Planning*, edited by P. E. Hart, G. Mills, and J. K. Whitaker. London: Butterworth and Co. Ltd., 25–54; Reprinted as chapter 8 in *Contributions to Econometrics: John Denis Sargan*, edited by Esfandiar Maasoumi, Cambridge University Press.

Sargent, Thomas J. 1971. "A Note on the Accelerationist Controversy." *Journal of Money, Credit and Banking* 3 (August): 721–25.

Sargent, Thomas J. 1981. "The Ends of Four Big Inflations." Federal Reserve Bank of Minneapolis Working Paper 158 (May).

Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton, N.J.: Princeton University Press.

Solow, Robert M. 1969. *Price Expectations and the Behavior of the Price Level*. Manchester, U.K.: Manchester University Press.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment and Monetary Policy." *The Journal of Economic Perspectives* 11 (Winter): 33–49.

Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." Carnegie-Rochester Conference Series on Public Policy 39 (November): 195–214.

Tobin, James. 1980. "Stabilization Policy Ten Years After." *Brookings Papers on Economic Activity* 11:19–78.

Yun, Tack. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37 (April): 345–70.

Economic Quarterly—Volume 94, Number 4—Fall 2008—Pages 361–395

The New Keynesian Phillips Curve: Lessons From Single-Equation Econometric Estimation

James M. Nason and Gregor W. Smith

The last decade has seen a renewed interest in the Phillips curve that might be an odd awakening for a macroeconomic Rip van Winkle from the 1980s or even the 1990s. Wasn't the Phillips curve tradition discredited by the oil prices shocks of the 1970s or by theoretical critiques of Friedman, Phelps, Lucas, and Sargent? It turns out that the New Keynesian Phillips curve (NKPC) is consistent with both the theoretical demands of modern macroeconomics and some key statistical properties of inflation. In fact, the NKPC can take a sufficient number of guises to accommodate a wide range of perspectives on inflation.

The NKPC originated in descriptions of price setting by firms that possess market power. For example, Rotemberg (1982) describes how a monopolist sets prices if it faces a cost of adjustment that rises with the scale of the price change. He shows that prices then gradually track a target price and also depend on expected, future price targets. Calvo (1983) instead describes firms that are monopolistic competitors. They change their prices periodically. Knowing that some time may pass before they next set prices, firms anticipate future cost and demand conditions, as well as current ones, in setting their price. Also, the staggering or nonsynchronization of price setting by firms

James M. Nason is a policy advisor and research economist at the Federal Reserve Bank of Atlanta. Gregor W. Smith is a professor of economics at Queen's University in Canada. The authors thank the Social Sciences and Humanities Research Council of Canada and the Bank of Canada Research Fellowship program for support of this research. The views in this paper represent those of the authors and not the Bank of Canada, the Federal Reserve Bank of Richmond, the Federal Reserve Bank of Atlanta, the Federal Reserve System, or any of the staff. The authors thank Nikolay Gospodinov for the use of his GMM code; Ellis Tallman, Ricardo Nunes, and Frank Schorfheide for helpful suggestions; and Andreas Hornstein and Thomas Lubik for detailed comments that greatly improved this paper.

creates an aggregate stickiness: The aggregate price level will react only partially on impact to an economy-wide shock, such as an unexpected change in monetary policy.

These theoretical models link prices to a targeted real variable such as a markup on the costs faced by the price-setting firm. Therefore, they also relate the change in prices over time (i.e., the inflation rate) to real variables. So it is natural to label them as Phillips curves. In fact, there are a range of setups called NKPCs that vary depending on (a) the information and price-setting behavior attributed to firms and (b) the measure of costs or demand that firms are assumed to target. Whether a specific version of the NKPC fits inflation data has implications for our understanding of recent macroeconomic history and for the design of good policy. For example, the parameters of the NKPC influence how monetary policy ideally should respond to external shocks. Schmitt-Grohé and Uribe, in this issue, make this connection clear. (King, also in this issue, describes the uses of an older Phillips curve tradition in policymaking.)

Yet, putting the NKPC to use for policy analysis requires that it has a good econometric track record in describing actual inflation dynamics. In this article we review this record using single-equation statistical methods that study the NKPC on its own. These methods stand in contrast to approaches that place the NKPC in larger economic models, sometimes referred to as systems methods, which are reviewed by Schorfheide in this issue. A disadvantage of single-equation methods is that they do not make use of everything known about the economy (e.g., the monetary policy regime), so they generally do not provide the greatest statistical precision. Their advantage is that they allow us to be agnostic about the rest of the economy, and so their findings remain valid and will not be affected by misspecification of other parts of a larger macroeconomic model.

This article asks the following questions: How can we estimate the NKPC and what do we find when we do so for the United States? Are its parameters stable over time and well-identified? Is there a relation between inflation and real activity? Do we reach similar conclusions about the NKPC regardless of the way in which we measure inflation, forecast future inflation, or model the costs or output gap that inflation tracks?

We focus on marginal cost as the real activity variable in the NKPC. We find that the single-equation statistical evidence for this relationship is mixed. Since 1955 there does seem to be a stable NKPC for the United States with positive parameter values as we would expect from economic theory. But our confidence intervals for these parameter values are somewhat wide, the findings depend on how we model expected future inflation, and further research is needed on the best way to represent the marginal cost variable to which price changes react. Before outlining the methods and findings, though, we begin by introducing the specific NKPC that we will estimate and the inflation history it aims to explain.

1. FOUNDATIONS AND INTERPRETATIONS

The New Keynesian Phillips curve arises from a description of staggered price setting, which is then linearized for ease of study. The result is an equation in which the inflation rate, π_t, depends on the expected inflation rate next period, $E_t\pi_{t+1}$, and a measure of marginal costs, denoted x_t:

$$\pi_t = \gamma_f E_t \pi_{t+1} + \lambda x_t.$$

Iterating this NKPC difference equation forward gives inflation as the present value of future marginal costs:

$$\pi_t = \lambda \sum_{i=0}^{\infty} \gamma_f^i E_t x_{t+i}.$$

This present-value relation shows that firms consider both their current marginal costs, x_t, and their expectations or forecasts of future costs when adjusting prices.

Lacker and Weinberg (2007) describe the history and derivation of this key building block of New Keynesian macroeconomic models. Dennis (2007) outlines a range of environments that can underpin a hybrid NKPC. Calvo's (1983) specific price-setting model is only one of several possible microfoundations for the NKPC. In a Calvo-based NKPC, a fraction, θ, of firms cannot change prices in a given period. Firms also have a discount factor, β. The reduced-form parameters of the NKPC, γ_f and λ, are related to these two underlying pricing parameters according to

$$\gamma_f = \beta, \quad \lambda = \frac{(1-\theta)(1-\beta\theta)}{\theta}.$$

Because β is a discount factor, both it and γ_f must range between zero and one. The same holds for θ because it represents the fraction of firms unable to move prices at any moment.

Many estimates of the NKPC find that lagged inflation helps to explain current inflation. We report much the same in this article. This has suggested to some economists that a better fit to inflation history can be obtained with this equation:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \pi_{t+1} + \lambda x_t,$$

which Galí and Gertler (1999) call the *hybrid* NKPC. They develop their NKPC by modifying Calvo's (1983) description of price-setting decisions. In this case, a fraction, ω, of firms can change prices, but do not choose this option. Define $\phi = \theta + \omega[1 - \theta(1 - \beta)]$. Then the mapping between these structural parameters and the reduced-form parameters is

$$\gamma_b = \frac{\omega}{\phi}, \quad \gamma_f = \frac{\beta\theta}{\phi}, \quad \lambda = \frac{(1-\omega)(1-\theta)(1-\beta\theta)}{\phi}.$$

Galí and Gertler (1999) note that this mapping between the structural, price-setting parameters ω, θ, and β and the reduced-form hybrid NKPC parameters γ_b, γ_f, and λ is unique if the former set of parameters lie between zero and one.

This form of the NKPC also is consistent with the incomplete indexing model of Woodford (2003). He assumes that those firms that cannot optimally alter their prices instead index to a fraction of the lagged inflation rate. This feature makes current inflation depend on lagged inflation, and so provides an alternative interpretation of the hybrid NKPC. Christiano, Eichenbaum, and Evans (2005), among others, study the implications of full indexation, which is equivalent to the restriction $\gamma_b + \gamma_f = 1$ in the Galí-Gertler hybrid NKPC.

Like the original NKPC, the hybrid version can be rewritten in present-value form:

$$\pi_t = \delta_1 \pi_{t-1} + \left(\frac{\lambda}{\delta_2 \gamma_f}\right) \sum_{k=0}^{\infty} \left(\frac{1}{\delta_2}\right)^k \mathrm{E}_t x_{t+k},$$

where δ_1 and δ_2 are stable and unstable roots, respectively, of the characteristic equation in the lag operator L:

$$-L^{-1} + \frac{1}{\gamma_f} - \frac{\gamma_b L}{\gamma_f} = 0.$$

The present-value version of the hybrid NKPC shows that inflation persistence can arise from the influence of lagged inflation or the slow evolution of the present value of marginal costs.

The NKPC is often derived by log-linearizing a typical firm's price-setting rule around a mean zero inflation rate. Ascari (2004) shows that non-zero mean inflation can affect the response of inflation to current and future marginal cost in the NKPC. However, we follow much of the empirical NKPC literature and demean the data.

Cogley and Sbordone (forthcoming) build on Ascari's work, among others, by log-linearizing the NKPC around time-varying trend inflation. This procedure assumes that inflation is nonstationary to obtain a NKPC with time-varying coefficients even though the underlying Calvo-pricing parameters are constant. The resulting NKPC is purely forward-looking, assigning no role to lagged inflation. Cogley and Sbordone estimate the structural coefficients of their NKPC using a vector autoregression. Hornstein (2007) assesses the implications of this approach for the stability of the NKPC. Since these studies use system estimators, we omit them from our review.

In sum, the hybrid NKPC is consistent with various pricing or information schemes. This suggests a focus on the reduced-form coefficients λ, γ_b, and

γ_f (rather than the structural price-setting parameters ω, θ, and β), which is our emphasis in this article. We also use single-equation estimators to explore the fit of the hybrid NKPC to U.S. data. After all, obtaining a good fit for the hybrid NKPC is a necessary first step in attributing a monetary transmission mechanism to staggered price setting by firms.

Economists have not yet reached a consensus on two key questions concerning the NKPC parameters. First, what is the mixture of forward (γ_f) and backward (γ_b) weights? If γ_f is large, events in the future (including changes in monetary policy) can influence the current inflation rate. If, instead, γ_b is large, inflation has considerable inertia independent of any slow movements in the cost variable. Such inertia affects the design of monetary policy (again, see Uribe and Schmitt-Grohé). Woodford (2007) reviews several explanations for inflation inertia and also discusses whether it could be stable over time.

Second, can we identify a significant $\hat{\lambda}$, the coefficient on marginal costs? In this case, identification simply refers to measuring a partial correlation coefficient in historical data rather than the possibility of misspecification (i.e., whether this coefficient necessarily measures the theoretical parameter studied in New Keynesian models.) Finding a significant value is a *sine qua non* for empirical work with the NKPC. If we cannot find a way to represent a price-setting target, we cannot hope to identify the adjustment process of inflation to its target. Consequently, much of the research on estimating the NKPC involves exploring the x-variable or how to measure marginal costs.

Before looking at formal econometric methods, let us look at the data. Figure 1 plots the U.S. inflation rate (the black line) and a measure of marginal cost (the gray line) from 1955:1 to 2007:4. We measure the inflation rate, π_t, as the quarter-to-quarter, annualized growth rate in the U.S. implicit GDP deflator (`GDPDEF` from `FRED` at the Federal Reserve Bank of St. Louis). Marginal cost, x_t, is the update of the series on real unit labor costs used by Galí and Gertler (1999) and Sbordone (2002). It is given by 1.0765 times the logarithm of nominal unit labor costs in the nonfarm business sector (the ratio of `COMPNFB` to `OPHNFB` from `FRED`) divided by the implicit GDP deflator. Both series have fallen since 1980, which in itself provides some statistical support for the idea that inflation tracks marginal cost. There also are some obvious divergences, for example around 2000–2001. However, it is possible that these occurred because inflation was tracking expected future marginal cost (as in the present-value model) or because it was linked to lagged inflation (as in the hybrid NKPC). We next describe the statistical tools economists have used to see if either of these explanations fits the facts.

2. ESTIMATION

The fundamental challenge with estimating the parameters of the hybrid NKPC is that expected inflation cannot be directly observed. The most popular

Figure 1 U.S. Inflation and Marginal Cost

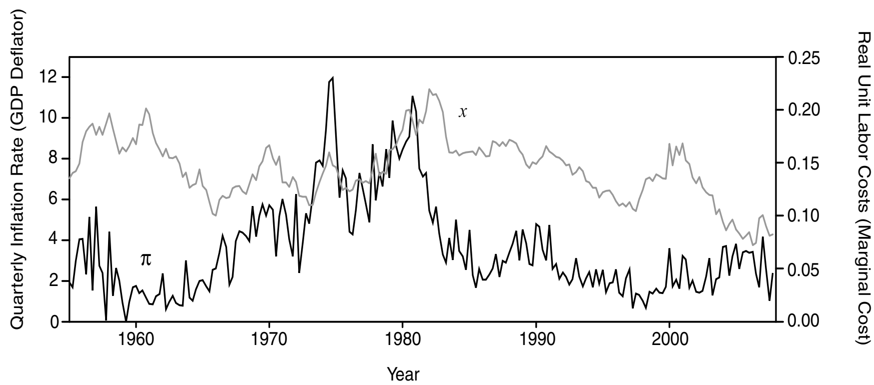

econometric method for dealing with this issue begins from the properties of forecasts. To see how this works, let us label the information used by price-setters to forecast inflation by I_t, so their forecast is $\mathrm{E}[\pi_{t+1}|I_t]$. Economists do not observe this forecast, but it enters the NKPC and influences the current inflation rate. Denote by $\mathrm{E}[\pi_{t+1}|z_t]$ an econometric forecast that uses some variables, z_t, to predict next period's inflation rate, π_{t+1}. Suppose that z_t is a subset of the information available to price setters. To construct our econometric forecast, we simply regress actual inflation on our set of variables (sometimes called instruments), z_t, like this:

$$\pi_{t+1} = bz_t + \epsilon_{t+1},$$

so that our forecast is simply the fitted value

$$\mathrm{E}[\pi_{t+1}|z_t] = \hat{b}z_t.$$

By construction it is uncorrelated with the residual term $\hat{\epsilon}_{t+1}$.

A key principle of forecasts (or rational expectations) is the law of iterated expectations. According to that law, our econometric prediction of price-setters' forecast of inflation is simply our forecast. Symbolically,

$$\mathrm{E}\left[\mathrm{E}[\pi_{t+1}|I_t]|z_t\right] = \mathrm{E}[\pi_{t+1}|z_t].$$

The idea is that our effort to predict what someone with better information will forecast simply gives us our own best forecast. With the law of iterated expectations in hand, we can also imagine regressing the unknown forecast on z_t:

$$\begin{aligned}
\mathrm{E}\left[\pi_{t+1} \mid I_t\right] &= \mathrm{E}\left[\pi_{t+1} \mid I_t \mid z_t\right] + \eta_t \\
&= \mathrm{E}\left[\pi_{t+1} \mid z_t\right] + \eta_t,
\end{aligned}$$

in which the residual, η_t, also is uncorrelated with the econometric forecast. The econometric forecast does not use all the information available to price setters when they construct forecasts, so it does not capture all the variation in their forecasts. Put differently, the unobserved, economic forecast has some added variation that appears in η_t.

With this statistical reasoning behind us, the hybrid NKPC can be rewritten:

$$\begin{aligned}
\pi_t &= \gamma_b \pi_{t-1} + \gamma_f \mathrm{E}\left[\pi_{t+1} \mid I_t\right] + \lambda x_t \\
&= \gamma_b \pi_{t-1} + \gamma_f \left(\mathrm{E}\left[\pi_{t+1} \mid z_t\right] + \eta_t\right) + \lambda x_t \\
&= \gamma_b \pi_{t-1} + \gamma_f \mathrm{E}\left[\pi_{t+1} \mid z_t\right] + \lambda x_t + \gamma_f \eta_t.
\end{aligned}$$

This is an econometric equation that can be used to estimate the parameters by least squares, for we now have measurements of the three variables on the right-hand side of the equation. In fact, this two-step procedure—forecast using predictors z_t, then substitute and apply least squares—is just two-stage least squares, familiar from econometrics textbooks. Provided that we include the other hybrid NKPC explanatory variables, π_{t-1} and x_t, in the list of first-stage regressors, the error term will be uncorrelated with them, too, and so least-squares will be valid in the second stage. (In contrast, simply estimating the NKPC by least squares, using π_{t+1} in place of $\mathrm{E}[\pi_{t+1}]$, yields inconsistent estimates of the parameters, because π_{t+1} is correlated with the residual, η_t.)

Two-stage least squares, in turn, is a special case of a method known as generalized instrumental variables, or generalized method-of-moments (GMM) estimation. To see how this works, take the hybrid NKPC and write it as follows:

$$\pi_t - \gamma_b \pi_{t-1} - \gamma_f \mathrm{E}[\pi_{t+1} | I_t] - \lambda x_t = 0.$$

Then imagine forecasting this entire combination of variables:

$$\begin{aligned}
\mathrm{E}\left[\pi_t - \gamma_b \pi_{t-1} - \gamma_f \mathrm{E}[\pi_{t+1} | I_t] - \lambda x_t | z_t\right] &= \\
\mathrm{E}\left[\pi_t - \gamma_b \pi_{t-1} - \gamma_f \pi_{t+1} - \lambda x_t | z_t\right] &= 0,
\end{aligned}$$

where we again have used the law of iterated expectations to replace the unobserved market forecast with our own econometric one. The last part of this equation is the basis for numerous studies of the NKPC. Simply put, there should be no predictable departures from the inflation dynamics implied by the hybrid NKPC or, equivalently, the residuals should have a mean of zero and be uncorrelated with predictor variables z_t. These properties allow for a diagnostic test of the NKPC.

Finding estimates of the hybrid NKPC parameters proceeds as follows. We collect data on inflation and marginal costs. Then we make a list of widely available information, z_t, that could include current and lagged values of these same variables, as well as other macroeconomic indicators such as interest rates. These instrumental variables (also known simply as instruments), z_t, must have two key properties. First, they must be at least as numerous as the parameters of the model (which here number three) and provide independent sources of variation (i.e., they cannot be perfectly correlated with one another). Intuitively, to measure three effects on inflation there must be at least three independent pieces of information or exogenous variables. Second, they must be uncorrelated with forecast errors that appear as residuals in the hybrid NKPC. Instruments with these properties are called valid.

Next, we use some econometric software to adjust the economic parameters $\{\gamma_b, \gamma_f, \lambda\}$ so that departures from the hybrid NKPC are uncorrelated with z_t, and so are unpredictable. The criterion guiding the adjustment is that the moment conditions that consist of the cross-products of the NKPC residuals with the instruments should be as close to zero as possible. Whenever we have at least three valid instruments in the set z_t, we can identify and solve for values of the three hybrid NKPC parameters using this criterion. In practice, the algorithm attempts this task by squaring the deviations of the moment conditions from zero and then minimizing the weighted square of this list of deviations. Cochrane (2001, chapters 10–11) provides a lucid introduction to GMM.

We make brief technical digressions on two details of GMM estimation. First, the distance of moment conditions (i.e., forecasts of departures from the hybrid NKPC) from zero is measured relative to their sampling variability, just as with any statistic. For GMM this involves calculating a heteroskedasticity-and-autocorrelation-consistent (HAC) covariance matrix. This article employs either a Newey and West (1994) or an Andrews (1991) quadratic-spectral HAC estimator with automatic lag-length selection. Second, some authors note that the way the hybrid NKPC is written matters for its estimation. For example, multiplying it by the Calvo parameter, ϕ, might seem to make it easier to estimate ϕ as the weight on π_t and ω as the weight on π_{t-1}. In this article, we use the continuously updated (CU-)GMM estimator of Hansen, Heaton, and Yaron (1996), from which estimates of the hybrid NKPC are independent of any normalization applied to it.

There are many macroeconomic indicators that could be included in z_t. We need to place at least three macroeconomic variables in z_t so that the three parameters of the hybrid NKPC are just-identified, in the jargon of econometrics. The parameters are said to be overidentified when four or more macroeconomic variables are included in z_t. It turns out that this possibility provides a test of the validity of the NKPC. According to econometric theory, any instrument set should yield the same coefficients except for some random sampling error. So by estimating with various sets of instruments and comparing the findings (or seeing if the NKPC departures are close to zero even when we use a long list of instruments) we can test whether the NKPC really holds or not. This diagnostic procedure is called a test of overidentifying restrictions. Informally, we refer to an NKPC that passes this test as fitting the data.

In practice, most researchers have used lagged macroeconomic variables as instruments. To see why, recall that an error term, $\gamma_f \eta_t$, arises in the estimating equation. Recall, however, that by including $\mathrm{E}[\pi_{t+1}|z_t]$ we in fact are trying to represent the regressor $\mathrm{E}[\pi_{t+1}|I_t]$ (a forecast of inflation made in the current period), not π_{t+1}. So one can think of the econometric equation as containing an error term dated t that reflects the difference between these two measures. Moreover, some economists have argued that there are unobserved cost shocks (components of x_t) that also can underlie an error term in the NKPC. Recall that a key property of instruments is that they be uncorrelated with the error term. Many researchers studying the NKPC, therefore, have used only lagged variables as instruments, labelled z_{t-1}, to try to ensure that this property holds.

Another way to think of this approach is that using instrumental variables is a classic way of dealing with the problem of a regressor that is subject to measurement error. If the marginal cost series, x_t, is measured with error, then including x_t as an instrument will lead to the attenuation bias (bias toward zero) familiar in this errors-in-variables problem. Using lagged instruments can avoid this bias.

Next, we present examples of GMM estimation. Instruments include lagged values of inflation and marginal costs. In addition, we also present results with a longer list of instruments. This list includes the term spread of the five-year Treasury bond over the 90-day Treasury bill, ts, which is a natural candidate for forecasting inflation. Table 1 gives the complete list of variables and instruments, with the symbols used to represent them and the sources for these data. Different econometricians might measure the output gap, y_t, differently. We used linear detrending of real per capita GDP to produce the output gap as an instrument, but the findings are very similar if we use other possible measures of the output gap as an instrument instead. The list of instruments is inspired by Galí and Gertler (1999). They use four lags of this list of six variables for their quarterly 1960–1997 sample. We adopt the same list, though updated to 2007, so that the reader can compare our findings to theirs.

Table 2 reports CU-GMM estimates of the hybrid NKPC on the 1955:1–2007:4 sample. The first column lists the instrument set, z_t. The next columns list estimates of the reduced-form parameters $\hat{\gamma}_b$, $\hat{\gamma}_f$, and $\hat{\lambda}$ over their standard errors, followed by the structural Calvo-pricing estimates $\hat{\omega}$, $\hat{\theta}$, and $\hat{\beta}$ over their standard errors. The final column gives a test statistic, denoted J, for the

Table 1 Measuring Variables and Instruments

Label	Definition	Code/Source
π	Inflation rate, implicit GDP deflator	GDPDEF/FRED
x	Log labor share of income	1.0765 ln[COMPNFB/OPHNFB] – ln[GDPDEF]/FRED
y	Linearly detrended Log per capita real GDP	ln[GDPC96/CNP16OV]/FRED
ts	Five-year Treasury constant-maturity interest rate minus 90-day Treasury bill rate, quarterly average	GS5 – TB3MS/FRED
wi	Wage inflation	ln[COMPNFB×HOANBS]/FRED
cp	PPI commodity price inflation	BLS

hypothesis that the overidentifying restrictions hold. The hypothesis implies the same parameters apply for any instrument set. We include the J statistic, along with its degrees of freedom (df), over its p-value. The J statistic is asymptotically distributed χ^2 with df equal to the number of overidentifying restrictions (the number of instruments minus the number of parameters). Fixing the number of overidentifying restrictions, a larger J statistic yields a smaller p-value, which indicates that the residual is predictable and constitutes a rejection of the hybrid NKPC.

The top row of Table 2 presents CU-GMM estimates based on four instruments. With three hybrid NKPC parameters to estimate, this gives one overidentifying restriction. Besides lagged inflation, the set of instruments in the first row contains only lags of marginal cost, x_t. Between the first and last rows of Table 2, we add instruments one by one from the longer Galí and Gertler (1999) list. The penultimate row of Table 2 includes the entire set of 24 instruments used by Galí and Gertler. (The last row presents least-squares estimates, discussed in Section 4.)

Table 2 shows that the hybrid NKPC estimates vary with the set of instruments. We return to that finding in Section 4. Meanwhile, once we include lags of marginal cost and of the output gap, y, we find that γ_b and γ_f are significant, positive fractions, with the weight on expected future inflation much greater than the weight on lagged inflation (see the second and sixth rows). Estimates based on these instrument sets indicate that there is little inflation inertia. These estimates also reveal a significant, positive impact of real unit labor costs, measured by $\hat{\lambda}$, but the scale of the inflation response is small. Sbordone (2002) and Eichenbaum and Fisher (2007) show that if firms have firm-specific capital it can lead to a low response to a cost shock, i.e., a small value for λ.

Table 2 U.S. New Keynesian Phillips Curve, 1955:1–2007:4

	$E\left[\pi_t - \gamma_b \pi_{t-1} - \gamma_f \pi_{t+1} - \lambda x_t \mid z_t\right]=0$						
Instruments (z_t)	$\hat{\gamma}_b$ (se)	$\hat{\gamma}_f$ (se)	$100 \times \hat{\lambda}$ (se)	$\hat{\omega}$ (se)	$\hat{\theta}$ (se)	$\hat{\beta}$ (se)	$J(df)$ (p)
$\pi_{t-1}, x_{t-1}, x_{t-2}, x_{t-3}$	−0.11 (0.50)	1.14 (0.58)	1.41 (1.00)	−0.08 (0.36)	0.88 (0.03)	1.03 (0.06)	0.30(1) (0.58)
$\pi_{t-1}, x_{t-1}, x_{t-2}, x_{t-3}, y_{t-1}, y_{t-2}$	0.29 (0.10)	0.68 (0.10)	0.67 (0.35)	0.36 (0.17)	0.91 (0.03)	0.94 (0.05)	3.42(3) (0.33)
$\pi_{t-1}, x_{t-1}, x_{t-2}, x_{t-3}, ts_{t-1}, ts_{t-2}$	−0.20 (0.38)	1.25 (0.44)	0.17 (0.82)	−0.15 (0.23)	0.89 (0.02)	1.04 (0.04)	1.05(3) (0.79)
$\pi_{t-1}, x_{t-1}, x_{t-2}, x_{t-3}, wi_{t-1}, wi_{t-2}$	0.17 (0.18)	0.81 (0.20)	0.89 (0.45)	0.18 (0.24)	0.90 (0.02)	0.97 (0.06)	2.26(3) (0.52)
$\pi_{t-1}, x_{t-1}, x_{t-2}, x_{t-3}, cp_{t-1}, cp_{t-2}$	0.02 (0.17)	1.00 (0.19)	1.07 (0.53)	0.02 (0.16)	0.88 (0.03)	1.02 (0.05)	1.45(3) (0.69)
$(\pi, x, y, ts, wi, cp)_{t-1}, \ldots$ $(\pi, x, y, ts, wi, cp)_{t-4}$	0.25 (0.05)	0.73 (0.05)	0.63 (0.31)	0.29 (0.08)	0.92 (0.02)	0.95 (0.03)	17.08(21) (0.71)
$\pi_{t-1}, \pi_{t+1}, x_t$ (NLLS)	0.25 (0.01)	0.75 (0.03)	2.65 (1.33)	0.27 (0.01)	0.81 (0.06)	0.99 (0.06)	

Notes: Data are de-meaned prior to estimation. The estimator is CU-GMM, except the last case is NLLS, and all use a Newey-West HAC correction and automatic plug-in lag length.

The second and sixth rows of Table 2 also contain estimates of the structural Calvo-pricing parameters that are positive fractions, significant, and in keeping with the theory. These estimates yield a discount factor of around 0.95 and indicate that about 90 percent of firms are unable to change prices in a given quarter. Of the 10 percent that can change prices, about a third decide against it.

Additional support for the NKPC is provided by the J-test statistics in the final column of Table 2. Across all instrument sets, we obtain a large p-value associated with this statistic, so that the overidentifying restrictions cannot be rejected.

3. STABILITY

So far, the updated empirical evidence supports the hybrid NKPC. However, as mentioned in the introduction, it also is natural to ask whether the hybrid NKPC parameters are stable over time. To test their stability, we divide the entire sample at a given quarter, called a break date, into two separate subsamples. We consider all possible break dates between 1963:1 and 1999:4, which trims 15 percent from the beginning and end of the entire sample. For each date in this range, we estimate the NKPC twice: first from 1955:1 to the break date quarter and second from one quarter beyond the break date to 2007:4. All CU-GMM estimates employ the instrument vector $\{\pi_{t-1}, x_{t-1}, x_{t-2}, y_{t-1}, y_{t-2}, wi_{t-1}, wi_{t-2}\}$, which mimics that of Galí, Gertler, and López-Salido (2005) minus $\{\pi_{t-2}, \pi_{t-3}, \pi_{t-4}\}$.

Figure 2 presents the results. The three panels plot estimates of the reduced-form parameters γ_b, γ_f, and λ, respectively, for break dates from 1963:1 to 1999:4. In each panel, the black line graphs estimates from the sample beginning in 1955:1 and ending at the break date shown, while the gray line graphs estimates from a sample beginning at the break date plus one quarter to 2007:4. For any date, the vertical distance between the two lines gives the difference between estimates from the "before" and "after" samples. The figure shows variation in the estimates that is limited for break dates since 1980, but noticeable for earlier break dates, particularly for $\hat{\gamma}_b$ and for $\hat{\lambda}$. The coefficient $\hat{\gamma}_b$ ranges from 0.29 to 0.40 estimating on the before sample and from 0.27 to 0.35 estimating on the after sample. For $\hat{\gamma}_f$, the corresponding ranges are from 0.64 to 0.73 and from 0.60 to 0.69, respectively. A glance at the vertical axes shows that these estimates confirm the earlier finding that the coefficient on expected future inflation, $\hat{\gamma}_f$, exceeds the coefficient on lagged inflation, $\hat{\gamma}_b$. Estimates of $100 \times \hat{\lambda}$ range from -1.31 to 1.03 on the before sample and from 0.08 to 0.87 on the after sample. However, ignoring estimates from the 1960s break dates constrains $100 \times \hat{\lambda}$ to range from about zero to around one.

Figure 2 Parameter Stability

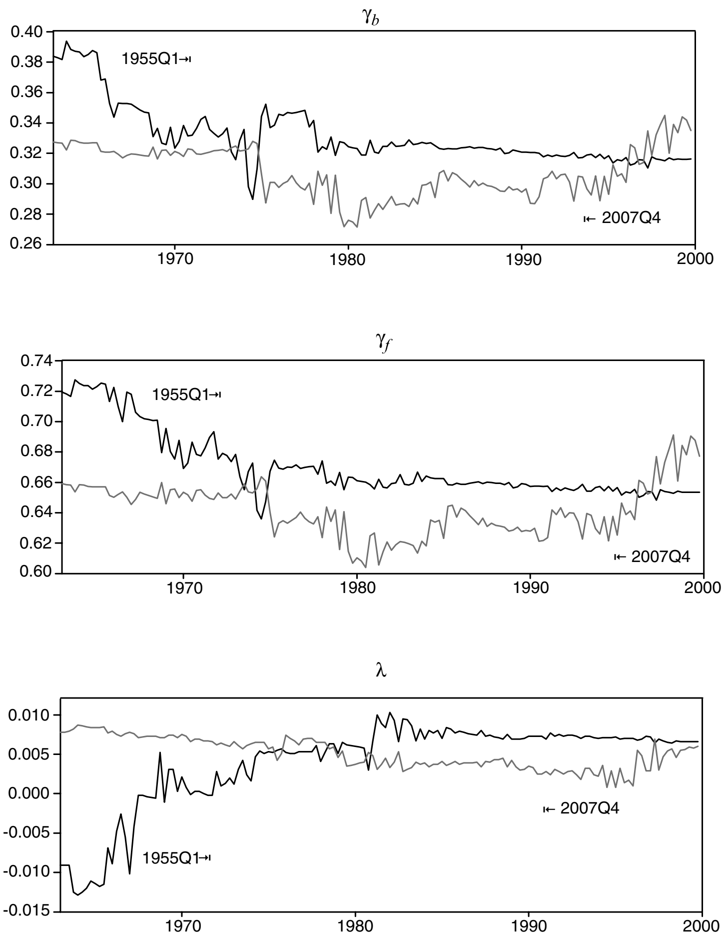

Next, we test one-by-one whether any of the three parameters $\{\gamma_b, \gamma_f, \lambda\}$ changes significantly from the first time period to the second time period. The method developed by Andrews (1993, 2003) allows a statistical test along these lines without pre-supposing knowledge of the exact date at which a break or shift in a parameter value took place. Following this method, we calculate the Wald test for the hypothesis that there is a significant difference between the "before" and "after" estimates for γ_b, γ_f, and λ over 1963–1999. We record the maximal value for each of these Wald statistics. Andrews (2003) gives critical values for this test statistic, while Hansen (1997) provides a method for computing p-values. For γ_b, the test statistic is 1.50 with a p-value of

0.91; for γ_f, the test statistic is 1.41 with a p-value of 0.92; and for λ, the test statistic is 3.30 with a p-value of 0.50. Since these p-values are far above conventional levels of statistical significance such as 0.05 or 0.10, the tests fail to reject the hypothesis that the parameters are stable.

In summary, our tests suggest that the reduced-form hybrid NKPC parameters are stable. This result is striking because the behavior of inflation has changed over time. For example, inflation was on average lower and less volatile after the mid-1980s. Yet despite this change in the statistical properties of inflation, the link between inflation and marginal cost has remained stable. The stability of this relationship is striking because it suggests a flat hybrid NKPC—with the same relatively low slope $\hat{\lambda}$—during the business cycles of the 1970s and the Great Moderation and disinflation that took hold in the mid-1980s. Whatever the sources of this Great Moderation in inflation, the single-equation stability tests suggest that they acted through real activity as measured by marginal costs.

4. WEAK IDENTIFICATION?

Section 2 referred to the need to find valid instruments in order to use single-equation methods. Instruments must satisfy two statistical criteria. First, they must be as numerous as the parameters and must help predict or forecast π_{t+1} so that a projection based on them can be reasonably substituted for the unobserved forecast on the right-hand side of the Phillips curve. Second, they must be uncorrelated with the error term in the econometric equation, just like any regressor.

Unfortunately, these two criteria sometimes can conflict. To see how this can come about, recall that researchers often have used lagged instruments, z_{t-1}. The rationale for this choice is that these past outcomes must be exogenous and, therefore, uncorrelated with unobserved shocks to today's inflation rate, thus satisfying the second criterion. But now satisfying the first criterion can be challenging. The researcher needs to find at least one variable, z_{t-1}, that helps forecast π_{t+1}. Also, the list of instruments must include something other than the other two variables that enter the hybrid NKPC, x_t and π_{t-1}. That is because the constructed forecast $E\pi_{t+1}|z_{t-1}$ has to exhibit some variation independent of x_t and π_{t-1}. Otherwise, there will be no possibility to measure or identify separately the effects of π_{t-1}, x_t, and $E_t\pi_{t+1}$ on current inflation. We want to identify these three effects on current inflation so, logically, we need an inflation forecast that sometimes varies separately from π_{t-1} and x_t.

Seen in this way, the problem of finding instruments is recast as the problem of trying to forecast inflation but with a twist. The statistical challenge is to predict next quarter's inflation rate, π_{t+1}, but without using this quarter's inflation rate, π_t (because it is the variable we seek to explain on the left-hand side of the hybrid NKPC), or last quarter's inflation rate, π_{t-1}, or this quarter's costs or aggregate demand, x_t (because they appear separately on the right-hand side of the hybrid NKPC). Forecasting inflation is difficult, even without one hand tied behind one's back in this way. The statistical studies by Stock and Watson (1999, 2007) and Ang, Bekaert, and Wei (2007) show that it is challenging to find a stable relationship that can be used to forecast U.S. inflation, especially over the past 15–20 years. Perhaps competent central bankers can take some credit for creating a low, stable inflation rate that has not displayed persistent swings or cycles, but that outcome inherently makes it difficult to isolate an inflation forecast that differs from current or lagged inflation.

The hybrid NKPC provides another perspective on how to forecast π_{t+1}. We lead the present-value version of the hybrid NKPC forward by one time period and forecast to obtain

$$\mathrm{E}_t\pi_{t+1} = \delta_1\pi_t + \left(\frac{\lambda}{\delta_2\gamma_f}\right)\sum_{k=0}^{\infty}\left(\frac{1}{\delta_2}\right)^k \mathrm{E}_t x_{t+1+k}.$$

Next, suppose that x_t can be forecasted only from its own, lagged value. Suppose that marginal costs follow a first-order autoregression, with coefficient ρ, so that its multistep forecast is

$$\mathrm{E}_t x_{t+1+k} = \rho^{1+k} x_t.$$

Combining the last two equations gives the forecasting equation

$$\mathrm{E}_t\pi_{t+1} = \delta_1\pi_t + \frac{\lambda\rho}{\gamma_f(\delta_2-\rho)}x_t.$$

In this case, the three reduced-form parameters cannot be identified by GMM because there is no source of variation in $\mathrm{E}_t\pi_{t+1}$ other than π_t and x_t (which already are included in the hybrid NKPC). Nason and Smith (2008) show that if x_t is an autonomous pth-order autoregression, p must equal 2 to just identify and be greater than 2 to overidentify the three hybrid NKPC parameters using GMM. In other words, higher-order dynamics are needed for identification if x_t is predicted from its own past. That is why the first row of Table 2 includes several lags of marginal cost.

Nason and Smith (2008) also show that setting the NKPC in a broader, New Keynesian model does not suggest sources of identification for single-equation estimation. They show analytically the problem facing an econometrician who tries to estimate the NKPC by GMM in a textbook world where the hybrid NKPC combines with a dynamic IS curve and a Taylor rule. It turns out that there may be no valid instruments available. The logic is that the econometrician must lag instruments to make sure that they are uncorrelated with the residual in the NKPC equation. But lagging them enough to satisfy that criterion for instrumental variables also makes them irrelevant for forecasting π_{t+1}.

We dwell on the challenges of forecasting inflation because of another statistical issue. For instrumental variables (GMM) estimation to be informative, it turns out that we need a significant amount of predictability. Imagine reconstructing a forecast equation by regressing π_{t+1} on z_{t-1}. The F-statistic for the joint significance of the variables z_{t-1} in this regression must be above some threshold in order for the full GMM estimation of the hybrid NKPC to yield meaningful results. If this F-statistic, or inflation predictability, is too low, then the econometrician is said to be using weak instruments. In that case, the subsequent estimates of the hybrid NKPC parameters will be imprecisely estimated (possess large standard errors). Also, hypothesis tests may have the wrong size (probability of type I error); for example, they may not reject often enough. These problems will persist even in large samples.

Another symptom of the syndrome of weak identification is that estimates may vary a great deal with changes to the instrument set. Two economists with the same hybrid NKPC may obtain disparate parameter estimates when they employ different, but apparently equally admissible, instrument sets, z_{t-1}. In Table 2, this sensitivity is apparent in the GMM estimates of the reduced-form and structural hybrid NKPC parameters that are grounded on different combinations of the Galí-Gertler instruments. Since different researchers have tended to apply different instrument sets, weak identification might help to explain the current lack of consensus on parameter estimates of the hybrid NKPC.

Recall from Section 2 that least-squares estimation of the NKPC yields inconsistent estimates; we cannot represent $E_t\pi_{t+1}$ simply by replacing it with the actual value, π_{t+1}. Another useful result from research on weak instruments is that instrumental-variables estimates converge to least-squares estimates as the econometrician adds more and more weak instruments. The last row of Table 2 shows what can happen in that case by reporting least-squares estimates. With the exception of the larger value for $\hat{\lambda}$, the least-squares estimates are similar to those in some previous rows. That similarity shows that finding these plausible values for the coefficients does not necessarily imply they have a sound statistical basis. And it raises the possibility that the GMM estimates are only weakly identified.

Ma (2002), Mavroeidis (2005), Dufour, Khalaf, and Kichian (2006), and Nason and Smith (2008) draw attention to the pitfalls of weak identification in GMM estimation of the hybrid NKPC. One response to this issue has been to reformulate the hybrid NKPC so that it involves fewer parameters. The idea is simply that by trying to measure a shorter list of effects, the investigator might have greater success in precisely measuring them. For example, one could set $\gamma_b = 0$ and so work with the original NKPC rather than the hybrid version. In that case, π_{t-1} also would become available as an instrument.

A number of investigators—including Henry and Pagan (2004) and Rudd and Whelan (2006)—suggest restricting the reduced-form, hybrid NKPC parameters so that $\gamma_b + \gamma_f = 1$. Imposing this restriction helps with identification by reducing the number of coefficients to be estimated by one. It turns out, though, that this restriction is inconsistent with one popular interpretation of the hybrid NKPC parameters, namely that they reflect an underlying Calvo-type model of staggered pricing. To show this, we use the earlier equations that Galí and Gertler (1999) outline to connect the hybrid NKPC parameters to those of the Calvo pricing model, namely ω, θ, and β. The proposed restriction gives

$$\gamma_b + \gamma_f = \frac{\omega}{\phi} + \frac{\beta\theta}{\phi} = 1,$$

where $\phi = \theta + \omega[1 - \theta(1 - \beta)]$. Some algebra reveals that this restriction implies that the fraction of firms that can change prices but choose not to is $\omega = 1$. However, this extreme result forces the reduced-form parameter on marginal costs,

$$\lambda = \frac{(1-\omega)(1-\theta)(1-\beta\theta)}{\phi},$$

to equal zero. Although Galí and Gertler point out that $\beta = 1$ also is consistent with $\gamma_b + \gamma_f = 1$, often this restriction is imposed without recourse to calibrating the firm's discount factor to one.

More generally, restricting the hybrid NKPC parameters can be problematic because we want to test hypotheses about all relevant values. We next explore statistical methods that apply even if identification is weak. An econometrician also can test the hybrid NKPC parameters (and compute their confidence intervals) using methods that are robust to weak identification, i.e., that remain valid whether the instruments are weak or not.

Many of these robust methods are based on a 60-year-old statistical insight from Anderson and Rubin (1949). Here is their idea, as applied to the hybrid NKPC. Rewrite the equation by taking the future value of inflation to the left-hand side (without forecasting it) and by adding some list of other variables, u_t, on the right-hand side:

$$\pi_t - \gamma_{f0}\pi_{t+1} = \gamma_b\pi_{t-1} + \lambda x_t + \delta u_t.$$

To create this composite variable on the left-hand side of the equation, we need to choose a value for γ_f, labelled γ_{f0}. We cannot use this regression to estimate that value. But it *can* be used to *test* any value for this weight on expected future inflation. To test the hypothesis that $\gamma_f = \gamma_{f0}$, we simply perform a traditional F-test of the hypothesis that $\delta = 0$ so that the auxiliary variables, u_t, are insignificant. The logic is that if we happen to select the correct value for γ_f, then the three explanatory variables in the hybrid NKPC will reproduce the time-series pattern in inflation π_t, and no systematic pattern in the residuals will be detected by including other macroeconomic variables, u_t.

Figure 3 Anderson-Rubin Tests

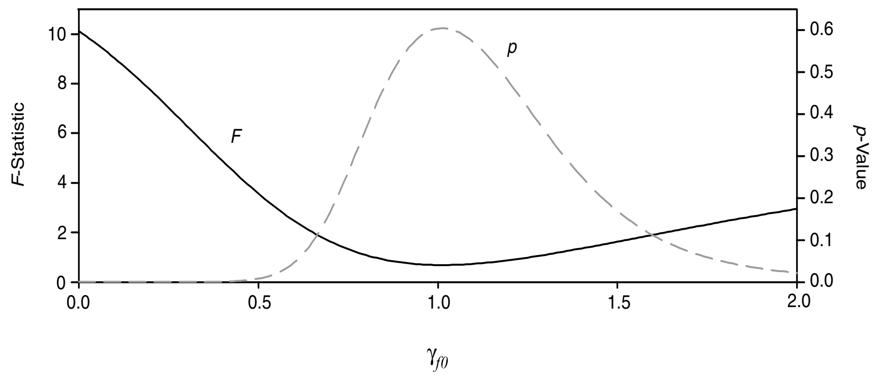

To illustrate the Anderson-Rubin (AR) test, we collect auxiliary variables that include the 90-day Treasury bill interest rate, r_t (again, a natural variable to consider in forecasting inflation), as well as extra lags of inflation and unit labor costs. The complete list is: $u_t = \{r_t, r_{t-1}, x_{t-1}, \pi_{t-2}\}$. The sample period is 1955:1–2007:4. We run the regression on a fine grid of values of γ_{f0} between 0 and 2. For each such value we record the F-statistic associated with the restriction that none of the variables in u_t enters the equation, and we calculate the corresponding p-value by locating the statistic in the $F(4, 204)$ distribution. Figure 3 graphs the candidate values, γ_{f0}, on the horizontal axis and the F-statistics (the solid black line) and their p-values (the dashed gray line) on the two vertical axes.

Figure 3 shows that the AR test rejects the restrictions for low values of the weight on expected future inflation and also for high values. In particular, when γ_{f0} is less than 0.5 or greater than 1.5, the F-statistics are high and the p-values are low. This means that $\hat{\delta}$ is far from zero, the auxiliary variables, u_t, enter the equation, and so the candidate values of γ_{f0} can be rejected. The test does not reject at intermediate values of γ_{f0}. The F-statistic reaches its minimum and the associated p-value its maximum for γ_{f0} around 1.0.

We already know that Table 2 has CU-GMM estimates of γ_f that are a large positive fraction (though the estimate depends on the instrument set) with a small standard error. Moreover, the J test did not reject the overidentifying restrictions. So what is gained from the AR approach? The answer is that tests in Table 2 may have been affected by weak identification, whereas statistics in Figure 3 apply whether identification is weak or not. To illustrate the effect of this robust method on inference, note that the range of values for which the F-statistics in Figure 3 fall below the α-percent critical value of the F-distribution (equivalently the p-values lie above α) constitutes the

$1-\alpha$-percent confidence interval for γ_f. In this case, the 90 percent confidence interval is (0.66, 1.62) and the 95 percent confidence interval is (0.61, 1.78). For comparison, the traditional, asymptotic confidence intervals for γ_f from the GMM estimates in the second-to-last row of Table 2 are (0.65, 0.81) at the 90 percent level and (0.63, 0.83) at the 95 percent level. These intervals understate the uncertainty, compared with the intervals that are robust to weak instruments. The AR test suggests a positive value for γ_f, but considerable uncertainty or imprecision remains, and values greater than 1 are possible.

How to draw inference with weak instruments is an active area of research by statisticians. Excellent surveys of inference under weak identification are provided by Dufour (2003) and Andrews and Stock (2006). The AR test assumes x_t is exogenous, whereas some more recent methods allow it to be endogenous. These methods allow tests of all the NKPC parameters, whereas we have focused only on γ_f. One important finding in this research is that the AR test also may lack power, especially when there is overidentification. In other words, it may fail to reject a false, assumed value γ_{f0} and so give too wide a confidence interval. This outcome is particularly likely if there are many auxiliary variables, u_t, and some are irrelevant as instruments.

Using identification-robust methods, Ma (2002) finds large confidence sets for the hybrid NKPC parameters, which suggests that they are weakly identified. Dufour, Khalaf, and Kichian (2006) apply the AR test and some more recent tests to the United States for a 1970–1997 sample. They too find wide confidence sets. Nason and Smith (2008) reject the hybrid NKPC for the United States—by finding empty confidence intervals—when testing either reduced-form parameters or the underlying ones ω, θ, and β. For no value of γ_{f0} does the hybrid NKPC produce unpredictable residuals, so the confidence intervals are empty. (They use a slightly different definition of x_t, described in Section 6.) Kleibergen and Mavroeidis (2008) use identification-robust methods and conclude that $\gamma_f > \gamma_b$. However, they find wide confidence intervals, especially for γ_b and for λ, where the confidence interval includes zero. They also apply a stability test devised by Caner (2007) that is robust to weak identification. This test suggests that the NKPC experienced a structural break around 1984 and subsequently became flatter. Overall, methods that are robust to weak identification suggest more skepticism about the NKPC than do traditional econometric tools. They reveal considerable uncertainty about the NKPC parameters or, in some cases, reject all reasonable values.

One way to gain power in tests like these (or to find more precise estimates of the hybrid NKPC parameters) is to utilize more information on the inflation forecasting equation and the evolution of the exogenous variable x_t, a conclusion that directs us to consider systems of econometric equations that set the hybrid NKPC within a broader economic/statistical model. In these systems, researchers supplement the hybrid NKPC either with (a) an explicit, statistical forecasting model that recognizes that x_t is most likely endogenous, or (b) additional equations like a policy rule and dynamic IS curve so as to form a coherent New Keynesian model. Either of these approaches can potentially provide more precision at the cost of introducing bias if the added assumptions are misspecified. Studies that use forecasting models (vector autoregressions) include those of Fuhrer (1997), Sbordone (2002, 2005), Kurmann (2005, 2007), and Rudd and Whelan (2005a, 2005b, 2006), while Lindé (2005) uses a three-equation New Keynesian model. On the other hand, Galí, Gertler, and López-Salido (2005) review these approaches and conclude that GMM estimation remains informative. Schorfheide's article in this issue provides a complete review of systems estimation of the hybrid NKPC.

5. FORECAST SURVEY DATA

As we have noted, many of the statistical challenges with estimating and testing the NKPC arise because inflation expectations cannot be directly observed. There is an alternative to constructing these forecasts with instrumental variables, though, and that is simply to ask some people what they expect the inflation rate to be in the next quarter. The Federal Reserve Bank of Philadelphia does just this in its Survey of Professional Forecasters (SPF). There are other measures of actual forecasts, but they tend to belong to forecasters either with (potentially) more information (in the case of the Federal Reserve's Greenbook forecasts) or different information (in the case of the Michigan household survey) than we might attribute to a typical, price-setting firm. These issues have helped to make the SPF the most widely used data source in this context. Another reason to favor the survey-based measures is that they are in real time. Unlike our typical, instrumental-variables estimates, their construction does not involve estimation with any data reported subsequent to the date of the forecast. Roberts (1995), Orphanides and Williams (2002, 2005), Adam and Padula (2003), Dufour, Khalaf, and Kichian (2006), Zhang, Osborn, and Kim (2008), and Brissimis and Magginas (2008) use forecasts to estimate the NKPC.

Next, we see what happens when we use the median forecast from the SPF in our estimator. The series on expected inflation, labeled π^s_{t+1}, is the median of the one-quarter-ahead forecasts of the GDP deflator growth rate quarter-to-quarter at annual rates, `dpgdp3` from the SPF file `MedianGrowth.xls`, and is available for 1968:4–2007:4. In fact, the SPF survey referred to the GNP deflator until the end of 1991. This matters for the actual inflation rate, π_t, used to estimate the hybrid NKPC when the median SPF inflation is equated with expected inflation. In this case, we measure π_t with `GNPDEF` from `FRED` for the period prior to 1991.

As a benchmark, we present CU-GMM results similar to those in Table 2, but with a 1969:1–2007:3 sample. The first row of Table 3 has these results.

Table 3 Forecast Surveys in the U.S. NKPC, 1969:1–2007:3

Forecast	$\hat{\gamma}_b$ (se)	$\hat{\gamma}_f$ (se)	$100\times\hat{\lambda}$ (se)	$\hat{\omega}$ (se)	$\hat{\theta}$ (se)	$\hat{\beta}$ (se)	$J(df)$ (p)
$\mathrm{E}\left[\pi_{t+1} \mid z_t\right]$	0.27	0.72	2.46	0.30	0.81	0.97	3.58
	(0.07)	(0.07)	(1.66)	(0.10)	(0.05)	(0.06)	(0.73)
$\mathrm{E}\left[\pi^s_{t+1} \mid z_t\right]$	0.38	0.56	1.36	0.50	0.83	0.83	11.63
	(0.12)	(0.17)	(4.56)	(0.15)	(0.24)	(0.23)	(0.07)
π^s_{t+1}	0.36	0.68	-0.14	0.56	0.91	1.16	
(NLLS)	(0.03)	(0.03)	(0.13)	(0.04)	(0.11)	(0.09)	

Notes: Data are de-meaned prior to estimation. The estimator is CU-GMM with a Newey-West HAC correction and automatic, plug-in lag length. The instrument vector is $\pi_{t-1}, x_{t-1}, x_{t-2}, wi_{t-1}, wi_{t-2}, y_{t-1}, y_{t-2}, cp_{t-1}, cp_{t-2}$ using a linearly detrended output gap.

The weight on future inflation is greater than the weight on past inflation, and both are estimated precisely. The estimated values for the underlying parameters ω, θ, and β, are similar to some of those found in Table 2. The J test does not reject, but the coefficient on marginal costs, $\hat{\lambda}$, while positive, is statistically insignificant.

The second row of Table 3 lists results when the median survey value is equated with expected inflation. We continue to estimate with CU-GMM to allow for the possibility that this inflation expectations measure is contaminated with measurement error that is correlated with the survey measure. By comparison, with the CU-GMM estimates in the top row, the weights on past and future inflation tilt, with a larger weight on lagged inflation and a smaller one on expected future inflation in the second row of Table 3. There is now no significant role for marginal costs in explaining the inflation series ($\hat{\lambda}$ is smaller than its standard error) and the J test rejects the overidentifying restrictions at conventional levels of significance.

Brissimis and Magginas (2008) perform a similar exercise but find that the $\hat{\gamma}_f$ and $\hat{\gamma}_b$ weights tilt in the opposite direction, with a large weight on expected future inflation and no statistically significant weight on lagged inflation. They use the Bureau of Labor Statistics measure of the labor share of output as x_t, whereas we use the adjusted Sbordone (2002) measure that also is adopted by Galí and Gertler (1999). This sensitivity of the findings with forecast survey data to the measure of marginal costs may show either that we need further research on modeling marginal costs or that this is not a fruitful way to model expectations in the hybrid NKPC.

Finally, we replace the unobservable $E[\pi_{t+1}|I_t]$ with π^s_{t+1} and estimate the hybrid NKPC by least squares. Taking this step does not mean assuming these two series coincide. Instead, it yields consistent estimates whenever the median survey is based on less information than that reflected in forecasts driving actual inflation, so that

$$E[\pi_{t+1}|I_t] = \pi^s_{t+1} + \eta_t,$$

in which η_t is uncorrelated with π^s_{t+1}. In other words, it assumes that the median forecast is an unbiased predictor of the broader-based inflation forecast that influences the behavior of Calvo price setters.

The third row of Table 3 contains the results of least-squares estimation with the median report from the SPF. The striking finding is that $\hat{\lambda}$ is negative so that real unit labor costs enter the equation with the wrong sign. However, the point estimate is small and statistically insignificant. This finding can be viewed as evidence against the use of the median survey measure. Perhaps there is an errors-in-variables problem associated with this representation of expected inflation. But it is not straightforward to explain a negative coefficient, albeit an insignificant one, which argues that this finding also can be viewed as evidence against the hybrid NKPC.

A resolution to the question of how to represent expected inflation, that is, with instrumental variables forecasts or survey forecasts, can be found by including both. Smith (2007) and Nunes (2008) include a linear combination of the two measures and ask which combination best explains current inflation. The estimating equation becomes

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f \left(\mu E[\pi_{t+1}|z_t] + (1-\mu)\pi^s_{t+1}\right) + \lambda x_t.$$

Nunes offers an economic interpretation of this mixture as reflecting price-setters' different forecasting methods. Smith instead has a purely statistical interpretation. Either way, the evidence is that both measures matter. Their estimates place a slightly greater weight on the survey measure than on the econometric measure. The estimated hybrid NKPC parameters $\{\hat{\gamma}_b, \hat{\gamma}_f, \hat{\lambda}\}$ in these studies resemble those in Table 2 and are consistent with theory.

6. WHAT DRIVES INFLATION?

Up to this point we have studied the hybrid NKPC in which inflation tracks real unit labor costs. But several authors have argued that measures of the output gap (i.e., the cyclical component of real GDP) are better explanatory variables for inflation. This section sets both types of x-variables in the hybrid NKPC to learn which might be most useful for explaining U.S. inflation. We first describe the properties of nine different candidate x-variables. We then use these series to estimate the hybrid NKPC.

We consider two measures of real unit labor costs. RULC1 is the Sbordone (2002) measure described earlier in Table 1. RULC2 measures real unit labor cost as the cointegration relation of the logarithm of nominal unit labor cost with the logarithm of the GDP deflator (allowing for an intercept and time trend). The estimated cointegrating coefficient is 1.03. Nason and Slotsve (2004) show that RULC2 is consistent with Calvo's staggered pricing mechanism. Nason and Smith (2008) use this variable in their estimated hybrid NKPC.

The first measure of the output gap, labeled CBO, is published by the Congressional Budget Office. The remaining measures are based on per capita output. LT (QT) is the series of residuals from linearly (quadratic) detrending per capita real GDP. Next, measures UC and BN are based on the unobserved-components model and Beveridge-Nelson decomposition, respectively. Both of these measures treat real per capita output as the sum of a permanent component and a transitory component. These time-series models assume the permanent component of output is a random walk with drift while the transitory component follows a second-order autoregression. The difference is that the UC model imposes a zero correlation between innovations to the permanent and transitory components. The BN decomposition estimates this correlation, which is -0.97. Maximum likelihood estimation of the UC and BN models is undertaken with the Kalman filter, and the associated output gap estimates are filtered, not smoothed. The UC and BN output gap measures rely on the work of Morley, Nelson, and Zivot (2003).

Measure BK is based on the Baxter and King (1999) bandpass filter. Since the technical details for this implementation are not straightforward, we refer statistically inclined readers to Harvey and Trimbur (2003). They show how to estimate the BK cycle or output gap with the Kalman filter. Finally, measure HP is the cycle that remains after applying the Hodrick and Prescott (1997) filter to output growth, as implemented by Cogley and Nason (1995).

Figure 4 plots the nine (demeaned) measures. All series are shown since 1955 (omitting the volatile Korean war period), but the vertical scale varies across the three panels. In the top panel the two measures of marginal cost have different trends, but RULC1 tends to be dominated by low-frequency movements. The middle panel shows the CBO output gap and the two deterministically detrended output gaps. These three generate more cycles than the marginal cost measures and are also more volatile than RULC1 and RULC2. The CBO, LT, and QT output gaps behave similarly except during the late 1960s and since 1999. The bottom panel of Figure 4 presents UC, BN, BK, and HP measures of the output gap. The BN and BK output gaps have most of their variation between two and four years per cycle, while relatively lower-frequency fluctuations produce most of the variation in the UC and HP output gaps. Volatility varies from one measure to another, with the UC and HP output gaps exhibiting the most variance.

Nason and Smith (2008) show that being able to predict future x_t can be key for the viability of single-equation approaches to the NKPC. Recall

Figure 4 Marginal Costs and Output Gaps

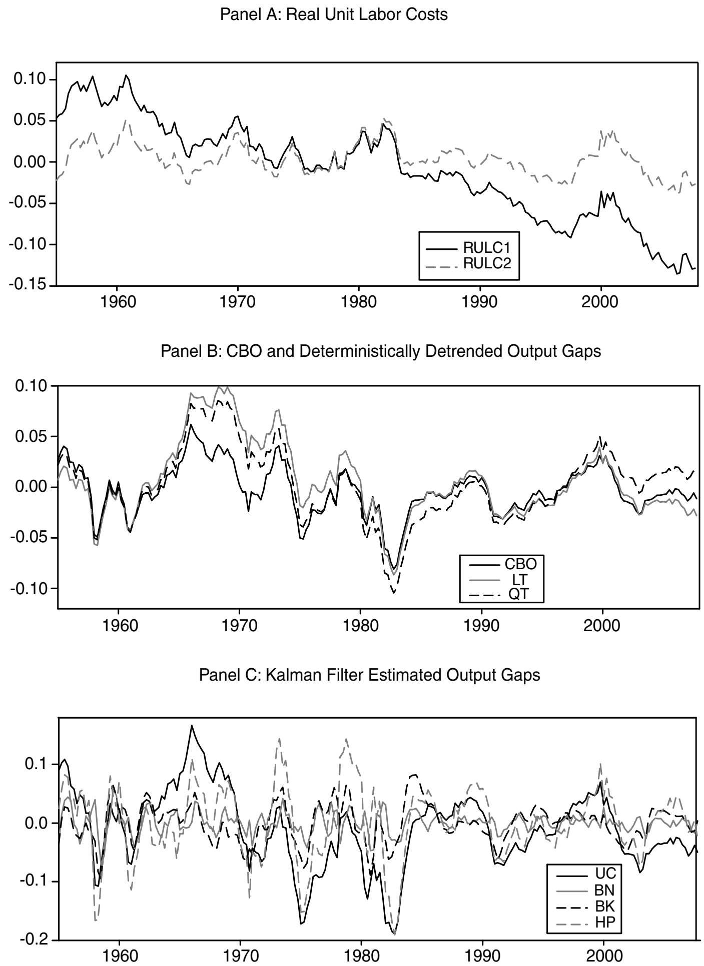

that (a) according to the hybrid NKPC, inflation is related to lagged inflation and to the present value of current and future x_t, and (b) finding instruments involves predicting next quarter's inflation rate, π_{t+1}. Combining these two facts means that we must predict future values of x_t in order to identify the NKPC.

One possibility discussed in Section 4 is that x_t can be forecasted using its own lagged values. In that case, higher-order dynamics are needed for identification. The idea that some complicated dynamics in x_t help us learn about the NKPC makes intuitive sense. If there are predictable movements in these fundamentals, they should be matched by swings in inflation. The extent to which they are matched can shed light on whether the NKPC is a good guide to inflation. If, instead, there are no predictable movements in x_t that inflation is theorized to be tracking, there will be no way to identify the response of inflation.

We test for lag length in univariate autoregressions for each of the nine x-variables using the Akaike information criterion, Hannan-Quinn information criterion, Schwarz or Bayesian information criterion, and likelihood ratio (LR) test. The evidence is that for most of these series there are complicated dynamics in which three to five lags contain forecasting information. The two measures of unit labor costs, RULC1 and RULC2, and the BN output gap appear to be exceptions, because LR tests suggest high-order dynamics that reach 10 to 12 lags. These results suggest the RULCs and output gaps have the requisite dynamics to overidentify the three structural price-setting or reduced-form parameters of the hybrid NKPC.

Of course, our main reason for studying RULC1 and RULC2 or the output gaps is to use them in the hybrid NKPC. Thus, the main goal of this section is to estimate this NKPC with the nine x-measures. Table 4 contains the results. When we estimate using RULC1 or RULC2 we find a small positive effect of marginal costs on inflation. The significance of this effect depends on the instrument set, as we documented in Section 2.

The remaining rows of Table 4 present NKPC estimates using the output gaps. We find no significant role for any of these x-measures. The coefficient on x_t, $\hat{\lambda}$, is small and negative (albeit statistically insignificant) for all the output gaps. However, most economists would predict the opposite effect: a positive output gap leading to a rise in prices. The other hybrid NKPC coefficients also take surprising values that are difficult to interpret. The coefficient on lagged inflation is negative, while the coefficient on expected future inflation is greater than one. These coefficients on past and future inflation most likely are affected by omitted-variables bias. Without some confidence in one's measure of the x-variable that inflation tracks in the NKPC, there cannot be much confidence in estimates of inflation inertia or other properties of inflation dynamics.

Investigators who work with an output gap might sometimes wonder whether their findings depend on the specific filtering or detrending procedure they use to measure this variable. We have used measures that are commonly adopted and that have been used in forecasting or explaining inflation, yet found no role for any of them. Our evidence suggests little support for the idea that the output gap drives U.S. inflation.

Some recent studies work with inflation and output gaps and do find statistical links between them. Harvey (2007) adopts an unobserved-components model of both inflation and output and finds a link between the cyclical components of the two series. Basistha and Nelson (2007) use the NKPC to define or measure the output gap so that it fits into the NKPC by construction. But

Table 4 U.S. New Keynesian Phillips Curve, 1955:1–2007:4

$$\mathrm{E}\left[\pi_t - \gamma_b\pi_{t-1} - \gamma_f\pi_{t+1} - \lambda x_t \mid z_t\right] = 0$$

x-**Measure**	$\hat{\gamma}_b$ (se)	$\hat{\gamma}_f$ (se)	$100\times\hat{\lambda}$ (se)	$J(df)$ (p)
RULC1	0.28	0.68	0.56	1.28(1)
	(0.16)	(0.18)	(0.27)	(0.26)
RULC2	0.34	0.61	0.70	1.64(1)
	(0.16)	(0.18)	(0.37)	(0.20)
CBO	−0.57	1.70	−3.90	0.02(1)
	(0.84)	(1.02)	(3.13)	(0.89)
LT	−0.58	1.76	−2.90	0.02(1)
	(0.70)	(0.88)	(2.20)	(0.90)
QT	−0.66	1.82	−2.60	0.11(1)
	(0.77)	(0.94)	(1.83)	(0.73)
UC	−0.51	1.59	−0.97	0.80(1)
	(0.55)	(0.65)	(0.65)	(0.37)
BN	−0.43	1.56	−1.78	0.39(2)
	(0.83)	(1.02)	(2.89)	(0.94)
BK	−0.66	1.83	−2.15	0.34(3)
	(1.68)	(2.03)	(2.88)	(0.84)
HP	−0.54	1.70	−1.84	0.06(1)
	(0.69)	(0.87)	(1.34)	(0.80)

Each equation includes a constant term and each instrument set includes a vector of ones. Instruments are $x_t, ..., x_{t-J+1}$and π_{t-1}, where J is the lag length estimated from the AIC. Estimation is by CU-GMM with a quadratic-spectral kernal HAC estimator.

these studies do not test for the role of conventionally measured output gaps in the standard NKPC. Conversely, there also is statistical work that, like ours, questions the links between measures of the output gap and inflation. For example, Orphanides and van Norden (2005) find that output gaps do not help forecast inflation when both are measured realistically in real time (rather than in revised data).

Marginal Costs Revisited

We have seen that the NKPC that uses the labor share to represent RULCs is relatively successful empirically. This measure of costs is easy to construct and has intuitive appeal. But some labor market arrangements imply that this measure is misspecified, so some recent research augments this model of marginal costs.

Macroeconomic models contain descriptions of the production technology that firms use. Models that contain different technologies will predict different ways to measure the marginal cost variable toward which firms adjust their prices. In particular, if a firm faces other frictions besides the costs of adjusting

Table 5 Augmenting the Labor Share

Study	Real Rigidity	Factor Added
Blanchard and Galí (2005)	Real wage stickiness	Unemployment rate
Christiano, Eichenbaum, and Evans (2005)	Multiple	Interest rate
Ravenna and Walsh (2006)	Financial friction	Interest rate
Chowdury, Hoffman, and Schabert (2006)	(cost channel)	
Krause, López-Salido, and Lubik (2008)	Labor market search	Hiring cost
Coenen, Levin, and Christoffel (2007)	Fixed and random length Calvo pricing	Inflation target
Batini, Jackson, and Nickell (2005) and Guerrieri, Gust, and López-Salido (2008)	Foreign competition	Import prices

prices, those may affect how it sets prices. For example, imagine a firm that must borrow from a bank to finance its wage bill. An increase in the interest rate it pays then will act like a cost shock and affect how it prices its goods. These additional frictions are sometimes called "real rigidities." They can include the financing constraint just mentioned, sticky real wages, or costs of hiring new employees.

Table 5 lists several recent studies that augment the labor-share measure of real unit labor costs with additional variables. Moreover, several of these studies estimate the NKPC by GMM with the revised measures of x_t and find statistical support for the added terms or right-hand-side variables. Few economists would argue that our model of firms' costs should be chosen according to how well it explains inflation in the NKPC, and these studies also examine other empirical evidence. But it is promising that a range of plausible modifications have improved the fit of the NKPC (its success in passing tests of overidentifying restrictions or stability tests) without significantly altering the findings about forward-looking and backward-looking weights.

7. MEASURES OF INFLATION

Conclusions about the NKPC also might depend on how the inflation rate is measured. The statistics so far have been based on the GDP deflator, so it seems natural to wonder whether they change if we measure inflation using another index such as the consumer price index (CPI), the deflator for personal consumption expenditure, or the producer price index. To check on the first of these alternatives, we average the monthly CPI (all items, all urban consumers, seasonally adjusted), `CPIAUCSL` from `FRED` to find the quarterly value, then construct the inflation rate as the annualized, quarter-to-quarter growth rate in percentage points.

Figure 5 shows this CPI inflation rate (the dashed gray line) and the inflation rate measured with the GDP deflator (the solid black line) used so far

Figure 5 Inflation in the GDP Deflator and CPI

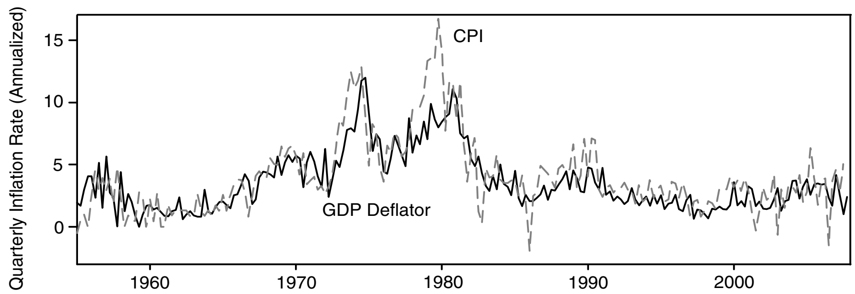

in this article. The figure shows a common, low-frequency cycle in the two measures of quarterly inflation. But the CPI inflation rate is more volatile. The only persistent difference between the two series occurred in the late 1970s when CPI inflation exceeded deflator inflation for several consecutive quarters.

When we estimate the NKPC with CPI inflation and RULC1 and RULC2, the results change modestly. The coefficient on lagged inflation, $\hat{\gamma}_b$, is slightly larger, and the coefficient on expected future inflation, $\hat{\gamma}_f$, is slightly smaller. The coefficient on marginal costs is smaller and is estimated less precisely. Finally, when we combine the CPI inflation rate with the seven output gaps, the results are quite negative for that approach, just as in the previous section. Overall, we conclude that the evidence summarized so far does not depend significantly on how the inflation rate is measured.

8. INTERNATIONAL EVIDENCE

Researchers also have used single-equation methods to study the NKPC in other countries. Galí, Gertler, and López-Salido (2001) find the hybrid NKPC fits well in quarterly Euro-area aggregate data for 1970–1998. As in the U.S. data, $\hat{\gamma}_f > \hat{\gamma}_b$, $\hat{\lambda}$ is statistically significant, and the J test does not reject overidentifying restrictions. Neiss and Nelson (2005) compare estimates of the NKPC for the United States, United Kingdom, and Australia. They also propose a new measure of the output gap that statistically explains inflation as well as measures of marginal costs. Leith and Malley (2007) estimate hybrid NKPCs for the G7 countries for 1960–1999, while Rumler (2007) does so for eight Euro-area countries for 1980–2003. Both of these studies discuss the role of the terms of trade, in addition to the labor share, in measures of marginal costs. They also report on differences in parameter estimates (such as those measuring price stickiness or inflation inertia) across countries. The international evidence on the hybrid NKPC may provide a guide to reform the measurement of marginal costs in open economies, in that the effects of foreign trade may be easier to detect in small, open economies than in the United States.

Batini, Jackson, and Nickell (2005) extend the model of marginal costs to reflect the relative price of imports, varying markups, and costs of adjusting employment. They arrive at a hybrid NKPC that nests the standard version and improves on its fit for the United Kingdom for 1972–1999. Bårsden, Jansen, and Nymoen (2004) also estimate more general statistical models of inflation, for the Euro area. They find that the hybrid NKPC can be improved on in terms of forecasting inflation even though it passes the J test.

Nason and Smith (2008) estimate the NKPC for the United Kingdom and Canada and provide tests that are robust to weak instruments. As in U.S. data, they find that the robust tests and traditional single-equation GMM estimation give different messages. The robust tests provide little evidence of forward-looking dynamics in these NKPCs. This international research thus conveys a similar message to the work on U.S. data.

9. CONCLUSION

This article outlines single-equation econometric methods for studying the NKPC and offers a progress report on the empirical evidence. How successful is the NKPC when estimated and tested on U.S. inflation? Enter the proverbial two-handed economist. On the one hand, the hybrid NKPC estimated by GMM on a quarterly 1955–2007 sample has coefficients that have signs and sizes that accord with economic theory and are statistically significant. The structural coefficients ($\hat{\omega}$, $\hat{\theta}$, and $\hat{\beta}$) are positive fractions, as are the reduced-form coefficients on past inflation and expected future inflation ($\hat{\gamma}_b$ and $\hat{\gamma}_f$), while the slope of the reduced-form Phillips curve ($\hat{\lambda}$) is positive. The hybrid NKPC also passes statistical tests based on the unpredictability of its residuals (the J test) and its stability over time (the sup-Wald test). The findings are not sensitive to alternative measures of inflation. Real unit labor costs are much better at statistically explaining inflation than are a plethora of output gap measures.

On the other hand, the t-statistic on real unit labor costs usually is not much above two. This indicates that there is not a close relationship between inflation and this measure of marginal costs. Estimates of the NKPC using surveys of forecasts give very different coefficients from those using instrumental-variables estimation. The confidence interval that is valid even with weak instruments gives a wide range of possible values for the parameter on expected future inflation. Moreover, other tests that are robust to weak identification often yield unreasonable values for the other hybrid NKPC parameters or reject the NKPC entirely. Our macroeconomic "Rip van Winkle," accustomed to the evidence against the Phillips curve garnered during the 1970s and 1980s, might find that the world has not changed much after all.

How will we learn more from single-equation methods? One promising and active avenue of research focuses on measurement of the cost variable, x_t, toward which prices adjust. Econometric tools for drawing inferences with weak identification also continue to advance. And the simple accumulation of macroeconomic data over time may help with precision, too.

Of course, systems methods of estimation also continue to be fruitful ways to identify and estimate the NKPC. Another complement to traditional, single-equation methods is to look at microeconomic data from individual firms or industries. Economists increasingly ask whether macroeconomic models of price stickiness are consistent with data on how prices are adjusted at the microlevel. It may be possible to measure cost shocks in microeconomic data and estimate pricing equations at that level, too.

The NKPC continues to be a key building block for macroeconomic models that require a monetary transmission mechanism. Our econometric work shows that marginal costs may be superior to many output gaps as a guide to inflation. We also obtain GMM estimates that give an important role to expected future inflation in explaining current inflation, while lagged inflation receives less weight. But measuring the effect of expected, future inflation on current inflation can be problematic because of weak instruments. Future research on this key response would be valuable because such forward-looking effects continue to have implications for the design of good monetary policy.

REFERENCES

Adam, Klaus, and Mario Padula. 2003. "Inflation Dynamics and Subjective Expectations in the United States." European Central Bank Working Paper No. 222.

Anderson, Theodore W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *Annals of Mathematical Statistics* 20: 46–63

Andrews, Donald W.K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59: 817–58.

Andrews, Donald W.K. 1993. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica* 61: 821–56.

Andrews, Donald W.K. 2003. "Tests for Parameter Instability and Structural Change with Unknown Change Point: A Corrigendum." *Econometrica* 71 (January): 395–7.

Andrews, Donald W.K., and James H. Stock. 2006. "Inference with Weak Instruments." In *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society,* vol. III, edited by R. Blundell, W.K. Newey, and T. Persson. Cambridge, UK: Cambridge University Press, 122–73.

Ang, Andrew, Geert Bekaert, and Min Wei. 2007. "Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?" *Journal of Monetary Economics* 54 (May): 1163–212.

Ascari, Guido. 2004. "Staggered Prices and Trend Inflation: Some Nuisances." *Review of Economic Dynamics* 7 (July): 642–67.

Bårsden, Gunnar, Eilev S. Jansen, and Ragnar Nymoen. 2004. "Econometric Evaluation of the New Keynesian Phillips Curve." *Oxford Bulletin of Economics and Statistics* 66: 671–86.

Basistha, Arabinda, and Charles R. Nelson. 2007. "New Measures of the Output Gap Based on the Forward-Looking New Keynesian Phillips Curve." *Journal of Monetary Economics* 54 (March): 498–511.

Batini, Nicoletta, Brian Jackson, and Stephen Nickell. 2005. "An Open-Economy New Keynesian Phillips Curve for the U.K." *Journal of Monetary Economics* 52 (September): 1061–71.

Baxter, Marianne, and Robert G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics* 81 (November): 575–93.

Blanchard, Olivier, and Jordi Galí. 2005. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit and Banking* 39: 35–65.

Brissimis, Sophocles N., and Nicholas S. Magginas. 2008. "Inflation Forecasts and the New Keynesian Phillips Curve." *International Journal of Central Banking* 4 (June): 1–22.

Calvo, Guillermo. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (September): 383–98.

Caner, Mehmet. 2007. "Boundedly Pivotal Structural Change Tests in Continuous Updating GMM with Strong, Weak Identification and Completely Unidentified Cases." *Journal of Econometrics* 137 (March): 28–67.

Chowdhury, Ibrahim, Mathias Hoffmann, and Andreas Schabert. 2006. "Inflation Dynamics and the Cost Channel of Monetary Transmission." *European Economic Review* 50 (May): 995–1016.

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

Cochrane, John H. 2001. *Asset Pricing*. Princeton, NJ: Princeton University Press.

Coenen, Günter, Andrew T. Levin, and Kai Christoffel. 2007. "Identifying the Influences of Nominal and Real Rigidities in Aggregate Price-Setting Behavior." *Journal of Monetary Economics* 54 (November): 2439–66.

Cogley, Timothy, and Argia M. Sbordone. Forthcoming. "Trend Inflation and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review.*

Cogley, Timothy, and James M. Nason. 1995. "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research." *Journal of Economic Dynamics and Control* 19: 253–78.

Dennis, Richard. 2007. "Fixing the New Keynesian Phillips Curve." Federal Reserve Bank of San Francisco Economic Letter Number 2007-35, November 30.

Dufour, Jean-Marie. 2003. "Identification, Weak Instruments, and Statistical Inference in Econometrics." *Canadian Journal of Economics* 36 (November): 767–808.

Dufour, Jean-Marie, Lynda Khalaf, Maral Kichian. 2006. "Inflation Dynamics and the New Keynesian Phillips Curve: An Identification-Robust Econometric Analysis." *Journal of Economic Dynamics and Control* 30: 1707–27.

Eichenbaum, Martin, and Jonas Fisher. 2007. "Estimating the Frequency of Price Re-Optimization in Calvo-Style Models." *Journal of Monetary Economics* 54 (October): 2032–47.

Fuhrer, Jeffrey C. 1997. "The (Un)Importance of Forward-Looking Behavior in Price Specifications." *Journal of Money, Credit and Banking* 29 (August): 338–50.

Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (October): 195–222.

Galí, Jordi, Mark Gertler, and David López-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45: 1237–70.

Galí, Jordi, Mark Gertler, and David López-Salido. 2005. "Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve." *Journal of Monetary Economics* 52 (September): 1107–18.

Guerrieri, Luca, Christopher Gust, and David López-Salido. 2008. "International Competition and Inflation: A New Keynesian

Perspective." Board of Governors of the Federal Reserve System, International Finance Discussion Paper No. 918 (January).

Hansen, Bruce E. 1997. "Approximate Asymptotic P-Values for Structural Change Tests." *Journal of Business and Economic Statistics* 15 (January): 60–7.

Hansen, Lars Peter, John Heaton, and Amir Yaron. 1996. "Finite-Sample Properties of Some Alternative GMM Estimators." *Journal of Business and Economic Statistics* 14 (July): 262–80.

Harvey, Andrew C. 2007. "Modeling the Philips Curve with Unobserved Components." Mimeo, Department of Economics, Cambridge University.

Harvey, Andrew C., and Thomas M. Trimbur. 2003. "General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series." *The Review of Economics and Statistics* 85: 244–55.

Henry, S. G. B., and Adrian R. Pagan. 2004. "The Econometrics of the New Keynesian Policy Model: Introduction." *Oxford Bulletin of Economics and Statistics* 66 (September): 581–607.

Hodrick, Robert J., and Edward C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking* 29 (February): 1–16.

Hornstein, Andreas. 2007. "Evolving Inflation Dynamics and the New Keynesian Phillips Curve." Federal Reserve Bank of Richmond *Economic Quarterly* 93 (Fall): 317–39.

Kleibergen, Frank, and Sophocles Mavroeidis. 2008. "Weak Instrument Robust Tests in GMM and the New Keynesian Phillips Curve." Mimeo, Department of Economics, Brown University.

Krause, Michael U., David López-Salido, and Thomas A. Lubik. 2008. "Inflation Dynamics with Search Frictions: A Structural Econometric Analysis." *Journal of Monetary Economics* 55 (July): 892–916.

Kurmann, André. 2005. "Quantifying the Uncertainty about a Forward-Looking, New Keynesian Pricing Model." *Journal of Monetary Economics* 52: 1119–34.

Kurmann, André. 2007. "VAR-Based Estimation of Euler Equations with an Application to New Keynesian Pricing." *Journal of Economic Dynamics and Control* 31 (March): 767–796.

Lacker, Jeffrey M., and John A. Weinberg. 2007. "Inflation and Unemployment: A Layperson's Guide to the Phillips Curve." Federal Reserve Bank of Richmond *Economic Quarterly* 93 (Summer): 201–27.

Leith, Campbell, and Jim Malley. 2007. "Estimated Open Economy New Keynesian Phillips Curves for the G7." *Open Economies Review* 18 (September): 405–26.

Lindé, Jesper. 2005. "Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach." *Journal of Monetary Economics* 52 (September): 1135–49.

Ma, Adrian. 2002. "GMM Estimation of the New Phillips Curve." *Economics Letters* 76 (August): 411–7.

Mavroeidis, Sophocles. 2005. "Identification Issues in Forward-Looking Models Estimated by GMM, with an Application to the Phillips Curve." *Journal of Money, Credit, and Banking* 37 (June): 421–48.

Morley, James C., Charles R. Nelson, and Eric Zivot. 2003. "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP so Different?" *Review of Economics and Statistics* 85: 235–43.

Nason, James M., and George A. Slotsve. 2004. "Along the New Keynesian Phillips Curve with Nominal and Real Rigidities." Federal Reserve Bank of Atlanta Working Paper 2004-09.

Nason, James M., and Gregor W. Smith. 2008. "Identifying the New Keynesian Phillips Curve." *Journal of Applied Econometrics* 23: 525–51.

Neiss, Katharine S., and Edward Nelson. 2005. "Inflation Dynamics, Marginal Cost, and the Output Gap: Evidence from Three Countries." *Journal of Money, Credit and Banking* 37 (December): 1019–45.

Newey, Whitney, and Kenneth D. West. 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies* 64 (October): 631–53.

Nunes, Ricardo. 2008. "Inflation Dynamics: The Role of Expectations." Mimeo, Federal Reserve Board of Governors.

Orphanides, Athanasios, and John C. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity* 33: 63–146.

Orphanides, Athanasios, and John C. Williams. 2005. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations." *Journal of Economic Dynamics and Control* 29 (November): 1927–50.

Orphanides, Athanasios, and Simon van Norden. 2005. "The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time." *Journal of Money, Credit and Banking* 37 (June): 583–601.

Ravenna, Federico, and Carl E. Walsh. 2006. "Optimal Monetary Policy with the Cost Channel." *Journal of Monetary Economics* 53 (March): 199–216.

Roberts, John M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27 (November): 975–84.

Rotemberg, Julio J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (October): 517–31.

Rudd, Jeremy, and Karl Whelan. 2005a. "Does Labor's Share Drive Inflation?" *Journal of Money, Credit and Banking* 37 (April): 297–312.

Rudd, Jeremy, and Karl Whelan. 2005b. "New Tests of the New Keynesian Phillips Curve." *Journal of Monetary Economics* 52 (September): 1167–81.

Rudd, Jeremy, and Karl Whelan. 2006. "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" *American Economic Review* 96 (March): 303–20.

Rumler, Fabio. 2007. "Estimates of the Open Economy New Keynesian Phillips Curve for Euro Area Countries." *Open Economies Review* 18 (September): 427–51.

Sbordone, Argia M. 2002. "Prices and Unit Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49: 235–56.

Sbordone, Argia M. 2005. "Do Expected Marginal Costs Drive Inflation Dynamics?" *Journal of Monetary Economics* 52: 1183–97.

Smith, Gregor W. 2007. "Pooling Forecasts in Linear Rational Expectations Models." Queen's Economics Department Working Paper 1129.

Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44 (October): 293–335.

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

Woodford, Michael. 2003. *Interest and Prices*. Princeton University Press: Princeton, NJ.

Woodford, Michael. 2007. "Interpreting Inflation Persistence: Comments on the Conference on 'Quantitative Evidence on Price Determination.'" *Journal of Money, Credit and Banking* 39: 203–10.

Zhang, Chengsi, Denise Osborn, and Dong Heon Kim. 2008. "The New Keynesian Phillips Curve: From Sticky Inflation to Sticky Prices." *Journal of Money, Credit and Banking* 40: 667–99.

Economic Quarterly—Volume 94, Number 4—Fall 2008—Pages 397–433

DSGE Model-Based Estimation of the New Keynesian Phillips Curve

Frank Schorfheide

An important building block in modern dynamic stochastic general equilibrium (DSGE) models is the price-setting equation for firms. In models in which the adjustment of nominal prices is costly, this equation links inflation to current and future expected real marginal costs and is typically referred to as the New Keynesian Phillips curve (NKPC). Its most popular incarnation can be derived from the assumption that firms face quadratic nominal price adjustment costs (Rotemberg 1982) or that firms are unable to re-optimize their prices with a certain probability in each period (Calvo 1983). The Calvo model has a particular appeal because it generates predictions about the frequency of price changes, which can be measured with microeconomic data (Bils and Klenow 2004, Klenow and Kryvtsov 2008). The slope of the NKPC is important for the propagation of shocks and determines the output-inflation tradeoff faced by policymakers. The Phillips curve relationship can also be used to forecast inflation.

This article reviews estimates of NKPC parameters that have been obtained by fitting fully specified DSGE models to U.S. data. By now, numerous empirical papers estimate DSGE models with essentially the same NKPC specification. In this literature, the Phillips curve implies that inflation can be expressed as the discounted sum of expected future marginal costs, where marginal costs equal the labor share. We document that the identification of

Ed Herbst and Maxym Kryshko provided excellent research assistance. I am thankful to Michael Dotsey, Andreas Hornstein, Thomas Lubik, and James Nason for helpful comments and suggestions. Support from the Federal Reserve Bank of Richmond and the National Science Foundation (Grant SES 0617803) is gratefully acknowledged. The views expressed in this article do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: schorf@ssc.upenn.edu.

the Phillips curve coefficients is tenuous and no consensus about its slope and the importance of lagged inflation has emerged from the empirical studies.

We begin by examining how the NKPC parameters are identified in a DSGE model-based estimation. This is a difficult question. Many estimates are based on a likelihood function, which is the model-implied probability distribution of a set of observables indexed by a parameter vector. The likelihood function peaks at parameter values for which the model-implied autovariance function of a vector of macroeconomic time series matches the sample autocovariance function. Unfortunately, this description is not particularly illuminating. More intuitively, the NKPC parameters are estimated by a regression of inflation on the sum of discounted future expected marginal costs. The likelihood function corrects the bias that arises from the endogeneity of the marginal cost regressor. We show that if one simply uses ordinary least-squares (OLS) to regress inflation on measures of expected marginal costs, the slope coefficient is very close to zero. This finding is quite robust to the choice of detrending method and marginal cost measure. Hence, much of the variation in the estimates reported in the literature is due to the multitude of endogeneity corrections that arise by fitting different DSGE models that embody essentially the same Phillips curve specification.

The review of empirical studies distinguishes between papers in which marginal costs are included in the observations and, hence, are directly used in the estimation and studies that treat marginal costs as a latent variable. In the latter case, NKPC estimates are more sensitive to the specification of the households' behavior, the conduct of monetary policy, and the law of motion of the exogenous disturbances. Estimates of the slope of the Phillips curve lie between 0 and 4. If the list of observables spans the labor share, then the slope estimates fall into a much narrower range of 0.005 to 0.135. No consensus has emerged with respect to the importance of lagged inflation in the Phillips curve. We compare estimates of the relative movement of inflation and output in response to a monetary policy shock, which captures an important tradeoff for monetary policymakers. We find that the estimates in the studies that are surveyed in this article range from 0.07 to 1.4. A value of 0.07 (1.4) implies that a 1 percent increase in output due to a monetary policy shock is accompanied by a quarter-to-quarter inflation rate of 7 (140) basis points.

The remainder of this paper is organized as follows. We discuss the derivation of the NKPC as well as our concept of DSGE model-based estimation in Section 1. In Section 2, a simple DSGE model that can be solved analytically is used to characterize various sources of NKPC parameter identification. Any particular DSGE model-based estimation might exploit some or all of these sources of information. Section 3 provides empirical evidence from least-squares regressions of inflation on the discounted sum of future marginal costs as well as evidence from a vector autoregression (VAR) on the relative movement of output and inflation in response to a monetary policy shock. We thereby characterize some features of the data that are important for understanding the DSGE model-based parameter estimates reviewed in Section 4. Finally, Section 5 concludes.

1. PRELIMINARIES

This section begins with a brief description of the price-setting problem that gives rise to a Phillips curve in New Keynesian DSGE models. We then discuss some of the defining characteristics of DSGE model-based estimation of NKPC parameters.

Price Setting in DSGE Models

New Keynesian DSGE models typically assume that production is carried out by two types of firms: final good producers and intermediate goods producers. The latter hire labor and capital services from the households to produce a continuum of intermediate goods. The final good producers purchase the intermediate goods and bundle them into a single aggregate good that can be used for consumption or investment. The intermediate goods are imperfect substitutes and, hence, each producer faces a downward-sloping demand curve. Price stickiness is introduced by assuming that it is costly to change nominal prices. Rotemberg (1982) assumed that the price adjustment costs are quadratic, whereas Calvo (1983) set forth a model of staggered price setting in which the costs are either zero or infinite with fixed probabilities, i.e., only a fraction of firms is able to change or, more precisely, re-optimize prices.

Aggregating the optimal price-setting decisions of firms leads to the following expression for inflation in the price of the final good, referred to as the New Keynesian Phillips curve:

$$\tilde{\pi}_t = \gamma_b \tilde{\pi}_{t-1} + \gamma_f \mathrm{E}_t \left[\tilde{\pi}_{t+1} \right] + \lambda \widetilde{MC}_t + \tilde{\xi}_t. \tag{1}$$

Here $\widetilde{\pi}_t$ represents inflation, $\widetilde{MC}_t$ is real marginal costs, and $\widetilde{\xi}_t$ is an exogenous disturbance that is often called a mark-up shock. We use $\widetilde{z}_t$ to denote percentage deviations of a variable, z_t, from its steady state. The coefficients γ_b, γ_f, and λ are functions of model-specific taste and technology parameters. For instance, in Calvo's (1983) model of price stickiness

$$\gamma_b = \frac{\omega}{1+\beta\omega}, \quad \gamma_f = \frac{\beta}{1+\beta\omega}, \text{ and } \quad \lambda = \frac{(1-\zeta)(1-\zeta\beta)}{\zeta(1+\beta\omega)},$$

where β is the households' discount factor and ζ is the probability that an intermediate goods producer is unable to re-optimize its price in the current period. In the derivation of (1), it was assumed that those firms that are unable to re-optimize their prices either adjust their past price by the steady-state inflation rate or by lagged inflation. The parameter ω represents the fraction of firms that indexes their prices to lagged inflation.

Assuming that $\beta = 0.99$, the sum of γ_b and γ_f is slightly less than 1 and the coefficient of lagged inflation lies between 0 (no dynamic indexation, $\omega = 0$) and 0.5 (full dynamic indexation, $\omega = 1$). If $\omega = 0$ and steady-state inflation is 0, then $1/(1-\zeta)$ can be interpreted as the expected duration between price changes. For instance, $\zeta = \frac{2}{3}$ implies that the expected duration of a price set by an intermediate goods producer is three quarters, which leads to a slope coefficient of $\lambda = 0.167$. On the other hand, if $\zeta = \frac{7}{8}$, which means that the duration of a price is eight quarters, then the NKPC is much flatter: $\lambda = 0.018$.

Our survey of the empirical literature will focus on coefficient estimates for γ_b, γ_f, and λ rather than the model-specific preference-and-technology parameters. The slope, λ, determines the output-inflation tradeoff faced by central banks and affects, for instance, the relative response of output and inflation in response to an unanticipated monetary policy shock. A detailed exposition of the role that the NKPC plays in the analysis of monetary policy is provided in an article by Stephanie Schmitt-Grohé and Martín Uribe in this issue. The coefficient γ_b affects the persistence of inflation and, for instance, the rate at which inflation effects of shocks to marginal costs die out. This is an important parameter, particularly for central banks that pursue a policy of inflation targeting. If we rearrange the terms in (1), such that expected inflation appears on the left-hand side and all other terms on the right-hand side, then the Phillips curve delivers a forecasting equation for inflation.

DSGE Model-Based Estimation

This article focuses on estimates of γ_b, γ_f, and λ that are obtained by exploiting the full structure of a model economy. Thus, we consider approaches in which the researcher solves not only the decision problems of the firms but also those of the other agents in the economy and imposes an equilibrium concept. If the economy is subject to exogenous stochastic shocks, the DSGE model generates a joint probability distribution for time series such as aggregate output, inflation, and interest rates. Suppose we generically denote the vector of time, t, observables by x_t and assume that the DSGE model has been solved by log-linear approximation techniques. Then the equilibrium law of motion takes the form of a vector autoregressive moving average (VARMA) process of the form (omitting deterministic trend components)

$$x_t = \Phi_1 x_{t-1} + \ldots \Phi_p x_{t-p} + R\epsilon_t + \Psi_1 R\epsilon_{t-1} + \ldots \Psi_q R\epsilon_{t-q}. \quad (2)$$

The matrices Φ_i, Ψ_j, and R are complicated functions of the Phillips curve parameters γ_b, γ_f, and λ, as well as the remaining DSGE model parameters, which we will summarize by the vector θ. The vector ϵ_t stacks the innovations to all exogenous stochastic disturbances and is often assumed to be normally and independently distributed.

A natural approach of exploiting (2) is likelihood-based estimation. Maximum likelihood (ML) estimation of optimization-based rational expectations models in macroeconomics dates back at least to Sargent (1989) and has been widely applied in the DSGE model literature (e.g., Altug [1989], Leeper and Sims [1994], and many of the papers reviewed in Section 4). The likelihood function is defined as the joint density of the observables conditional on the parameters, which can be derived from (2). Let $X^t = \{x_1, \ldots, x_t\}$, then

$$p(X^T|\gamma_b, \gamma_f, \lambda, \theta) = p(x_1|\gamma_b, \gamma_f, \lambda, \theta) \prod_{t=2}^{T} p(x_t|X^{t-1}, \gamma_b, \gamma_f, \lambda, \theta). \quad (3)$$

The evaluation of the likelihood function typically requires the use of numerical methods to solve for the equilibrium dynamics and to integrate out unobserved elements from the joint distribution of the model variables (see, for instance, An and Schorfheide [2007]). A numerical optimization routine can then be used to find the maximum of the (log-)likelihood function. The potential drawback of the ML approach is that identification problems can make it difficult to find the maximum of the likelihood function and render standard large sample approximations to the sampling distribution of the ML estimator and likelihood ratio statistics inaccurate.

A popular alternative to the frequentist ML approach is Bayesian inference. Bayesian analysis tends to interpret the likelihood function as a density function for the parameters given the data. Let $p(\gamma_b, \gamma_f, \lambda, \theta)$ denote a prior density for the DSGE model parameters. Bayesian inference is based on the posterior distribution characterized by the density

$$p(\gamma_b, \gamma_f, \lambda, \theta|X^T) = \frac{p(X^T|\gamma_b, \gamma_f, \lambda, \theta)p(\gamma_b, \gamma_f, \lambda, \theta)}{\int p(X^T|\gamma_b, \gamma_f, \lambda, \theta)p(\gamma_b, \gamma_f, \lambda, \theta)d(\gamma_b, \gamma_f, \lambda, \theta)}. \quad (4)$$

Notice that the denominator does not depend on the parameters and simply normalizes the posterior density so that it integrates to one. The controversial ingredient in Bayesian inference is the prior density as it alters the shape of the posterior, in particular if the likelihood function does not exhibit much curvature. On the upside, the prior allows the researcher to incorporate additional information in the time series analysis that can help sharpen inference. Many of the advantages of Bayesian inference in the context of DSGE model estimation are discussed in Lubik and Schorfheide (2006) and An and Schorfheide (2007). The implementation of Bayesian inference typically relies on Markov-chain Monte Carlo methods that allow the researcher to generate random draws of the model parameters from their posterior distribution. These draws can then be transformed—one by one—into statistics of interest. Sample moments computed from these draws provide good approximations to the corresponding population moments of the posterior distribution.

Notwithstanding all the desirable statistical properties of likelihood-based estimators, the mapping of particular features of the data into parameter estimates is not particularly transparent. Superficially, the likelihood function peaks at parameter values for which a weighted discrepancy between DSGE model-implied autocovariances of x_t and sample autocovariances is minimized. The goal of the next section is to explore the extent to which this matching of autocovariances can identify the parameters of the New Keynesian Phillips curve.

2. IDENTIFYING THE NKPC PARAMETERS

The identification of DSGE model parameters through likelihood-based methods tends to be a black box because the relationship between structural parameters and autocovariances or other reduced-form representations is highly nonlinear. This section takes a look inside this black box to develop some understanding about particular features of the DSGE model that contribute to the identifiability of NKPC parameters. Rather than asking whether there is enough variation in postwar data to estimate the NKPC parameters reliably, for now we focus on sources of identification in infinite samples. In practice, the estimation of a particular model might exploit several of these sources of information simultaneously.

Since the Phillips curve provides a relationship between marginal costs and inflation, the measurement of marginal costs is important for the identification of the NKPC parameters. A key feature of likelihood-based inference—as opposed to the single-equation methods reviewed by James Nason and Gregor Smith in this issue—is the exploitation of model-implied restrictions of contemporaneous correlations between variables, as well as the use of information from impulse responses. In many instances, higher-order autocovariances of inflation and marginal costs are an additional source of information.

While this section focuses on identifying the slope, λ, we also offer some insights into identifying γ_b and γ_f. For now we assume that $\gamma_b = 0$. In the context of the Calvo model this assumption implies that the fraction, ω, of firms that engage in dynamic indexation is zero. In this case, $\gamma_f = \beta$. Since β in a fully specified DSGE model is related to the steady-state real interest rate, the coefficient γ_f can be determined, for instance, by averaging interest rate data, and its identification is not a concern. Under our simplifications, the Phillips curve takes the form

$$\widetilde{\pi}_t = \beta \mathrm{E}_t[\widetilde{\pi}_{t+1}] + \lambda \widetilde{MC}_t + \widetilde{\xi}_t. \tag{5}$$

Solving this difference equation forward we find that today's inflation is a function of future expected marginal costs:

$$\widetilde{\pi}_t = \sum_{j=0}^{\infty} \beta^j \mathrm{E}_t[\lambda \widetilde{MC}_{t+j} + \widetilde{\xi}_{t+j}]. \tag{6}$$

Observed Versus Latent Marginal Costs

The identification of λ crucially depends on whether real marginal costs are treated as directly observable or as a latent variable. If $\widetilde{MC}_t$ is directly observed and, hence, is an element of the vector x_t in (2) and (4), then the main obstacle to the identification of λ is the endogeneity problem caused by the potential correlation between the mark-up shock, ξ_t, and marginal costs. The estimation of future expected marginal costs in (6) poses no real challenge because $\mathrm{E}_t[\widetilde{MC}_{t+j}]$ can be obtained from the reduced-form representation associated with the law of motion (2), which is always identifiable. The downside of including a direct measure of marginal costs in the set of observables is that measurement errors pertaining to the marginal cost series can potentially distort the inference about the NKPC parameters. Yet, identifying λ is more tenuous if marginal costs are not included in the vector x_t.

To make the discussion more concrete, imagine an economy in which labor is the only factor of production and, in log-linear terms,

$$\widetilde{Y}_t = \widetilde{Z}_t + \widetilde{H}_t.$$

Z_t is an unobserved total factor productivity process and H_t is hours worked. Marginal costs are given by

$$\widetilde{MC}_t = \widetilde{W}_t - \widetilde{Z}_t,$$

where W_t are wages. Moreover, suppose that the households' instantaneous utility function is of the form

$$U(C_t, H_t) = \frac{C_t^{1-1/\tau}}{1-1/\tau} - \phi H_t,$$

and μ_t denotes the marginal utility of consumption. Under these preferences labor supply is infinitely elastic, the wage has to satisfy $W_t = 1/\mu_t$, and the marginal utility of consumption is given by $\mu_t = C_t^{-1/\tau}$. Finally, assume that output is entirely used for household consumption such that $C_t = Y_t$. Then we obtain the following link between marginal costs and output:

$$\widetilde{MC}_t = \frac{1}{\tau}(\widetilde{Y}_t - \tau \widetilde{Z}_t).$$

If the vector of observables, x_t, contains output, wages, and hours worked, then the marginal costs are directly observed because

$$\widetilde{MC}_t = \widetilde{lsh}_t = \widetilde{W}_t + \widetilde{H}_t - \widetilde{Y}_t.$$

More generally, in models with Cobb-Douglas technology the vector x_t spans marginal costs as long as one can construct the labor share, $\widetilde{lsh}_t$, from the observables. If, however, the vector x_t only contains observations on output in addition to inflation and interest rates, then marginal costs are latent because they depend on the observed output as well as the unobserved technology process, $\widetilde{Z}_t$, and the unknown parameter, τ. Rewriting (5) in terms of inflation and output yields

$$\widetilde{\pi}_t = \beta \mathrm{E}_t[\widetilde{\pi}_{t+1}] + \frac{\lambda}{\tau}\widetilde{Y}_t - \widetilde{Z}_t + \widetilde{\xi}_t.$$

Two challenges arise. First, the presence of $\widetilde{Z}_t$ exacerbates the endogeneity problem that arises in the NKPC estimation. Moreover, the coefficient associated with $\widetilde{Y}_t$ in itself does not identify the original slope parameter, λ, since it also depends on the utility function parameter, τ, which needs to be identified from other equilibrium relationships.

In practice, likelihood-based estimation of DSGE models relies on the so-called state-space representation of the DSGE model, rather than the VARMA representation in (2). Omitting deterministic trend components, the state-space representation takes the form

$$x_t = A s_t, \quad s_t = B_1 s_{t-1} + B_\epsilon \epsilon_t, \tag{7}$$

where x_t is the vector of observables, s_t is a vector of latent variables, and the matrices A, B_1, and B_ϵ are functions of the DSGE model parameters. The likelihood function associated with (7) can be computed with the Kalman filter. If the information in the vector x_t does not span marginal costs directly, then the Kalman filter constructs an estimate of the latent marginal costs (and technology, $\widetilde{Z}_t$, in our example) based on x_t and the parameters λ and θ. To the extent that the Kalman filter inference for the latent variables is sensitive to the assumed law of motion of the unobserved exogenous processes, inference about the slope of the Phillips curve is also sensitive to these auxiliary assumptions.

Identifying Information in Contemporaneous Correlations

Fully-specified DSGE models impose strong restrictions on the contemporaneous interactions of macroeconomic variables. We will show in the context of a simple example that these restrictions enter the likelihood function and potentially provide important identifying information that is not used in the single-equation approaches reviewed by James Nason and Gregor Smith in this issue. For the remainder of Section 2 we adopt the convention that all variables are measured in percentage deviations from a deterministic steady state and omit tildes to simplify the notation.

Consider the log-linear approximation of the Euler equation associated with the households' problem in the previous subsection:

$$Y_t = \mathrm{E}_t[Y_{t+1}] - \tau(R_t - \mathrm{E}_t[\pi_{t+1}]) + \epsilon_{\phi,t}. \tag{8}$$

$R_t - \mathrm{E}_t[\pi_{t+1}]$ is the expected real return from holding a one-period nominal bond. The parameter, τ, can be interpreted as the intertemporal substitution elasticity of the household and $\epsilon_{\phi,t}$ is an exogenous preference shifter. To complete the model, we characterize monetary policy by an interest rate feedback rule of the form

$$R_t = \psi \pi_t + \epsilon_{R,t}, \tag{9}$$

where $\epsilon_{R,t}$ is a monetary policy shock.

We now substitute the marginal cost expression derived in the previous subsection into the NKPC and obtain

$$\pi_t = \beta \mathrm{E}_t[\pi_{t+1}] + \frac{\lambda}{\tau}(Y_t - \tau Z_t) + \xi_t. \tag{10}$$

Since the unobserved technology shock, Z_t, and the mark-up shock, ξ_t, affect the equilibrium law of motion in a similar manner in this simple model, we set $Z_t = 0$ and let $\xi_t = \epsilon_{\xi,t}$. Moreover, we define $\kappa = \frac{\lambda}{\tau}$ and will direct our attention to the estimation of the output inflation tradeoff, κ, rather than λ. Thus, we are essentially abstracting from the two additional difficulties that arise if marginal costs are treated as a latent variable. Finally, it is assumed that the three exogenous shocks, $\epsilon_{R,t}$, $\epsilon_{\phi,t}$, and $\epsilon_{\xi,t}$ are independently and identically distributed zero mean normal random variables with standard deviations σ_R, σ_ϕ, and σ_ξ, respectively.

The linear rational expectations (LRE) model comprised of (8) to (10) can be solved with standard methods such as the one described in Sims (2002). To ensure that the LRE system has a unique stable solution, we impose $\psi > 1$, which implies that the central bank raises the real interest rate in response to an inflation rate that exceeds its steady-state level. Lubik and Schorfheide (2004) show that the equilibrium law of motion for the three observables is of the form

$$\begin{bmatrix} Y_t \\ \pi_t \\ R_t \end{bmatrix} = \frac{1}{1+\kappa\tau\psi} \begin{bmatrix} -\tau & 1 & -\tau\psi \\ -\kappa\tau & \kappa & 1 \\ 1 & \kappa\psi & \psi \end{bmatrix} \begin{bmatrix} \epsilon_{R,t} \\ \epsilon_{\phi,t} \\ \epsilon_{\xi,t} \end{bmatrix}. \tag{11}$$

Since our model lacks both endogenous and exogenous propagation mechanisms, output, inflation, and interest rates—the three variables observed by the econometrician—are serially uncorrelated in equilibrium. Thus, all the information about the slope of the Phillips curve must come from the contemporaneous correlations among the three observables.

The single-equation approach to the estimation of the NKPC reviewed by James Nason and Gregor Smith in this issue can be interpreted in two ways. First, one can write the NKPC as a regression of the form

$$\pi_{t+1} = \frac{1}{\beta}\pi_t - \frac{\kappa}{\beta}Y_t - \frac{1}{\beta}\eta_{t+1} - \frac{1}{\beta}\epsilon_{\xi,t} = \alpha_1\pi_t + \alpha_2 Y_t + resid_{t+1}. \tag{12}$$

Here we replaced the conditional expectation of inflation, $\mathrm{E}_t[\pi_{t+1}]$, by π_{t+1} and a forecast error $\eta_{t+1} = \pi_{t+1} - \mathrm{E}_t[\pi_{t+1}]$. The lack of serial correlation in the equilibrium dynamics implies that least-squares estimates of α_1 and α_2 converge in probability to zero. Hence, based on a large sample, an econometrician concludes that the slope of the Phillips curve is zero. The estimation of (12) with an instrumental variable estimator that tries to correct a potential bias due to the correlation between Y_t and $\epsilon_{\xi,t}$ is also bound to fail because in equilibrium, lagged values of output and inflation are uncorrelated with the regressors.

Alternatively, one can express the Phillips curve as a regression of the form

$$\pi_t = \alpha_1\mathrm{E}_t[\pi_{t+1}] + \alpha_2 Y_t + resid_t. \tag{13}$$

However, even if the econometrician realizes that $\mathrm{E}_t[\pi_{t+1}] = 0$ and excludes the expected inflation regressor, it is not possible to estimate the slope of the Phillips curve consistently. The least-squares estimator of α_2 provides a biased estimate of κ because of the correlation between output and the mark-up shock, which is subsumed in the residual. Instrumental variable estimation is also uninformative because lagged endogenous variables are uncorrelated with current output. Notice that this failure of single-equation estimation is not directly apparent from (10). It is a consequence of the auxiliary assumptions about the other sectors in the economy and the law of motion of the exogenous disturbances. Nason and Smith (2008) show that the identification problems associated with single-equation methods prevail, even if the DSGE model is enriched with serially correlated exogenous disturbances.

DSGE model-based estimation of the Phillips curve parameters utilizes the information in the contemporaneous relationship between output, inflation, and interest rates.[1] Let $\theta = [\tau, \psi, \sigma_R, \sigma_\phi, \sigma_\xi]'$ and factorize the joint density of the observables as

[1] It is tempting to check identification by comparing the number of structural parameters to the number of free parameters in the covariance matrix of Y_t, π_t, and R_t. In the DSGE model described by the law of motion (11), the two parameter counts are equal to six. Unfortunately, having at least as many estimable reduced-form parameters as structural parameters is neither sufficient for identification, nor does it provide interesting insights about the sources of identification.

$$p(Y_t, \pi_t, R_t|\kappa, \theta) = p(Y_t|\kappa, \theta)p(\pi_t|Y_t, \kappa, \theta)p(R_t|Y_t, \pi_t, \theta). \tag{14}$$

The first term represents the marginal density of output and the third term is generated by the monetary policy rule. Key to understanding the DSGE model-based estimation of κ is the second term, that is, the conditional distribution of inflation given output. Since all the shocks are normally distributed,

$$\pi_t|Y_t \sim \mathcal{N}\left(\mathrm{E}\left[\pi_t \mid Y_t\right] \ , \ var[\pi_t|Y_t]\right)$$

and we can focus our attention on the conditional mean and variance.

We begin with the derivation of $\mathrm{E}[\pi_t \mid Y_t]$. Solving the Phillips curve relationship forward as in (6) leads to

$$\pi_t = \kappa Y_t + \epsilon_{\xi,t}. \tag{15}$$

Taking expectations conditional on Y_t of the left-hand side and right-hand side of (15) yields

$$\mathrm{E}\left[\pi_t \mid Y_t\right] = \kappa Y_t + \mathrm{E}\left[\epsilon_{\xi,t} \mid Y_t\right].$$

Using (11) and the formula for the conditional moments of a joint normal distribution,[2] we obtain

$$\mathrm{E}\left[\epsilon_{\xi,t} \mid Y_t\right] = \mu_{\xi|y}(\theta)Y_t = -\frac{1}{\tau\psi}\underbrace{\frac{\tau^2\psi^2\sigma_\xi^2}{\tau^2\sigma_R^2 + \sigma_\phi^2 + \tau^2\psi^2\sigma_\xi^2}}_{sh(\sigma_y^2, \epsilon_\xi)} Y_t. \tag{16}$$

The conditional expectation depends on the intertemporal elasticity of substitution, the policy rule coefficient, and all the shock variances. Here $sh(\sigma_y^2, \epsilon_\xi)$ is the fraction of the variance of output that is due to the mark-up shock, $\epsilon_{\xi,t}$. We now turn to the calculation of the conditional variance of inflation. Notice that $var[\pi_t|Y_t] = var[\epsilon_{\xi,t}|Y_t]$. Thus,

$$var[\pi_t|Y_t] = \sigma^2_{\xi|y}(\kappa, \theta) = \sigma_\xi^2 - \frac{(\tau\psi\sigma_\xi^2)^2}{(1 + \kappa\tau\psi)(\tau^2\sigma_R^2 + \sigma_\phi^2 + \tau^2\psi^2\sigma_\xi^2)}.$$

We deduce that

$$\begin{aligned} p(\pi_t|Y_t, \kappa, \theta) \quad \propto \quad & |\sigma^2_{\xi|y}(\kappa, \theta)|^{-1/2}\exp \\ & \left\{-\frac{1}{2\sigma^2_{\xi|y}(\kappa, \theta)}\left(\pi_t - [\kappa + \mu_{\xi|y}(\theta)]Y_t\right)^2\right\}, \end{aligned} \tag{17}$$

[2] Suppose that X and Y are jointly normally distributed with means μ_x and μ_y, variances v_{xx} and v_{yy}, and covariance v_{xy}; then the conditional mean $\mathrm{E}[X \mid Y = y] = \mu_x + v_{xy}v_{yy}^{-1}(y - \mu_y)$ and the conditional variance is $var[X|Y = y] = v_{xx} - v_{xy}^2/v_{yy}$.

where $\propto$ denotes proportionality.

We can draw several important conclusions from (17). First, the term $\mu_{\xi|y}$ given in (16) corrects for the endogeneity bias that arises in a regression of inflation and marginal costs. Suppose we set $\psi = 1.5$, which is Taylor's (1993) value, assume that $\tau = \frac{2}{3}$, which makes the agents slightly more risk-averse than agents with log preferences, and assume that 20 percent of the variation in output is due to mark-up or cost-push shocks. Then (16) implies that a simple least-squares regression of inflation on marginal costs, i.e. output, in our example model, would underestimate the slope, κ, by 0.2. Second, (17) implies that knowledge of the conditional distribution of inflation given output does not identify the slope of the Phillips curve. Moreover, the joint distribution of output and inflation is also not sufficient, because the marginal distribution of output only provides information about the variance of output, $\sigma_y^2(\kappa, \theta)$, which is insufficient to disentangle the values of all the θ elements. We will show below, however, that κ is identifiable with knowledge of the monetary policy reaction function.

To summarize, our simple example has a number of startling implications. First, a single-equation estimation based on (12) or (13) is unable to deliver a consistent estimate of κ. Second, an OLS regression of inflation on the sum of discounted future expected marginal costs generates a downward-biased estimate of κ. The magnitude of the bias is a function of central bank behavior, households' preferences, and, more generally, the importance of mark-up shocks for output fluctuations. Third, DSGE model-based estimation is promising but might require a prior that is informative about other model parameters, for instance those that control the law of motion of exogenous shocks or the conduct of monetary policy. We will subsequently elaborate on this last point.

Identifying Information in Impulse Response Functions

If the DSGE model embodies enough restrictions to identify a structural shock other than ξ_t from the observables, then one can potentially infer the Phillips curve slope from the impulse response function (IRF) associated with this shock. Consider the model analyzed in the previous subsection. Suppose that the policy rule coefficient, ψ, is known, which means that the sequence of monetary policy shocks can be directly obtained from interest rate and inflation data: $\epsilon_{R,t} = R_t - \psi\pi_t$. Recall from (15) that the forward solution of the Phillips curve takes the form

$$\pi_t = \kappa Y_t + \epsilon_{\xi,t}.$$

We previously showed that the correlation between the mark-up shock, $\epsilon_{\xi,t}$, and the regressor, Y_t, creates an endogeneity problem that complicates the identification of κ. The monetary policy shock can serve as an instrumental variable in the identification of κ. By assumption, the monetary policy shock is uncorrelated with $\epsilon_{\xi,t}$ but correlated with the regressor Y_t.

The argument can be formalized as follows. Suppose we factorize the likelihood function into[3]

$$\begin{aligned} p(Y_t, \pi_t, R_t|\kappa, \theta) &= p(R_t - \psi\pi_t|\kappa, \theta)p(Y_t|R_t - \psi\pi_t, \kappa, \theta) \\ & \quad p(\pi_t|Y_t, R_t - \psi\pi_t, \kappa, \theta). \end{aligned} \tag{18}$$

$R_t - \psi\pi_t$ measures the monetary policy shock, $\epsilon_{R,t}$, and the first term corresponds to its density. The second factor captures the distribution of output given the monetary policy shock. The third conditional density represents the Phillips curve. From this factorization it is apparent that, in a linear Gaussian environment, the following conditional expectations (we replace $R_t - \psi\pi_t$ by $\epsilon_{R,t}$) are identifiable:

$$\begin{aligned} \mathrm{E}\left[Y_{t|}R_t - \psi\pi_t, \kappa, \theta\right] &= \alpha_{11}\epsilon_{R,t} \text{ and} \\ \mathrm{E}\left[\pi_t \mid R_t - \psi\pi_t, Y_t\right] &= \alpha_{21}\epsilon_{R,t} + \alpha_{22}Y_t, \end{aligned}$$

where α_{ij} is a function of κ and θ. Since

$$\frac{\partial Y_t}{\partial \epsilon_{R,t}} = \alpha_{11}, \quad \frac{\partial \pi_t}{\partial \epsilon_{R,t}} = \alpha_{21} + \alpha_{22}\alpha_{11},$$

it follows from (11) that κ is identified by the ratio of the output and inflation response $\alpha_{21}/\alpha_{11} + \alpha_{22}$.

In our simple example the identification of the monetary policy shock depends on the assumed knowledge of the parameter ψ, which the reader might find unconvincing. More interestingly, there are a number of papers that estimate DSGE models that are specified such that monetary policy shocks can be identified from exclusion restrictions. Most notably, Rotemberg and Woodford (1997), Christiano, Eichenbaum, and Evans (2005), and Boivin and Giannoni (2006) consider models in which the private sector is unable to respond to monetary policy shocks contemporaneously.[4] In a Gaussian vector autoregressive system, this exclusion restriction is sufficient to identify monetary policy shocks and the associated impulse response functions independently of the DSGE model parameters.

---

[3] The Jacobian associated with the transformation of $[R_t - \psi\pi_t, Y_t, \pi_t]'$ into $[R_t, Y_t, \pi_t]'$ is equal to one. We maintain that θ is defined as $\theta = [\tau, \psi, \sigma_R, \sigma_\phi, \sigma_\xi]'$ and, hence, includes ψ.

[4] Rather than conducting likelihood-based inference, all three papers use an estimation method that exclusively relies on the identification of model parameters from IRF dynamics. The structural parameters are directly estimated by minimizing the discrepancy between the model-implied impulse responses to a monetary policy shock and those obtained from estimating a structural VAR.

Identifying Information in the Reduced-Form Dynamics

The absence of equilibrium dynamics in (11) is clearly at odds with reality. Aggregate output, inflation, and interest rates tend to exhibit fairly strong serial correlation. This serial correlation opens up another avenue for identification as lagged endogenous variables can serve as instruments to correct endogeneity biases. In fact, it is this serial correlation that single-equation approaches rely on.

Suppose that the vector x_t contains inflation, a measure of marginal costs as well as other variables, denoted by z_t: $x_t = [\pi_t, MC_t, z_t']'$. Moreover, assume that the mark-up shock, ξ_t, is independently distributed and that the DSGE model-implied law of motion for x_t has a VAR(1) representation:

$$x_t = \Phi_1(\lambda, \theta)x_{t-1} + u_t, \quad \text{where} \quad u_t = R(\lambda, \theta)\epsilon_t. \tag{19}$$

The matrices Φ_1 and R are functions of the DSGE model parameters, the vector ϵ_t stacks the innovations to the exogenous driving processes of the model economy, and u_t can be interpreted as reduced-form, one-step-ahead forecast errors. While the assumption that ξ_t is serially uncorrelated is crucial for the subsequent argument, the VAR(1) representation is not.

Define the selection vectors M_1 and M_2 such that $M_1'x_t = \pi_t$ and $M_2x_t = MC_t$. Equation (15) implies that the slope of the Phillips curve has to solve the following restriction:

$$M_1{\prime}\Phi_1 x_t - \lambda M_2{\prime}(I - \beta\Phi_1)^{-1}\Phi_1 x_t = 0 \quad \text{for all } x_t. \tag{20}$$

Recall that under the assumption that ξ_t is independently distributed, the forward solution of the Phillips curve takes the form

$$\pi_t = \lambda \sum_{j=0}^{\infty} \beta^j \mathrm{E}_t[MC_{t+j}] + \xi_t.$$

Thus, the first term in (20) can be interpreted as the one-step-ahead VAR forecast of inflation. The second term in (20) corresponds to the one-step-ahead forecast of the sum of discounted expected future marginal costs, scaled by the Phillips curve slope. As long as ξ_t is serially uncorrelated, the two forecasts have to be identical. Notice that although it might be impossible to uniquely determine λ and θ conditional on the VAR coefficient matrix Φ_1, Φ_1 is always identifiable based on the autocovariances of x_t, provided that $\mathrm{E}\left[x_t x_t'\right]$ is invertible: $\Phi_1 = \mathrm{E}\left[x_{t-1}x_t'\right](\mathrm{E}\left[x_t x_t'\right])^{-1}$. Hence, provided that inflation is serially correlated, the restriction (20) identifies λ.

Sbordone (2002, 2005) and Kurmann (2005, 2007) use (20) in conjunction with reduced-form VAR estimates of Φ to obtain estimates of the NKPC parameters. A system-based DSGE model estimation with serially uncorrelated mark-up shocks can be interpreted as simultaneously minimizing the discrepancy between an unrestricted, likelihood-based estimate of Φ_1 and the DSGE model-implied restriction function $\Phi_1(\lambda, \theta)$ and imposing the condition (20).

Identification of Backward-Looking Terms

Achieving identification becomes more difficult if we relax the restriction that $\gamma_b = 0$. Since insightful, analytical derivations are fairly complex, we offer a heuristic argument and point to some empirical evidence. Three factors contribute to the persistence of inflation: the backward-looking term $\gamma_b \widetilde{\pi}_{t-1}$, the persistence of marginal costs, and the persistence of the mark-up shock, ξ_t. Roughly speaking, we can measure inflation and marginal cost persistence from the data (provided observations on marginal costs are available). Hence, the challenge is to disentangle the relative contribution of γ_b and the mark-up shock to the persistence of inflation. Del Negro and Schorfheide (2006, Figure 8) display plots of the joint posterior distribution of γ_b and the autocorrelation of a latent mark-up shock obtained from the estimation of a DSGE model that is similar to the one developed by Smets and Wouters (2003). Not surprisingly, there is a strong negative correlation, suggesting that without strong *a priori* restrictions, it is difficult to measure the magnitude of γ_b. One widely used *a priori* restriction is the assumption that the mark-up shock is either absent or serially uncorrelated.

3. A (CRUDE) LOOK AT U. S. DATA

Before reviewing the DSGE model-based NKPC estimates reported in the literature, we will take a crude look at U.S. inflation, labor share, and output data. In view of the analysis presented in Section 2, two potentially important sources of variation in DSGE model-based estimates are (1) detrending methods for inflation data and marginal cost proxies and (2) endogeneity corrections. Thus, in the first subsection we construct different measures of steady-state deviations and compare the stochastic properties of the resulting $\widetilde{\pi}_t$, $\widetilde{MC}_t$, and $\widetilde{Y}_t$ series. We established that the estimation of the NKPC parameters amounts to a regression of inflation on future expected marginal costs. This regression, hidden within a complicated likelihood function, is plagued by an endogeneity problem, which, according to the simple model in Section 2, leads to a negative bias of least-squares estimates of the Phillips curve slope. It turns out that these least-squares estimates are relatively insensitive to data definitions (second subsection), which suggests that much of the variation across empirical studies is attributable to differences in the endogeneity correction.

We also showed that impulse response dynamics provide useful information about the NKPC coefficients. To the extent that a well-specified DSGE model is comparable in fit to a more densely parameterized VAR, evidence (reported in the third subsection) on the propagation of a monetary policy shock can be helpful to understand DSGE model-based estimates of NKPC parameters. Finally, the autocovariance restrictions exploited in the DSGE model-based estimation tend to nest those used by Sbordone (2002) to construct a VAR-based minimum distance estimator. Hence, we briefly review these minimum distance estimates in the fourth subsection.

Measures of Inflation and Marginal Costs

Most authors use the gross domestic product (GDP) deflator as a measure of inflation when estimating New Keynesian DSGE models. Our subsequent review focuses on estimates obtained with DSGE models in which marginal costs equal the labor share. These estimates either include the labor share in the vector of observables or treat marginal costs as a latent variable. In the latter case, deviations of aggregate output from a trend or natural level are implicitly used as a marginal cost proxy. To study the stochastic properties of these series, we compile a small data set with quarterly U.S. observations. The raw data are taken from Haver Analytics. Real output is obtained by dividing the nominal series (GDP) by population 16 years and older and by deflating using the chained-price GDP deflator. Inflation rates are defined as log differences of the GDP deflator. The labor share is computed by dividing total compensation of employees (obtained from the National Income and Product Accounts) by nominal GDP. We take logs of real per capita output and the labor share. Our sample ranges from 1960:Q1 to 2005:Q4.

We will consider three measures of $\widetilde{\pi}_t$. First, $\widetilde{\pi}$ (mean) is obtained by subtracting the sample mean computed over the period 1960 to 2005 from the GDP deflator inflation. This calculation assumes that the target inflation rate has essentially stayed constant for the past 45 years. Second, we compute separate means for the subsamples 1960–69, 1970–1982, and 1983–2005. The break points are broadly consistent with the regime estimates obtained in Schorfheide (2005). The resulting measure of inflation deviations is denoted by $\widetilde{\pi}$ (break) and reflects the view that the target inflation rate rose in the 1970s because policymakers perceived an exploitable long-run output inflation tradeoff. Finally, we consider $\widetilde{\pi}$ (HP), which can be interpreted as deviations from a drifting target inflation rate.

We plot the inflation rate as well as the three versions of the target inflation in Figure 1. It is apparent from the figure that views about target inflation significantly affect the stochastic properties of $\widetilde{\pi}_t$. For instance, the first-order autocorrelations (see Table 1) are 0.88, 0.68, and 0.49 for $\widetilde{\pi}$ (mean), $\widetilde{\pi}$ (break), and $\widetilde{\pi}$ (HP), respectively. The two panels of Figure 2 depict $\widetilde{MC}_t$ as approximated by output movements or measured by labor share fluctuations. In models that treat marginal cost as a latent variable, the most common

Figure 1 Inflation and Measures of Trend Inflation

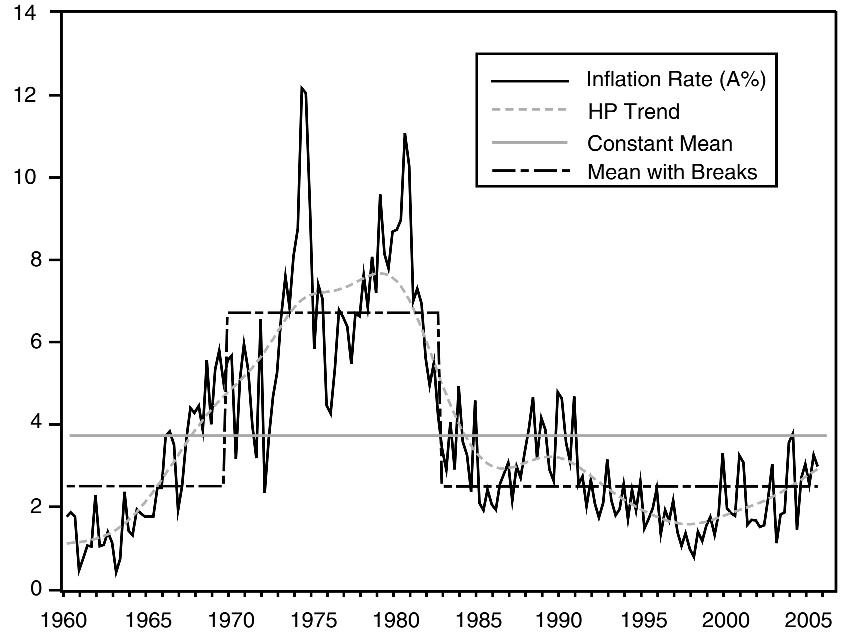

Notes: Inflation is measured as quarter-to-quarter changes in the log GDP deflator, scaled by 400 to convert it into annualized percentages. The sample ranges from 1960:Q1 to 2005:Q4.

marginal cost proxies are given by linearly detrended output, output deviations from a quadratic trend, and HP-filtered output. Since the potential output series produced by the Congressional Budget Office closely resembles the HP trend, we are not considering it separately.[5] Panel A clearly indicates that output deviations from a deterministic trend tend to be more volatile and persistent than deviations from the HP trend, since the HP filter removes more of the low frequency variation from the output series. Panel B shows time series for labor share deviations from a constant mean and an HP trend. As before, deviations from an HP trend tend to be smoother. First-order autocorrelations for the marginal cost measures are reported in Table 1. They range from 0.7 (HP-filtered labor share) to 0.97 (linearly detrended output).

[5] In some DSGE models, e.g., Schorfheide (2005), technology evolves according to a unit root process and the output term that appears in the Phillips curve refers, strictly speaking, to deviations from a latent stochastic trend. We do not consider this case in the regressions reported in this section.

Figure 2 Measures of Marginal Cost Deviations

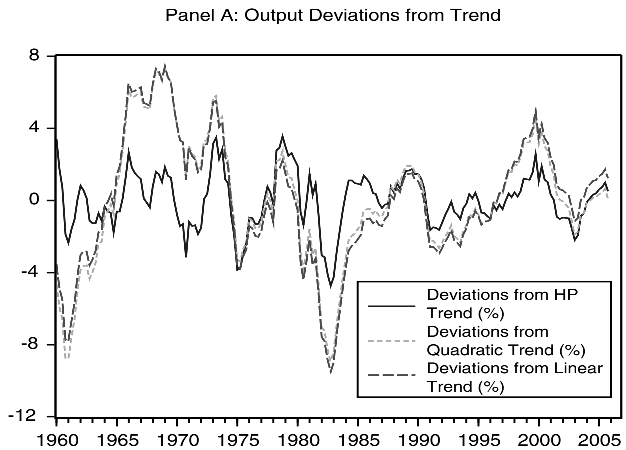

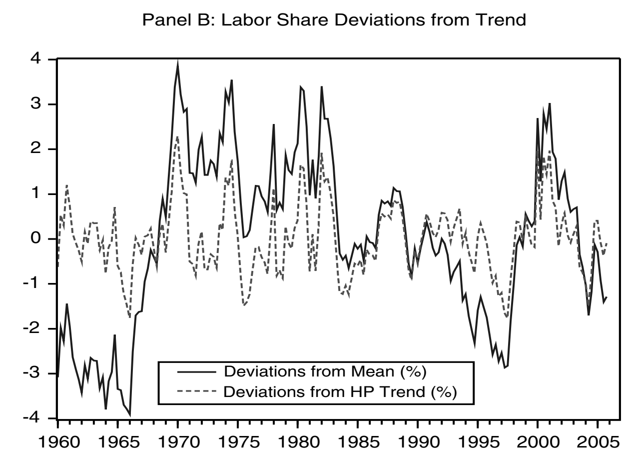

Inflation and Marginal Cost Regressions

Under the assumptions that $\gamma_b = 0$, $\widetilde{\xi}_t$ is serially uncorrelated, $\beta = 0.993$, and marginal cost dynamics are well approximated by an AR(1) process with coefficient $\hat{\rho}$, one can express the forward solution of (6) as

Table 1 Persistence of Marginal Cost and Inflation Measures

Series	AR(1)
$\tilde{\pi}$ (mean)	0.88
$\tilde{\pi}$ (break)	0.68
$\tilde{\pi}$ (HP)	0.49
$\tilde{Y}$ (linear trend)	0.97
$\tilde{Y}$ (quadratic trend)	0.96
$\tilde{Y}$ (HP)	0.85
$l\tilde{s}h$ (mean)	0.94
$l\tilde{s}h$ (HP)	0.70

Notes: The table reports AR(1) coefficient estimates based on a sample from 1960:Q1 to 2005:Q4.

$$\widetilde{\pi}_t = \kappa \left(\frac{1}{1 - 0.993\hat{\rho}_Y} \right) \widetilde{Y}_t + \tilde{\xi}_t, \quad \text{or} \quad \widetilde{\pi}_t = \lambda \left(\frac{1}{1 - 0.993\hat{\rho}_{lsh}} \right) \widetilde{lsh}_t + \tilde{\xi}_t, \tag{21}$$

where lsh denotes the labor share. As in Section 2, the parameter, κ, confounds the slope, λ, and the elasticity of marginal costs with respect to output. Least squares regression results for (21) are summarized in Table 2. We report point estimates of κ and λ in (12) and R^2 statistics in parenthesis for the full sample as well as three subsamples: 1960–1969, 1970–1982, and 1983–2005.

Since there is no guarantee that the mean of inflation and marginal cost deviations is zero in the subsamples, we also include an intercept in the regression. As in other studies, e.g., Rudd and Whelan (2007), we find that the slope estimates and the R^2 statistics tend to be small. The largest estimate of κ is 0.03, obtained by regressing demeaned inflation on the HP-filtered output using the 1960–1969 subsample. If one regresses inflation on the labor share, the largest slope estimate is 0.05, which is obtained by using an HP-filter on both inflation and the labor share and restricting the sample to 1970–1982. The median slope estimate reported in the table is 0.002. The R^2 values range from nearly zero to 66 percent. If we assume that the target inflation rate has shifted in early 1970 and 1982 and use the demeaned labor share as a measure of marginal cost, then $\hat{\lambda} = .003$ and R^2 is 6 percent. The Durbin-Watson statistics (not reported in the table) for the OLS regressions indicate that the least-squares residuals have substantial positive serial correlation.

We draw two broad conclusions for the subsequent review of DSGE model-based estimates. First, since the least-squares estimates range from 0 to 0.03 for κ and 0 to 0.05 for λ, any variation beyond this range is most likely caused by the endogeneity correction. Second, for the Phillips curve to

Table 2 Inflation and Marginal Cost Regressions

Inflation	**Marginal Cost** Measure $\widetilde{MC}_t$	**Sample Period** 1960–2005	1960–1969	1970–1982	1983–2005
$\widetilde{\pi}$ (mean)	$\widetilde{Y}$ (lin trend)	1E-4 (4E-4)	.002 (0.64)	−8E-4 (0.02)	−.001 (0.08)
	$\widetilde{Y}$ (quad trend)	8E-4 (0.01)	.002 (0.65)	−8E-4 (0.01)	−.001 (0.05)
	$\widetilde{Y}$ (HP)	.006 (.008)	0.03 (0.35)	.009 (0.06)	.003 (.008)
	$\widetilde{lsh}$ (mean)	0.01 (0.40)	0.01 (0.66)	.007 (0.03)	.002 (0.03)
	$\widetilde{lsh}$ (HP)	0.04 (0.03)	0.03 (0.03)	0.03 (0.04)	.002 (.001)
$\widetilde{\pi}$ (break)	$\widetilde{Y}$ (lin trend)	5E-4 (0.02)			
	$\widetilde{Y}$ (quad trend)	8E-4 (0.03)			
	$\widetilde{Y}$ (HP)	0.01 (0.07)		same as $\widetilde{\pi}$ (mean)	
	$\widetilde{lsh}$ (mean)	.003 (0.06)			
	$\widetilde{lsh}$ (HP)	0.02 (0.02)			
$\widetilde{\pi}$ (HP)	$\widetilde{Y}$ (lin trend)	2E-4 (.005)	2E-4 (0.03)	−4E-4 (.007)	.001 (0.11)
	$\widetilde{Y}$ (quad trend)	2E-4 (.006)	2E-4 (0.03)	−4E-4 (.006)	.001 (0.11)
	$\widetilde{Y}$ (HP)	.004 (0.02)	.007 (0.12)	.002 (.005)	.007 (0.07)
	$\widetilde{lsh}$ (mean)	.002 (0.03)	.003 (0.16)	0.01 (0.17)	5E-4 (.002)
	$\widetilde{lsh}$ (HP)	0.03 (0.08)	0.02 (0.07)	0.05 (0.15)	0.01 (0.03)

Notes: We are reporting coefficient estimates, $\hat{\alpha}_1$ and R^2, in parenthesis for a regression of the form $\widetilde{\pi}_t = \alpha_0 + \alpha_1[\widetilde{MC}_t/(1-\hat{\beta}\hat{\rho}_{MC})]$, where $\hat{\beta} = 0.993$, $\hat{\rho}_{MC}$ is the first-order autocorrelation of the marginal cost measure, and the marginal cost measure is either $\widetilde{Y}_t$ or $\widetilde{lsh}_t$.

capture the inflation persistence well, it has to be the case that lagged inflation enters the NKPC, that the mark-up shock is fairly persistent, or that inflation deviations are computed relative to a time-varying target inflation rate.

VAR-IRF Evidence

We explained in Section 2 that if the DSGE model imposes enough restrictions to unambiguously identify, say, a monetary policy shock, then the response of output and marginal costs to this shock provides useful information about the NKPC parameters. To the extent that we would expect a well-specified DSGE model to generate impulse responses that are similar to those obtained from a structural VAR analysis, it is informative to examine prototypical VAR responses to a monetary policy shock and to determine a range of NKPC parameterizations that are consistent with these responses.

Under the assumption that lagged inflation does not enter the NKPC and that marginal costs are proportional to output, the impulse responses to a monetary policy shock have to satisfy

$$\frac{\partial \tilde{\pi}_{t+h}}{\partial \epsilon_t^R} = \kappa \sum_{j=0}^{\infty} \beta^j \mathrm{E}_{t+h}\left[\frac{\partial \tilde{Y}_{t+h+j}}{\partial \epsilon_t^R}\right].$$

As in Section 2, we use κ to denote the slope of the Phillips curve with respect to output. The parameter, κ, absorbs the elasticity of marginal costs with respect to output into the definition of the slope. Suppose that the impulse responses are monotonic and the output response decays approximately exponentially at rate δ in response to a monetary policy shock. Then

$$\frac{\partial \tilde{\pi}_{t+h}}{\partial \epsilon_t^R} \approx \frac{\kappa}{1-\delta\beta} \frac{\partial \tilde{Y}_{t+h}}{\partial \epsilon_t^R}.$$

While a large literature exists (see Christiano, Eichenbaum, and Evans [1999] and Stock and Watson [2001] for surveys) that uses structural VARs to measure the effect of monetary policy shocks, we focus on a prominent recent study by Christiano, Eichenbaum, and Evans (2005).

The authors estimate a nine-variable VAR using data on real GDP, real consumption, the GDP deflator, real investment, the real wage, labor productivity, the federal funds rate, real profits, and the growth rate of M2. Christiano, Eichenbaum, and Evans (2005) find that a 15-basis point (bp) drop in the federal funds rate (quarterly percentage points) leads to a 5-bp increase in the quarterly inflation rate after 12 quarters and a 50-bp increase of output after nine quarters.[6] Hence, according to the mean impulse responses, κ should be about 0.1 if we set the decay factor, δ, to zero and 0.05 if we

[6] These numbers are approximate, based on Figure 1 in Christiano, Eichenbaum, and Evans (2005).

set $\delta = 0.5$. Suppose now that we ignore the dependence in the sampling distribution of the impulse response function estimators and let $\delta = 0$ again. Combining the lower bound of the reported 95 percent confidence band of the inflation response with the upper bound of the confidence band for the output response suggests that κ could be as low as 0.01. Combining the upper bound for the inflation response with the lower bound for the output response leads to a value of $\kappa = 0.5$. If we consider the labor share instead of the output response, we can obtain an estimate of λ instead of κ. Along the mean impulse response estimated by Christiano, Eichenbaum, and Evans (2005), the labor share appears to drop by about 25 bp, which for $\delta = 0$ and $\delta = 0.5$ leads to values of $\lambda = 0.2$ and $\lambda = 0.1$, respectively.

Evidence from Inflation and Marginal Cost Dynamics

Several papers, e.g., Sbordone (2002, 2005) and Kurmann (2005, 2007), exploit the restriction (20) to construct minimum-distance estimates of the NKPC parameters from the estimates of an unrestricted VAR that includes inflation and a measure of marginal costs. Using data from 1951 to 2002 on the labor share and inflation, Sbordone (2005) obtains an estimate of $\hat{\lambda} = 0.025$ in the purely forward-looking specification, and $\hat{\lambda} = 0.014$ and $\hat{\gamma}_b = 0.18$ if she allows lagged inflation to enter the NKPC. To the extent that the restriction (20) is also embodied in a DSGE model likelihood function, the DSGE model-based estimates of the NKPC parameters should be similar, provided that the same measure of marginal costs is used, the mark-up shock is assumed to be i.i.d., and the vector autoregressive approximation to the law of motion of the estimated DSGE model resembles the unrestricted VAR estimates.

4. REVIEW OF EMPIRICAL RESULTS

Broadly speaking, the empirical papers reviewed in this section fall into two categories: either marginal costs are treated as a latent variable or the set of observables spans the labor share and, hence, the model-implied measure of marginal costs. Consider once again the simple model of Section 2 and let us denote the labor share as lsh. Abstracting from inference about γ_b and γ_f, a study that estimates λ in

$$\tilde{\pi}_t = \beta \mathrm{E}_t[\tilde{\pi}_{t+1}] + \frac{\lambda}{\tau}\tilde{Y}_t - \tilde{Z}_t + \tilde{\xi}_t, \tag{22}$$

based on observations of $\tilde{\pi}_t$ and $\tilde{Y}_t$, falls in the first category. Identification of λ in (22) is tenuous because the presence of $\tilde{Z}_t$ exacerbates the endogeneity problem and the parameter, τ, has to be separately estimable from the observables for λ to be identifiable. On the upside, the use of (22) is more robust to the presence of measurement errors in the labor share (marginal cost) series. For some of the papers that fall into the first category, we will report estimates of the output coefficient, κ, which corresponds to $\frac{\lambda}{\tau}$ in the example, rather than λ. A paper that estimates λ in

$$\tilde{\pi}_t = \beta \mathrm{E}_t[\tilde{\pi}_{t+1}] + \lambda \widetilde{lsh}_t + \tilde{\xi}_t, \tag{23}$$

with observations on $\tilde{\pi}_t$ and $\widetilde{lsh}_t$, belongs to the second category.

Since the literature on estimated DSGE models is growing rapidly, we had to strike a balance between scope and depth. This survey is limited to models in which firms' price-setting equations are derived either under quadratic adjustment costs or under the Calvo mechanism. Ongoing research explores alternative sources of nominal rigidities that are not included in the subsequent review, for instance, menu costs and state-dependent pricing models (Dotsey, King, and Wolman 1999, Gertler and Leahy 2006), models with labor market search frictions (Gertler and Trigari 2006, Krause and Lubik 2007), and models with information processing frictions (Sims 2003, Mackowiak and Wiederholt 2007, Mankiw and Reis 2007, and Woodford 2008). Moreover, we focus on models in which the labor share is the model-implied measure of marginal costs.[7]

The numerical values reported in Tables 3–5 refer to point estimates that are obtained with one of four methods. In addition to papers that use Bayesian[8] and maximum likelihood methods as discussed in Section 1, we consider studies that estimate the DSGE model parameters by minimizing the discrepancy between impulse responses implied by the DSGE model and those obtained from the estimation of a structural VAR, or by minimizing the distance between sample moments obtained from U.S. data and DSGE model-implied population moments. The remainder of this section is organized as follows. We review estimates that are obtained by treating marginal costs as a latent variable. We examine studies in which the authors treat marginal costs as observable. For monetary policy analysis, the relationship between inflation and output is at least as important as the relationship between inflation and marginal costs. So we examine DSGE model-based estimates of the relative movements of inflation and output in response to a monetary policy shock. Finally, we discuss the role of wage stickiness.

[7] Krause, López-Salido, and Lubik (2008) show that in a model with labor market search frictions, marginal costs are also affected by the labor market tightness. However, empirically they find that matching frictions in the labor market appear to affect the cyclical behavior of marginal costs only slightly in terms of co-movement, persistence, and volatility.

[8] Bayesian inference combines information contained in the likelihood function with prior information to form posterior estimates. Since it is difficult to disentangle the contribution of various sources of information ex post, we restrict our attention to the posterior estimates without examining the priors that were used to generate these posteriors.

Latent Marginal Costs

Table 3 summarizes parameter estimates of a Phillips curve specification in which marginal costs are replaced by output or a measure of the output gap:

$$\tilde{\pi}_t = \gamma_b \tilde{\pi}_{t-1} + \gamma_f \mathrm{E}_t[\tilde{\pi}_{t+1}] + \kappa \tilde{Y}_t + \tilde{\xi}_t, \tag{24}$$

where $\tilde{\xi}_t$ represents the latent variables that enter the NKPC in any particular model. These estimates are obtained by fitting New Keynesian DSGE models to observations of output, inflation, and interest rates. The models implicitly share the following features: household preferences are linear in labor and capital is not a factor of production. Estimates for κ range from values less than 0.001 (Cho and Moreno 2006) to 4.15 (Canova forthcoming). While the studies differ with respect to sample period as well as the detrending of inflation and output, our least-squares analysis in Section 3 suggests that most of the differences in $\hat{\kappa}$ are probably due to the treatment of latent variables. We showed that the likelihood function corrects for the endogeneity problem that arises in a regression of inflation on future expected output due to the correlation of the latent variables with expected output. This endogeneity correction appears to be very sensitive to the assumed correlation among the exogenous disturbances that enter the Phillips curve, the Euler equation, and the monetary policy rule. Models in which the shocks in the Euler equation and the Phillips curve are forced to be or allowed to be correlated tend to deliver larger estimates of κ than models in which these disturbances are assumed to be uncorrelated.[9]

We now turn to estimates of New Keynesian Phillips curves that are expressed in terms of marginal costs instead of output:

$$\tilde{\pi}_t = \gamma_b \tilde{\pi}_{t-1} + \gamma_f \mathrm{E}_t[\tilde{\pi}_{t+1}] + \lambda \widetilde{MC}_t + \tilde{\xi}_t. \tag{25}$$

These estimates are reported in Table 4. Rabanal and Rubio-Ramírez (2005) fit a canonical New Keynesian DSGE model without capital and habit formation using a data set that contains, in addition to inflation, interest rates, and detrended output, a measure of the real wage. For specifications in which γ_b is restricted to be zero, the authors obtain estimates of λ of about 0.015. If γ_b is estimated subject to the restriction that $\gamma_b + \gamma_f = 0.99$, the estimate of

[9] Correlation arises either because the structural model implies that, say, the government spending shock enters both the Phillips curve and the Euler equation (e.g., Schorfheide 2005), or because authors attach reduced-form disturbances to the Phillips curve and the Euler equation and assume that these disturbances are correlated (e.g., Lubik and Schorfheide 2004). In small-scale models, it is often "testable" whether the exogenous disturbances are correlated. In large DSGE models, parameters associated with the endogenous propagation mechanism and auxiliary parameters that generate correlation between exogenous disturbances are often not separately identifiable.

Table 3 Published NKPC Estimates: Latent Labor Share (Part 1)

Study	**Sample Period**	π_{t-1}	$\mathrm{E}_t[\pi_{t+1}]$	Y_t	**Method**
No capital, no habit formation, output coefficient in Phillips curve					
Canova (forthcoming), Table 1	1955:Q1–2002:Q1		0.98	4.150	Bayes
Cho and Moreno (2006), Table 2	1980:Q4–2000:Q1	0.44	0.56	0.001	MLE
Cho and Moreno (2006), Table 2	1980:Q4–2000:Q1	0.44	0.56	0.001	MLE
Cho and Moreno (2006), Table 2	1980:Q4–2000:Q1	0.43	0.57	0.000	MLE
Del Negro and Schorfheide (2004), Table 2	1959:Q3–1979:Q2		1.00	0.314	Bayes
Del Negro and Schorfheide (2004), Table 2	1959:Q3–1979:Q2		1.00	0.249	Bayes
Lindé (2005), Table 5	1960:Q1–1997:Q4	0.72	0.28	0.048	MLE
Lindé (2005), Table 5	1960:Q1–1997:Q4	0.54	0.46	0.048	MLE
Lubik and Schorfheide (2004), Table 3	1960:Q1–1979:Q2		.997	0.770	Bayes
Lubik and Schorfheide (2004), Table 3	1960:Q1–1979:Q2		.997	0.750	Bayes
Lubik and Schorfheide (2004), Table 3	1982:Q4–1997:Q4		.993	0.580	Bayes
Rotemberg and Woodford (1997), Page 321	1980:Q1–1995:Q2		0.99	0.024	IRF-MD
Salemi (2006), Table 2	1965:Q1–2001:Q4	0.62	0.00	0.055	MLE
Salemi (2006), Table 2	1965:Q1–2001:Q4	0.43	0.57	0.003	MLE
Schorfheide (2005), Table 2	1960:Q1–1997:Q4		0.99	0.370	Bayes
Schorfheide (2005), Table 2	1960:Q1–1997:Q4		1.00	0.360	Bayes

Notes: We are providing point estimates of the New Keynesian Phillips curve, $\pi_t = \gamma_b \pi_{t-1} + \gamma_f \mathrm{E}_t[\pi_{t+1}] + \kappa Y_t + \xi_t$, based on the information provided in the cited studies. Estimation methods: MLE = maximum likelihood estimation; Bayes = Bayesian analysis; and IRF-MD minimize discrepancy between impulse responses estimated with a structural VAR and those implied by a DSGE model.

λ drops to 0.004, whereas the weight on lagged inflation in the Phillips curve is 0.43.

The canonical three equation New Keynesian DSGE model that underlies, for instance, the analysis in Rabanal and Rubio-Ramirez (2005) lacks persistent dynamics, which makes it difficult to capture the serial correlation in U.S. data. The lack of persistence can be overcome in part by using household preferences that exhibit habit formation, that is, the instantaneous utility is a function of current consumption relative to some habit stock, which in turn depends on past consumption. Habit formation enriches the endogenous propagation mechanism of the model and enhances the model's ability to capture the persistence in output and consumption. More importantly for us, it changes the relationship between (observed) output and (latent) marginal costs. The marginal utility of consumption, and thereby marginal costs, depends not just on the current level of output, but also on past and expected future levels as well as the parameters that determine the degree of habit formation. The estimates of λ reported in the second section of Table 4 range from 0.004 to 0.437.

If capital is treated as a variable input, marginal costs remain equal to the labor share as long as the production function is of the Cobb-Douglas form. However, if labor share observations are not used directly in the estimation, the presence of capital affects the decomposition of marginal costs into an observed and an unobserved component. In other words, if the DSGE model is estimated based on observations of output, inflation, and interest rates, introducing variable capital changes the stochastic properties of $\tilde{\xi}_t$ in (24) and the relationship between κ in (24) and λ in (25). The third section of Table 4 reports NKPC estimates from six studies, ranging from 0.008 to 0.112. Among these papers, Fernandez-Villaverde and Rubio-Ramirez (2007) allow the parameters of the monetary policy rule and the parameters that determine the degree of price and wage stickiness to vary over time. This allows the authors to obtain a time series of the Phillips curve coefficient. If the slope estimates of the Phillips curve are converted into the probability that a firm is unable to change its price in a Calvo model (see Section 1), then the estimates can be summarized as follows. Prices stayed constant for an average of four quarters in the 1960s and 1970s, while inflation was relatively high and became a bit more rigid after the Volcker disinflation. Based on a casual inspection of the smoothed time series of the Phillips curve coefficients, λ appeared to be, on average, around 0.06 before 1979 and subsequently dropped to 0.03. The average estimate of γ_b pre-1979 is about 0.35 and decreased to 0.3 after 1979. This pattern is broadly consistent with the notion that the NKPC is not structural in the following sense: If a high target inflation rate makes it very costly for firms not to change their prices—and, hence, more attractive to incur the costs of adjusting the prices—we should observe a steeper Phillips curve relationship.

Table 4 Published NKPC Estimates: Latent Labor Share (Part 2)

Study	Sample Period	π_{t-1}	$E_t[\pi_{t+1}]$	MC_t	Method
	No capital, no habit formation				
Rabanal and Rubio-Ramírez (2005), Table 2	1960:Q1–2001:Q4		0.99	0.015	Bayes
Rabanal and Rubio-Ramírez (2005), Table 2	1960:Q1–2001:Q4	0.43	0.56	0.004	Bayes
Rabanal and Rubio-Ramírez (2005), Table 2	1960:Q1–2001:Q4		0.99	0.016	Bayes
Rabanal and Rubio-Ramírez (2005), Table 2	1960:Q1–2001:Q4		0.99	0.017	Bayes
	No capital, with habit formation				
Andres, López-Salido, and Nelson (2004), Table 1	1980:Q1–1999:Q2		0.99	0.014	MLE
Andres, López-Salido, and Nelson (2005), Table 2	1979:Q3–2003:Q3	0.50	0.50	0.437	MLE
Boivin and Giannoni (2006), Table 2	1959:Q2–1979:Q2	0.50	0.50	0.006	IRF-MD
Boivin and Giannoni (2006), Table 2	1979:Q3–2002:Q2	0.50	0.50	0.004	IRF-MD
Galí and Rabanal (2004), Table 4	1948:Q1–2002:Q4	0.02	0.97	0.413	Bayes
Lubik and Schorfheide (2006), Table 2	1983:Q1–2002:Q4		1.00	0.200	Bayes
Milani (2007), Table 2	1960:Q1–2004:Q2		0.99	0.024	Bayes
	Models with capital				
Bouakez, Cardia, and Ruge-Murcia (2005), Table 1	1960:Q1–2001:Q2	0.50	0.50	0.015	MLE
Bouakez, Cardia, and Ruge-Murcia (forthcoming), Table 4	1964:Q1–2002:Q4		1.00	0.223	MD
Christensen and Dib (2008)	1979:Q3–2004:Q3		.993	0.092	MLE
Fernández-Villaverde and Rubio-Ramírez (2007), Table 6.1	1955:Q1–2000:Q4	0.13	0.87	0.008	Bayes
Fernández-Villaverde and Rubio-Ramírez (2007), Table 6.1	1955:Q1–1979:Q4	0.26	0.74	0.056	Bayes
Fernández-Villaverde and Rubio-Ramírez (2007), Table 6.1	1980:Q1–2000:Q4	0.13	0.87	0.030	Bayes
Laforte (2007), Table 3	1983:Q1–2003:Q1	0.40	0.59	0.112	Bayes
Rabanal (2007), Table 2	1959:Q1–2004:Q4	0.50	0.50	0.018	Bayes

Notes: We are providing point estimates of the New Keynesian Phillips curve, $\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t[\pi_{t+1}] + \lambda MC_t + \xi_t$, based on the information provided in the cited studies. Estimation methods: MLE = maximum likelihood estimation; Bayes = Bayesian analysis; IRF-MD minimize discrepancy between impulse responses estimated with a structural VAR and those implied by a DSGE model; and MD minimize discrepancy between sample moments and DSGE model-implied moments.

Observed Marginal Costs

We now turn to the Bayesian estimation of New Keynesian DSGE models based on a larger set of observables that spans the labor share and, hence, marginal costs as they appear in the Phillips curve. Intuitively, the use of labor share observations should lead to a sharper identification of λ. Table 5 summarizes empirical estimates from seven studies. Most estimates are based on a variant of the Smets and Wouters (2003) model, which augments a DSGE model by Christiano, Eichenbaum, and Evans (2005) with additional shocks to make it amenable to likelihood-based estimation. Smets and Wouters (2005), Levin et al. (2006), Del Negro et al. (2007), Smets and Wouters (2007), and Justiniano and Primiceri (2008) obtain estimates of λ of 0.01, 0.03, 0.10, 0.02, and 0.01, respectively. The estimates of the coefficient γ_b on lagged inflation are 0.25, 0.07, 0.43, 0.19, and 0.46, respectively. Compared to the numbers reported in Tables 3 and 4, the variation across studies is much smaller.

Impulse Response Dynamics

Much of our previous discussion focused on the marginal cost coefficient in the Phillips curve relationship. However, from a monetary policy perspective, equally important is the output-inflation tradeoff in the estimated DSGE model. This tradeoff not only depends on λ but also on the elasticity of marginal costs with respect to output. Thus, we will examine the relative movements of output and inflation in response to a monetary policy shock, that is, an unanticipated deviation from the systematic component of the monetary rule. Of course, these impulse responses do not merely depend on the slope of the NKPC, they also depend on other aspects of the model, such as labor market frictions and wage stickiness and the behavior of the central bank. Not all the papers for which we have reported estimates of the NKPC parameters in Tables 3 to 4 present impulse response functions. Those that do typically represent them in graphical form. The subsequent results are based on an inspection of impulse response plots and are summarized in Table 6.[10] We report the magnitude of the peak responses of the interest rate, inflation rate, and the output deviation from steady state. The interest rate response is measured in annualized percentages; that is, an entry of 0.25 implies that the monetary policy shock raises the interest rate 25 bp above its steady-state level. The inflation rate is not annualized and represents a quarter-to-quarter difference in the log price level, scaled by 100 to convert it into percentages. Output deviations are also reported in percentages. Since the length of a period in a DSGE model is typically assumed to be one quarter, in the context of the

[10] In a number of studies, it turned out to be difficult to determine whether interest rates and inflation rates are annualized. We tried to resolve this ambiguity.

Table 5 Published NKPC Estimates: Observed Labor Share

Study	**Sample Period**	π_{t-1}	$\mathrm{E}_t[\pi_{t+1}]$	MC_t	**Method**
Avouyi-Dovi and Matheron (2007), Tables 3–4	1955:Q1–1979:Q2	0.27	0.73	0.008	IRF-MD
Avouyi-Dovi and Matheron (2007), Tables 3–4	1982:Q3–2002:Q4	0.20	0.80	0.010	IRF-MD
Christiano, Eichenbaum, and Evans (2005), Table 2	1965:Q3–1995:Q3	0.50	0.50	0.135	IRF-MD
Del Negro et al. (2007), Table 1	1974:Q2–2004:Q1	0.43	0.57	0.100	Bayes
Justiniano and Primiceri (2008), Table 1	1954:Q3–2004:Q4	0.46	0.54	0.007	Bayes
Justiniano and Primiceri (2008), Table 1	1954:Q3–2004:Q4	0.46	0.54	0.005	Bayes
Levin et al. (2006), Table 1	1955:Q1–2001:Q4	0.07	0.92	0.033	Bayes
Smets and Wouters (2005), Table 1	1983:Q1–2002:Q2	0.25	0.74	0.007	Bayes
Smets and Wouters (2007), Table 1A/B	1966:Q1–2004:Q4	0.19	0.82	0.020	Bayes

Notes: We provide point estimates of the New Keynesian Phillips curve, $\pi_t = \gamma_b \pi_{t-1} + \gamma_f \mathrm{E}_t[\pi_{t+1}] + \lambda MC_t + \xi_t$, based on the information provided in the cited studies. Estimation methods: Bayes = Bayesian analysis; IRF-MD minimize discrepancy between impulse responses estimated with a structural VAR and those implied by a DSGE model.

"back-of-the-envelope" calculation in Section 3, the ratio of the inflation and output response, denoted by $\partial\pi/\partial y$, would correspond to $\kappa/(1 - \delta\beta)$, where δ is the factor at which the output response decays to zero.

In Table 6 we report the number of periods it takes for the responses to reach their respective peaks, the ratio of the peak response of inflation and output, and the estimate of $\hat{\kappa}$ in the underlying model. Models without capital and with little endogenous propagation typically generate monotonic impulse response functions. For the models without capital, the relative responses of inflation and output range from 0.07 to 2.00. Once capital is included and the model is augmented by additional frictions, this range narrows to 0.08 to 0.17, which seems consistent with the VAR evidence provided by Christiano, Eichenbaum, and Evans (2005). Comparing the estimates reported in Del Negro et al. (2007) and Smets and Wouters (2007), it appears that these tradeoffs can be obtained with quite differently priced Phillips curve slopes, λ: 0.002 and 0.10.

Wage Versus Price Rigidity

This article has focused on estimates of the degree of price rigidity in New Keynesian DSGE models. Many authors believe that inflexible wages are another important source of nominal rigidities. In fact, the DSGE models that are based on the work of Smets and Wouters (2003), and Christiano, Eichenbaum, and Evans (2005) incorporate both price and wage stickiness. Following work by Erceg, Henderson, and Levin (2000), in order to generate wage stickiness in DSGE models, one typically assumes that households supply differentiated labor services that are aggregated by labor packers into homogenous labor services. These homogeneous labor services are in turn utilized by the intermediate goods-producing firms. Households act as monopolistically competitive suppliers and are subjected to a Calvo (1983) friction: only a fraction of households is allowed to re-optimize nominal wage. To clear the labor market ex post, one must assume that each household has to satisfy the demand for its labor service at the posted price.

For a joint estimation of price and wage rigidity to be meaningful, the set of observables needs to span inflation, labor share, and wages. The joint dynamics of inflation and the labor share provide information about the price Phillips curve, and the wage series, together with an implicit measure of the marginal disutility of work, contains information about the degree of wage stickiness. Del Negro and Schorfheide (2008) estimate a variant of the Smets and Wouters (2003) under three priors that differ with respect to a priori beliefs about nominal rigidities. The *low rigidities* prior assumes that the price and wage Calvo parameters have a beta-distribution centered at 0.45 with a standard deviation of 0.10. The *high rigidities* prior is centered at 0.75

Table 6 Impulse Responses to a Monetary Policy Shock

Study	Interest Rate [Annualized %]		Inflation [Quarterly %]		Output [% Dev. from Trend]		$\partial\pi/\partial y$
	Peak	After x Q's	Peak	After x Q's	Peak	After x Q's	
No capital, no habit formation, marginal costs function of current output deviations							
Cho and Moreno (2006), Figure 2	−0.80	0	0.01	9	0.12	9	0.08
Del Negro and Schorfheide (2004), Figure 2	−0.25	0	0.05	0	0.05	0	1.00
Ireland (2004a), Figure 1	−0.20	0	0.20	0	0.60	0	0.33
Ireland (2004b), Figure 1	−1.00	0	0.07	0	0.40	0	0.18
Ireland (2007), Figure 2	−0.40	0	0.10	0	0.30	0	0.33
Lubik and Schorfheide (2004), Figure 3	−0.70	0	0.12	0	0.16	0	0.75
Lubik and Schorfheide (2004), Figure 3	−0.60	0	0.17	0	0.16	0	0.76
Lubik and Schorfheide (2004), Figure 3	−0.60	0	0.20	0	0.16	0	1.25
Rotemberg and Woodford (1997), Figure 1	−0.80	0	0.03	2	0.38	2	0.07
Salemi (2006), Figure 3	−1.00	0	0.020	10	0.40	9	0.05
Salemi (2006), Figure 3	−1.00	0	0.002	8	0.30	8	0.01
No capital, with habit formation, marginal costs are function of current, past, and future output deviations							
Andres, López-Salido, and Nelson (2004), Figure 2	−0.30	0	0.17	0	0.40	3	0.43
Andres, López-Salido, and Nelson (2005), Figure 1	−0.50	0	0.15	0	0.08	1	2.00
Andres, López-Salido, and Nelson (2005), Figure 1	−0.70	0	0.06	0	0.40	1	0.15
Boivin and Giannoni (2006), Figure 1	−1.00	0	0.14	6	1.10	4	0.13
Boivin and Giannoni (2006), Figure 1	−1.00	0	0.02	4	0.30	4	0.07
With capital, no direct observations on labor share							
Christensen and Dib (2008), Figure 1	−0.48	2	0.14	2	0.50	2	0.28
Laforte (2007), Figure 2	−0.75	0	0.15	0	0.50	2	0.30
Rabanal (2007), Figure 4	−1.00	0	0.10	4	0.50	3	0.20
With capital, with direct observations on labor share							
Christiano, Eichenbaum, and Evans (2005), Figure 1	−0.60	4	0.05	11	0.60	6	0.08
Del Negro et al. (2007), Figure 3	−1.10	0	0.05	2	0.33	2	0.15
Smets and Wouters (2005), Figure 5	−0.70	4	0.05	3	0.45	5	0.11
Smets and Wouters, (2007) Figure 6	−0.72	0	0.05	2	0.30	3	0.17

Notes: Based on the graphical information provided in the cited studies, we determined the peak responses for interest rates (annualized percentage points), inflation (quarter-to-quarter percentage points), and output (percentage deviations from trend/steady state) to an unanticipated loosening of monetary policy.

with a standard deviation of 0.1. Finally, the *agnostic* prior is centered at 0.6 and is more diffuse—its standard deviation is 0.2.

Posterior inference based on these priors can be summarized as follows: both under the *agnostic* and the *low rigidities* prior, the posterior estimate of the wage stickiness is small. The Calvo parameter is around 0.25, which means that the households re-optimize their wages, on average, every four months. The estimated price stickiness translates into a value of λ of about 0.22. Under the *high rigidities* prior, the estimates of both the wage and the price Calvo parameter turn out to be substantially larger, namely about 0.8. Most interestingly, the time series fit of all three specifications is very similar, yet the policy implications are quite different. The results presented in Del Negro and Schorfheide (2008) suggest that the macro time series we typically consider is not informative enough to precisely measure the degree of nominal rigidity. This conclusion is consistent with the literature survey conducted in this section: the variation of parameter estimates reported in the literature is substantial. No clear consensus has emerged as of now.

5. CONCLUSION

While the literature on DSGE model-based estimation of the NKPC is still fairly young, a wide variety of results have been published in academic journals already. In most of these studies, the Phillips curve estimation is not a goal but rather a byproduct of the empirical analysis. DSGE model-based NKPC estimates tend to be fragile and sensitive to model specification and data definitions, in particular if marginal costs are treated as a latent variable. If the observations span the labor share, which is the model-implied measure of marginal costs in the studies that we reviewed, then the slope estimates are more stable. No consensus has emerged on the importance of lagged inflation in the Phillips curve. Estimates are sensitive to detrending methods for inflation and assumptions about the autocovariance structure of the exogenous disturbances in the DSGE model. Thus, from a policymaker's perspective, accounting for parameter and model uncertainty is important for prediction and decision making.

We attempted to understand the identification of Phillips curve parameters in estimated DSGE models. Unlike single-estimation approaches, DSGE model-based estimates are able to extract information about the structural parameters from the contemporaneous correlations of output, inflation, interest rates, and other variables, as well as from impulse responses to structural shocks that are identifiable based on exclusion restrictions hard-wired in model specifications. Unfortunately, the data do not speak loudly and clearly to us, and many DSGE models imply that if the model is "true," it is difficult to identify the NKPC parameters and the output-inflation tradeoff with only 20 to 40 years of observations.

Identification in the context of simultaneous equations models is well understood. To identify the slope of a supply curve we need variation in exogenous demand shifters. Identification in DSGE models is much more complicated. Variation in the data is created by unobserved shocks that in most cases shift both demand and supply. Our reading of the early literature on estimated DSGE models is that there was hope that the model-implied cross-coefficient restrictions were so tight that identification was not a concern. Over time the profession learned that, despite tight cross-equation restrictions, identification should not be taken for granted, in particular in New Keynesian DSGE models. While currently ongoing research is developing econometric techniques to try to diagnose identification problems, it might be time to go back to the drawing board and develop future DSGE models with parameter identifiability in mind.

REFERENCES

Altug, Sumru. 1989. “Time-to-Build and Aggregate Fluctuations: Some New Evidence.” *International Economic Review* 30 (November): 889–920.

An, Sungbae, and Frank Schorfheide. 2007. “Bayesian Analysis of DSGE Models.” *Econometric Reviews* 26: 113–72.

Andres, Javier, J. David López-Salido, and Edward Nelson. 2004. “Tobin’s Imperfect Asset Substitution in Optimizing General Equilibrium.” *Journal of Money, Credit and Banking* 36 (August): 665–90.

Andres, Javier, J. David López-Salido, and Edward Nelson. 2005. “Sticky-price Models and the Natural Rate Hypothesis.” *Journal of Monetary Economics* 52 (July): 1025–53.

Avouyi-Dovi, Sanvi, and Julien Matheron. 2007. “Technology Shocks and Monetary Policy: Revisiting the Fed’s Performance.” *Journal of Money, Credit and Banking* 39: 471–507.

Bils, Mark, and Pete Klenow. 2004. “Some Evidence on the Importance of Sticky Prices.” *Journal of Political Economy* 112 (October): 947–85

Boivin, Jean, and Marc P. Giannoni. 2006. “Has Monetary Policy Become More Effective?” *Review of Economics and Statistics* 8 (October): 445–62.

Bouakez, Hafedh, Emanuela Cardia, and Francisco J. Ruge-Murcia. 2005. “Habit Formation and the Persistence of Monetary Shocks.” *Journal of Monetary Economics* 52 (September): 1073–88.

Bouakez, Hafedh, Emanuela Cardia, and Francisco J. Ruge-Murcia. Forthcoming. “The Transmission of Monetary Policy in a Multi-sector Economy.” *International Economic Review*.

Calvo, Guillermo. 1983. “Staggered Prices in a Utility-Maximizing Framework.” *Journal of Monetary Economics* 12 (September): 383–98.

Canova, Fabio. Forthcoming. “What Explains the Great Moderation in the U.S.? A Structural Analysis.” *Journal of European Economic Association*.

Cho, Seonghoon, and Antonio Moreno. 2006. “A Small-Sample Study of the New Keynesian Macro Model.” *Journal of Money, Credit and Banking* 39 (September): 471–507.

Christensen, Ian, and Ali Dib. 2008. “The Financial Accelerator in an Estimated New Keynesian Model.” *Review of Economic Dynamics* 11 (January): 155–78.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1999. “Monetary Policy Shocks: What have We Learned and to What End?” In *Handbook of Macroeconomics*, edited by John Taylor and Micheal Woodford. North Holland, Amsterdam: Elsevier, 65–148.

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 (February): 1–45.

Del Negro, Marco, and Frank Schorfheide. 2004. “Priors from General Equilibrium Models for VARS.” *International Economic Review* 45 (May): 643–73.

Del Negro, Marco, and Frank Schorfheide. 2006. “Forming Priors for DSGE Models (and How it Affects the Assessment of Nominal Rigidities).” Federal Reserve Bank of Atlanta Working Paper 2006-16.

Del Negro, Marco, and Frank Schorfheide. 2008. “Forming Priors for DSGE Models (and How it Affects the Assessment of Nominal Rigidities).” *Journal of Monetary Economics* 55 (October): 1191–208.

Del Negro, Marco, Frank Schorfheide, Frank Smets, and Rafael Wouters. 2007. “On the Fit of New Keynesian Models.” *Journal of Business and Economic Statistics* 25 (April): 123–62.

Dotsey, Michael, Robert King, and Alexander Wolman. 1999. “State Dependent Pricing and the General Equilibrium Dynamics of Money and Output.” *Quarterly Journal of Economics* 114 (May): 655–90.

Erceg, Chris, Dale Henderson, and Andrew Levin. 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics* 46 (October): 281–313.

Fernández-Villaverde, Jesús, and Juan F. Rubio-Ramírez. 2007. "How Structural are Structural Parameters?" *NBER Macroeconomics Annual 2007,* edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. Chicago: University of Chicago Press, 83-137.

Galí, Jordi, and Pau Rabanal. 2004. "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?" *NBER Macroeconomics Annual 2004*: 225–88.

Gertler, Mark, and Antonella Trigari. 2006. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." Working Paper 12498. Cambridge, Mass.: National Bureau of Economic Research. (August).

Gertler, Mark, and John Leahy. 2006. "A Phillips Curve with an Ss Foundation." Working Paper 11971. Cambridge, Mass.: National Bureau of Economic Research. (January).

Ireland, Peter. 2004a. "Technology Shocks in the New Keynesian Model." *Review of Economic and Statistics* 86: 923–36.

Ireland, Peter. 2004b. "Money's Role in the Monetary Business Cycle." *Journal of Money, Credit, and Banking* 36 (December): 969–83.

Ireland, Peter. 2007. "Changes in the Federal Reserve's Inflation Target: Causes and Consequences." *Journal of Money, Credit, and Banking* 39 (December): 1851–82.

Justiniano, Alejandro, and Giorgio E. Primiceri. 2008. "The Time-Varying Volatility of Macroeconomic Fluctuations." *American Economic Review* 98 (June): 604–41.

Klenow, Pete, and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (August): 863–904.

Krause, Michael, and Thomas Lubik. 2007. "The (Ir)relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions." *Journal of Monetary Economics* 54 (April): 706–27.

Krause, Michael, David López-Salido, and Thomas Lubik. 2008. "Inflation Dynamics with Search Frictions: A Structural Econometric Analysis." *Journal of Monetary Economics* 55 (July): 892–916.

Kurmann, André. 2005. "Quantifying the Uncertainty about a Forward-Looking, New Keynesian Pricing Model." *Journal of Monetary Economics* 52: 1119–34.

Kurmann, André. 2007. "VAR-Based Estimation of Euler Equations with an Application to New Keynesian Pricing." *Journal of Economic Dynamics & Control* 31 (March): 767–96.

Laforte, Jean-Philippe. 2007. “Pricing Models: A Bayesian DSGE Approach for the U.S. Economy.” *Journal of Money, Credit, and Banking* 39: 127–54.

Leeper, Eric, and Christopher Sims. 1994. “Toward a Modern Macroeconomic Model Usable for Policy Analysis.” In *NBER Macroeconomics Annual 1994*, edited by Stanley Fischer and Julio Rotemberg. Cambridge, Mass.: MIT Press, 81–117.

Levin, Andrew, Alexei Onatski, John Williams, and Noah Williams. 2006. “Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models.” In *NBER Macroeconomics Annual 2005*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, Mass.: MIT Press, 229–88.

Lindé, Jesper. 2005. “Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach.” *Journal of Monetary Economics* 52 (September): 1135–49.

Lubik, Thomas, and Frank Schorfheide. 2004. “Testing for Indeterminacy: An Application to U.S. Monetary Policy.” *American Economic Review* 94 (March): 190–217.

Lubik, Thomas, and Frank Schorfheide. 2006. “A Bayesian Look at New Open Economy Macroeconomics.” In *NBER Macroeconomics Annual 2005*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, Mass.: MIT Press, 313–66.

Mackowiak, Bartosz, and Mirko Wiederholt. 2007. “Business Cycle Dynamics under Rational Inattention.” Manuscript, Northestern University and European Central Bank.

Mankiw, Gregory, and Ricardo Reis. 2007. “Sticky Information in General Equilibrium.” *Journal of The European Economic Association* 5: 603–14.

Milani, Fabio. 2007. “Expectations, Learning and Macroeconomic Persistence.” *Journal of Monetary Economics* 54 (October): 2065–83.

Nason, James, and Gregor Smith. 2008. “Identifying the New Keynesian Phillips Curve.” *Journal of Applied Econometrics* 23: 525–51.

Rabanal, Pau. 2007. “Does Inflation Increase after a Monetary Policy Tightening? Answers Based on an Estimated DSGE Model.” *Journal of Economics Dynamics & Control* 31 (March): 906–37.

Rabanal, Pau, and Juan F. Rubio-Ramírez. 2005. “Comparing New Keynesian Models of the Business Cycle: A Bayesian Approach.” *Journal of Monetary Economics* 52 (September): 1151–66.

Rotemberg, Julio J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies*. 49 (October): 517–31.

Rotemberg, Julio J., and Michael Woodford. 1997. "An Optimization-Based Econometric Model for the Evaluation of Monetary Policy." *NBER Macroeconomics Annual 1997,* edited by Ben S. Bernanke and Julio Rotemberg. Cambridge, Mass.: MIT Press, 297–361.

Rudd, Jeremy, and Karl Whelan. 2007. "Modelling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit, and Banking* 39: 155–70.

Salemi, Michael K. 2006. "Econometric Policy Evaluation and Inverse Control." *Journal of Money, Credit, and Banking* 38 (October): 1737–64.

Sargent, Thomas. 1989. "Two Models of Measurements and the Investment Accelerator." *Journal of Political Economy* 97 (April): 251–87.

Sbordone, Argia. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49 (March): 265–92.

Sbordone, Argia. 2005. "Do Expected Future Marginal Costs Drive Inflation Dynamics?" *Journal of Monetary Economics* 52 (September): 1183–97.

Schorfheide, Frank. 2005. "Learning and Monetary Policy Shifts." *Review of Economic Dynamics* 8 (April): 392–419.

Sims, Christopher A. 2002. "Solving Linear Rational Expectations Models." *Computational Economics* 20 (October): 1–20.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (April): 665–90.

Smets, Frank, and Raf Wouters. 2003. "An Estimated Stochastic Dynamic General Equilibrium Model for the Euro Area." *Journal of the European Economic Association* 1: 1123–75.

Smets, Frank, and Raf Wouters. 2005. "Comparing Shocks and Frictions in U.S. and Euro Area Business Cycles: A Bayesian DSGE Approach." *Journal of Applied Econometrics* 20: 161–83.

Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (June): 586–606.

Stock, James H., and Mark W. Watson. 2001. "Vector Autoregressions." *Journal of Economic Perspectives* 15: 101–15.

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

Woodford, Michael. 2008. "Information-Constrained State-Dependent Pricing." Manuscript, Columbia University.

Economic Quarterly—Volume 94, Number 4—Fall 2008—Pages 435–465

Policy Implications of the New Keynesian Phillips Curve

Stephanie Schmitt-Grohé and Martín Uribe

The theoretical framework within which optimal monetary policy was studied before the arrival of the New Keynesian Phillips curve (NKPC), but after economists had become comfortable using dynamic, optimizing, general equilibrium models and a welfare-maximizing criterion for policy analysis, was one in which the central source of nominal nonneutrality was a demand for money. At center stage in this literature was the role of money as a medium of exchange (as in cash-in-advance models, money-in-the-utility-function models, or shopping-time models) or as a store of value (as in overlapping-generations models). In the context of this family of models a robust prescription for the optimal conduct of monetary policy is to set nominal interest rates to zero at all times and under all circumstances. This policy implication, however, found no fertile ground in the boardrooms of central banks around the world, where the optimality of zero nominal rates was dismissed as a theoretical oddity, with little relevance for actual central banking. Thus, theory and practice of monetary policy were largely disconnected.

The early 1990s witnessed a profound shift in monetary economics away from viewing the role of money primarily as a medium of exchange and toward viewing money—sometimes exclusively—as a unit of account. A key insight was that the mere assumption that product prices are quoted in units of fiat money can give rise to a theory of price level determination, even if money is physically nonexistent and even if fiscal policy is irrelevant for price

Stephanie Schmitt-Grohé is affiliated with Columbia University, CEPR, and NBER. She can be reached at stephanie.schmittgrohe@columbia.edu. Martín Uribe is affiliated with Columbia University and NBER. He can be reached at martin.uribe@columbia.edu. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. The authors would like to thank Andreas Hornstein and Alexander Wolman for many thoughtful comments.

level determination.[1] This theoretical development was appealing to those who regard modern payment systems as operating increasingly cashlessly. At the same time, nominal rigidities in the form of sluggish adjustment of product and factor prices gained prominence among academic economists. The incorporation of sticky prices into dynamic stochastic general equilibrium models gave rise to a policy tradeoff between output and inflation stabilization that came to be known as the New Keynesian Phillips curve.

The inessential role that money balances play in the New Keynesian literature, along with the observed actual conduct of monetary policy in the United States and elsewhere over the past 30 years, naturally shifted theoretical interest away from money growth rate rules and toward interest rate rules: In the work of academic monetary economists, Milton Friedman's celebrated k-percent growth path for the money supply gave way to Taylor's equally influential interest rate feedback rule.

In this article, we survey recent advancements in the theory of optimal monetary policy in models with a New Keynesian Phillips curve. Our survey identifies a number of important lessons for the conduct of monetary policy. First, optimal monetary policy is characterized by near price stability. This policy implication is diametrically different from the one that obtains in models in which the only nominal friction is a transactions demand for money. Second, simple interest rate feedback rules that respond aggressively to price inflation deliver near-optimal equilibrium allocations. Third, interest rate rules that respond to deviations of output from trend may carry significant welfare costs. Taken together, lessons one through three call for the adherence to an inflation targeting objective. Fourth, the zero bound on nominal interest rates does not appear to be a significant obstacle for the actual implementation of low and stable inflation. Finally, product price stability emerges as the overriding goal of monetary policy even in environments where not only goods prices but also factor prices are sticky.

Before elaborating on the policy implications of the NKPC, we provide some perspective by presenting a brief account of the state of the literature on optimal monetary policy before the advent of the New Keynesian revolution.

1. OPTIMAL MONETARY POLICY PRE-NKPC

Within the pre-NKPC framework, under quite general conditions, optimal monetary policy calls for a zero opportunity cost of holding money, a result known as the Friedman rule. In fiat money economies in which assets used for transactions purposes do not earn interest, the opportunity cost of holding money equals the nominal interest rate. Therefore, in the class of models

[1] This is the case, for instance, when the monetary stance is active and the fiscal stance is passive, which is the monetary/fiscal regime most commonly studied.

commonly used for policy analysis before the emergence of the NKPC, the optimal monetary policy prescribed that the riskless nominal interest rate—the return on federal funds, say—be set at zero at all times.

In the early literature, a demand for money is motivated in a variety of ways, including a cash-in-advance constraint (Lucas 1982), money in the utility function (Sidrauski 1967), a shopping-time technology (Kimbrough 1986), or a transactions-cost technology (Feenstra 1986). Regardless of how a demand for money is introduced, the intuition for why the Friedman rule is optimal in this class of model is straightforward: A zero nominal interest rate maximizes holdings of a good—real money balances—that has a negligible production cost. Another reason why the Friedman rule is optimal is that a positive interest rate can distort the efficient allocation of resources. For instance, in the cash-in-advance model with cash and credit goods, a positive interest rate distorts the allocation of private spending across these two types of goods. In models in which money ameliorates transaction costs or decreases shopping time, a positive interest rate introduces a wedge in the consumption-leisure choice.

To illustrate the optimality of the Friedman rule, we augment a neoclassical model with a transaction technology that is decreasing in real money holdings and increasing in consumption spending. Specifically, consider an economy populated by a large number of identical households. Each household has preferences defined over processes of consumption and leisure and described by the utility function

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, h_t), \tag{1}$$

where c_t denotes consumption, h_t denotes labor effort, $\beta \in (0, 1)$ denotes the subjective discount factor, and E_0 denotes the mathematical expectation operator conditional on information available in period 0. The single period utility function, U, is assumed to be increasing in consumption, decreasing in effort, and strictly concave.

Final goods are produced using a production function, $z_t F(h_t)$, that takes labor, h_t, as the only factor input and is subject to an exogenous productivity shock, z_t.

A demand for real balances is introduced into the model by assuming that money holdings, denoted M_t, facilitate consumption purchases. Specifically, consumption purchases are subject to a proportional transaction cost, $s(v_t)$, that is decreasing in the household's money-to-consumption ratio, or consumption-based money velocity,

$$v_t = \frac{P_t c_t}{M_t}, \tag{2}$$

where P_t denotes the nominal price of the consumption good in period t. The transaction cost function, $s(v)$, satisfies the following assumptions: (a) $s(v)$

is nonnegative and twice continuously differentiable; (b) there exists a level of velocity, $\underline{v} > 0$, to which we refer as the satiation level of money, such that $s(\underline{v}) = s'(\underline{v}) = 0$; (c) $(v - \underline{v})s'(v) > 0$ for $v \neq \underline{v}$; and (d) $2s'(v) + vs''(v) > 0$ for all $v \geq \underline{v}$. Assumption (b) ensures that the Friedman rule (i.e., a zero nominal interest rate) need not be associated with an infinite demand for money. It also implies that both the transaction cost and the distortion it introduces vanish when the nominal interest rate is zero. Assumption (c) guarantees that in equilibrium money velocity is always greater than or equal to the satiation level. Assumption (d) ensures that the demand for money is decreasing in the nominal interest rate.

Households are assumed have access to risk-free pure discount bonds, denoted B_t. These bonds are assumed to carry a gross nominal interest rate of R_t when held from period t to period $t+1$. The flow budget constraint of the household in period t is then given by

$$P_t c_t[1 + s(v_t)] + P_t \tau_t^L + M_t + \frac{B_t}{R_t} = M_{t-1} + B_{t-1} + P_t z_t F(h_t), \qquad (3)$$

where τ_t^L denotes real lump sum taxes. In addition, it is assumed that the household is subject to a borrowing limit that prevents it from engaging in Ponzi-type schemes. The government is assumed to follow a fiscal policy whereby it rebates any seigniorage income it receives from the creation of money in a lump sum fashion to households.

A stationary competitive equilibrium can be shown to be a set of plans $\{c_t, h_t, v_t\}$, satisfying the following three conditions:

$$v_t^2 s'(v_t) = \frac{R_t - 1}{R_t}, \qquad (4)$$

$$-\frac{U_h(c_t, h_t)}{U_c(c_t, h_t)} = \frac{z_t F'(h_t)}{1 + s(v_t) + v_t s'(v_t)}, \text{ and} \qquad (5)$$

$$[1 + s(v_t)]c_t = z_t F(h_t), \qquad (6)$$

given monetary policy $\{R_t\}$, with $R_t \geq 1$, and the exogenous process $\{z_t\}$. The first equilibrium condition can be interpreted as a demand for money or liquidity preference function. Given our maintained assumptions about the transactions technology, $s(v_t)$, the implied money demand function is decreasing in the gross nominal interest rate, R_t. Further, our assumptions imply that as the interest rate vanishes, or R_t approaches unity, the demand for money reaches a finite maximum level given by $c_t/\underline{v}$. At this level of money demand, households are able to perform transactions costlessly, as the transactions cost, $s(v_t)$, becomes nil. The second equilibrium condition shows that a level of money velocity above the satiation level, $\underline{v}$, or, equivalently, an interest rate greater than zero, introduces a wedge between the marginal rate of substitution of consumption for leisure and the marginal product of labor. This wedge, given by $1 + s(v_t) + v_t s'(v_t)$, induces households to move away from consumption and toward leisure. The wedge is increasing in the nominal interest rate, implying that the larger is the nominal interest rate, the more distorted is the consumption-leisure choice. The final equilibrium condition states that a positive interest rate entails a resource loss in the amount of $s(v_t)c_t$. This resource loss is increasing in the interest rate and vanishes only when the nominal interest rate equals zero.

We wish to characterize optimal monetary policy under the assumption that the government has the ability to commit to policy announcements. This policy optimality concept is known as Ramsey optimality. In the context of the present model, the Ramsey optimal monetary policy consists in choosing the path of the nominal interest rate that is associated with the competitive equilibrium that yields the highest level of welfare to households. Formally, the Ramsey policy consists in choosing processes R_t, c_t, h_t, and v_t to maximize the household's utility function given in equation (1) subject to the competitive equilibrium conditions given by equations (4) through (6).

When one inspects the three equilibrium conditions above, it is clear that if the policymaker sets the monetary policy instrument, which we take to be the nominal interest rate, such that velocity is at the satiation level, ($v_t = \underline{v}$), then the equilibrium conditions become identical to an economy without the money demand friction, i.e., $c_t = z_t F(h_t)$ and $-U_h(c_t, h_t)/U_c(c_t, h_t) = z_t F'(h_t)$. Because the real allocation in the absence of the monetary friction is Pareto optimal, the proposed monetary policy must be Ramsey optimal. By a Pareto optimal allocation we mean a feasible real allocation (i.e., one satisfying $c_t = z_t F[h_t]$) with the property that any other feasible allocation that makes at least one agent better off also makes at least one agent worse off. It follows from equation (4) that setting the nominal interest rate to zero ($R_t = 1$) ensures that $v_t = \underline{v}$. For this reason, optimal monetary policy takes the form of a zero nominal interest rate at all times.

Under the optimal monetary policy, the rate of change in the aggregate price level varies over time. Because, to a first approximation, the nominal interest rate equals the sum of the real interest rate and the expected rate of inflation, and because under the optimal monetary policy the nominal interest rate is held constant, the degree to which the inflation rate fluctuates depends on the equilibrium variations in the real rate of interest. In general, optimal monetary policy in a model in which a role for monetary policy arises solely from the presence of money demand is not characterized by inflation stabilization.

A second important consequence of optimal monetary policy in the context of the present model is that inflation is, on average, negative. This is because, with a zero nominal interest rate, the inflation rate equals, on average, the negative of the real rate of interest.

2. THE NKPC AND OPTIMAL MONETARY POLICY

The New Keynesian Phillips curve can be briefly defined as the dynamic output-inflation tradeoff that arises in a dynamic general equilibrium model populated by utility-maximizing households and profit-maximizing firms—such as the one laid out in the previous section—and augmented with some kind of rigidity in the adjustment of nominal product prices. The foundations of the NKPC were laid by Calvo (1983) and Rotemberg (1982). Woodford (1996) and Yun (1996) completed its development by introducing optimizing behavior on the part of firms facing Calvo-type dynamic nominal rigidities.

The most important policy implication of models featuring a New Keynesian Phillips curve is the optimality of price stability (see Goodfriend and King [1997] for an early presentation of this result). We will discuss the price stability result in a variety of theoretical models, including ones with a realistic set of real and nominal rigidities, policy instruments and policy constraints, and sources of aggregate fluctuations. We start, however, with the simplest structure within which the price stability result can be obtained. To this end, we strip the model presented in the previous section from its money demand friction and instead introduce costs of adjusting nominal product prices. In the resulting model, sticky prices represent the sole source of nominal friction.

To introduce sticky prices into the model of the previous section, assume that the consumption good, c_t, is a composite good made of a continuum of intermediate differentiated goods. The aggregator function is of the Dixit-Stiglitz type. Each household/firm unit is the monopolistic producer of one variety of intermediate goods. In turn, intermediate goods are produced using a technology like the one given in the previous section. The household/firm unit hires labor, $\tilde{h}_t$, from a perfectly competitive market.

The demand faced by the household/firm unit for the intermediate input that it produces is of the form $Y_t d(\tilde{P}_t/P_t)$, where Y_t denotes the level of aggregate demand, which is taken as exogenous by the household/firm unit; $\tilde{P}_t$ denotes the nominal price of the intermediate good produced by the household/firm unit; and P_t is the price of the composite consumption good. The demand function, $d(\cdot)$, is assumed to be decreasing in the relative price, $\tilde{P}_t/P_t$, and is assumed to satisfy $d(1) = 1$ and $-d'(1) \equiv \eta > 1$, where η denotes the price elasticity of demand for each individual variety of intermediate goods that prevails in a symmetric equilibrium. The restrictions on $d(1)$ and $d'(1)$ are necessary for the existence of a symmetric equilibrium. The monopolist sets the price of the good it supplies, taking the level of aggregate demand as given, and is constrained to satisfy demand at that price, that is, $z_t F(\tilde{h}_t) \geq Y_t d(\tilde{P}_t/P_t)$.

Price adjustment is assumed to be sluggish, à la Rotemberg (1982). Specifically, the household/firm unit faces a resource cost of changing prices that is quadratic in the inflation rate of the good it produces:

$$\text{Price adjustment cost} = \frac{\theta}{2}\left(\frac{\tilde{P}_t}{\tilde{P}_{t-1}} - 1\right)^2. \tag{7}$$

The parameter θ measures the degree of price stickiness. The higher is θ, the more sluggish is the adjustment of nominal prices. When θ equals zero, prices are fully flexible. The flow budget constraint of the household/firm unit in period t is then given by

$$c_t + \tau_t^L \le (1-\tau_t^D)w_t h_t + \left[\frac{\tilde{P}_t}{P_t}Y_t d\left(\frac{\tilde{P}_t}{P_t}\right) - w_t\tilde{h}_t - \frac{\theta}{2}\left(\frac{\tilde{P}_t}{\tilde{P}_{t-1}} - 1\right)^2\right],$$

where τ_t^D denotes an income tax/subsidy rate. We introduce this fiscal instrument as a way to offset the distortions arising from the presence of monopolistic competition. We restrict attention to a stationary symmetric equilibrium in which all household/firm units charge the same price for the intermediate good they produce. Letting $\pi_t \equiv P_t/P_{t-1}$ denote the gross rate of inflation, the complete set of equilibrium conditions is then given by

$$\pi_t(\pi_t - 1) = \frac{\eta c_t}{\theta}\left[\frac{w_t}{z_t F'(h_t)} - \frac{\eta-1}{\eta}\right] + \beta E_t \frac{U_c(c_{t+1}, h_{t+1})}{U_c(c_t, h_t)}\pi_{t+1}(\pi_{t+1} - 1), \tag{8}$$

$$-\frac{U_h(c_t, h_t)}{U_c(c_t, h_t)} = (1-\tau_t^D)w_t, \text{ and} \tag{9}$$

$$z_t F(h_t) - \frac{\theta}{2}(\pi_t - 1)^2 = c_t. \tag{10}$$

The above three equations provide solutions for the equilibrium processes of consumption, c_t, hours, h_t, and the real wage, w_t, given processes for the rate of inflation, π_t, and for the tax rate, τ_t^D, which we interpret to be outcomes of the monetary and fiscal policies in place, respectively.

The first equilibrium condition, equation (8), represents the NKPC, to which the current volume is devoted. It describes an equilibrium relationship between current inflation, π_t, the current deviation of marginal cost, $w_t/z_t F'(h_t)$, from marginal revenue, $(\eta - 1)/\eta$, and expected future inflation. Under full price flexibility, the firm would always set marginal revenue equal to marginal cost. However, in the presence of price adjustment costs, this practice is costly. To smooth out price changes over time, firms set prices to equate an average of current and future expected marginal costs to an average of current and future expected marginal revenues. This optimal price-setting behavior gives rise to a relation whereby, given expected inflation, current inflation is an increasing function of marginal costs. Intuitively, this relation is steeper the more flexible are prices (i.e., the lower is θ), the more competitive are product markets (i.e., the higher is η), and the higher is the current level of demand (i.e., the larger is c_t). At the same time, given marginal cost, current inflation is increasing in expected future inflation. This is because, with quadratic costs of changing nominal prices, a firm expecting higher inflation in the future would like to smooth out the necessary price adjustments over time by beginning to raise prices already in the current period.

We have derived the New Keynesian Phillips curve in the context of the Rotemberg (1982) model of price stickiness. However, a similar relationship emerges under other models of nominal rigidity, such as those due to Calvo (1983), Taylor (1993), Woodford (1996), and Yun (1996). For instance, in the Calvo-Woodford-Yun model, price stickiness arises because firms are assumed to receive an idiosyncratic random signal each period indicating whether they are allowed to reoptimize their posted prices. A difference between the Rotemberg and the Calvo-Woodford-Yun models is that the latter displays equilibrium price dispersion across firms even in the absence of aggregate uncertainty. However, up to first order, the NKPCs implied by the Rotemberg and Calvo-Woodford-Yun models are identical. Indeed, much of the literature on the NKPC focuses on a log-linear approximation of this key relationship, as in equation (11).

The second equilibrium condition presented in equation (9) states that the marginal rate of substitution of consumption for leisure is equated to the after-tax real wage rate. The third equilibrium condition, equation (10), is a resource constraint requiring that aggregate output net of price adjustment costs equal private consumption.

It is straightforward to establish that, in this economy, the optimal monetary policy, that is, the policy that maximizes the welfare of the representative household, is one in which the inflation rate is nil at all times. Formally, the optimal monetary policy must be consistent with an equilibrium in which

$$\pi_t = 1$$

for all $t \geq 0$. This result holds exactly provided the fiscal authority subsidizes labor income to a point that fully offsets the distortion arising from the existence of imperfect competition among intermediate goods producers. Specifically, the income tax rate, τ_t^D, must be set at a constant and negative level given by

$$\tau_t^D = \frac{1}{1-\eta}$$

for all $t \geq 0$.

To see that the proposed policy regime is optimal, we demonstrate that it implies a set of equilibrium conditions that coincide with the one that arises in an economy with fully flexible prices ($\theta = 0$) and perfect competition in product markets ($\eta = \infty$), such as the one analyzed in Section 1 in the absence of a money-demand distortion. In effect, when $\pi_t = 1$ for all t, equilibrium condition (10) collapses to $c_t = z_t F(h_t)$. In addition, under zero inflation the NKPC (equation [8]) reduces to $w_t = z_t F'(h_t)(\eta - 1)/\eta$. Using this expression along with the proposed optimal level for the income tax rate in the equilibrium labor supply, equation (9), yields the efficiency conditions $-U_h(c_t, h_t)/U_c(c_t, h_t) = z_t F'(h_t)$, which, together with the resource constraint $c_t = z_t F(h_t)$, constitute the equilibrium conditions of a perfectly competitive flexible-price economy. As we show in Section 1, the resource allocation in this economy is Pareto optimal.

3. THE OPTIMAL INFLATION RATE

At this point, it is of interest to summarize and compare the results in this section and in previous ones. We have shown that when prices are fully flexible and the only nominal friction is a demand for money, then optimal monetary policy takes the form of complete stabilization of the interest rate at a value of zero ($R_t = 1$ for all t). We have also established that in a cashless economy in which the only source of nominal friction is given by product price stickiness, optimal monetary policy calls for full stabilization of the rate of inflation at a value of zero ($\pi_t = 1$ for all t). Under optimal policy in the monetary flexible-price economy, inflation is time varying and equal to the negative of the real interest rate on average, whereas in the cashless sticky-price economy, inflation is constant and equal to zero at all times. Also, in the monetary flexible-price economy, optimal policy calls for a constant nominal interest rate equal to zero at all times, whereas in the cashless sticky-price economy, it calls for a time-varying nominal interest rate equal to the real interest rate on average.

These results raise the question of what the characteristics of optimal monetary policy are in a more realistic economic environment in which both a demand for money and price stickiness coexist. In particular, in such an environment a policy tradeoff emerges between the benefits of targeting zero inflation—i.e., minimizing price-adjustment costs—and the benefits of deflating at the real rate of interest—i.e., minimizing the opportunity cost of holding money. In the canonical economies with only one nominal friction studied in this and previous sections, the characterization of the optimal rate of inflation is relatively straightforward. As soon as both nominal frictions are incorporated jointly, it becomes impossible to determine the optimal rate of inflation analytically. One is therefore forced to resort to numerical methods.

The resolution of the Friedman-rule-versus-price-stability tradeoff was studied by, among others, Khan, King, and Wolman (2003) and Schmitt-Grohé and Uribe (2004a, 2007b). As one would expect, when both the money demand and sticky-price frictions are present, the optimal rate of inflation falls between zero and the one called for by the Friedman rule. The question of interest, however, is where exactly in this interval the optimal inflation rate lies. Khan, King, and Wolman find, in the context of a stylized model calibrated to match aspects of money demand and price dynamics in the postwar United States, that the optimal rate of inflation is -0.76 percent per year. By comparison, in their model the Friedman rule is associated with a deflation rate of 2.93 percent per year. Thus, in the study by Khan, King, and Wolman, the optimal policy is closer to price stability than to the Friedman rule. Taking these numbers at face value, one might conclude that price stickiness is the dominant friction in shaping optimal monetary policy. However, Khan, King, and Wolman (2003) and Schmitt-Grohé and Uribe (2004a, 2007b) show that the resolution of the tradeoff is quite sensitive to plausible changes in the values taken by the structural parameters of the model.

In Schmitt-Grohé and Uribe (2007b), we find that a striking characteristic of the optimal monetary regime is the high sensitivity of the welfare-maximizing rate of inflation with respect to the parameter governing the degree of price stickiness for the range of values of this parameter that is empirically relevant. The model underlying the analysis of Schmitt-Grohé and Uribe (2007b) is a medium-scale model of the U.S. economy featuring, in addition to money demand by households and sticky product prices, a number of real and nominal rigidities including wage stickiness, a demand for money by firms, habit formation, capital accumulation, variable capacity utilization, and investment adjustment costs. The structural parameters of the model are assigned values that are consistent with full- as well as limited-information approaches to estimating this particular model.

In the Schmitt-Grohé and Uribe (2007b) model, the degree of price stickiness is captured by a parameter denoted α, measuring the probability that a firm is not able to optimally set the price it charges in a particular quarter. The average number of periods elapsed between two consecutive optimal price adjustments is given by $1/(1-\alpha)$. Available empirical estimates of the degree of price rigidity using macroeconomic data vary from two to five quarters, or $\alpha \in [0.5, 0.8]$. For example, Christiano, Eichenbaum, and Evans (2005) estimate α to be 0.6. By contrast, Altig et al. (2005) estimate a marginal-cost-gap coefficient in the Phillips curve that is consistent with a value of α of around 0.8. Both Christiano, Eichenbaum and Evans (2005) and Altig et al. (2005) use an impulse-response matching technique to estimate the price-stickiness parameter α. Bayesian estimates of this parameter include Del Negro et al. (2004), Levin et al. (2006), and Smets and Wouters (2007), who report posterior means of 0.67, 0.83, and 0.66, respectively, and 90 percent posterior probability intervals of (0.51,0.83), (0.81,0.86), and (0.56,0.74), respectively.

Recent empirical studies have documented the frequency of price changes using microdata underlying the construction of the U.S. consumer price index. These studies differ in the sample period considered, in the disaggregation of the price data, and in the treatment of sales and stockouts. The median frequency of price changes reported by Bils and Klenow (2004) is four to five

Figure 1 Price Stickiness, Fiscal Policy, and Optimal Inflation

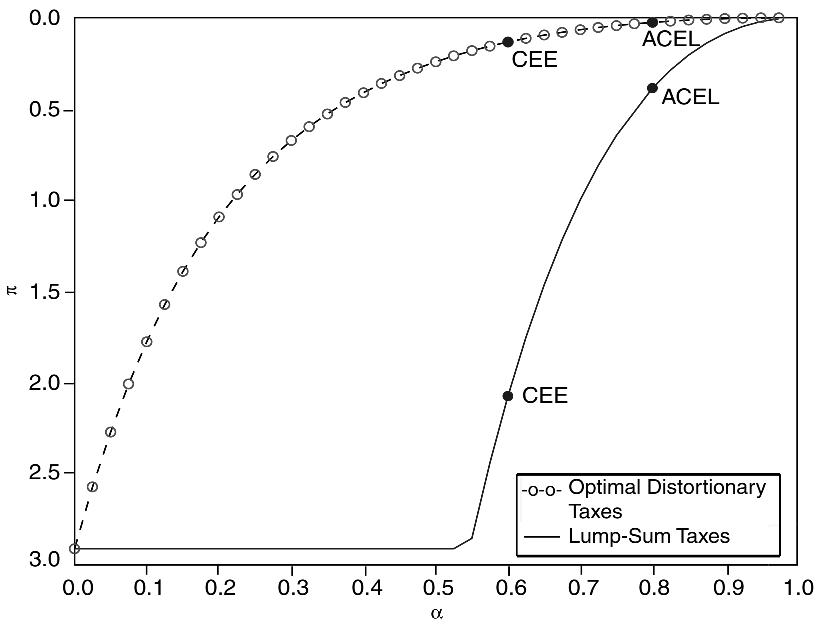

Notes: CEE and ACEL indicate, respectively, the values for the parameter, α, estimated by Christiano, Eichenbaum, and Evans (2005) and Altig et al. (2005).

months, the one reported by Klenow and Kryvtsov (2005) is four to seven months, and the one reported by Nakamura and Steinsson (2007) is eight to 11 months. However, there is no immediate interpretation of these frequency estimates to the parameter, α, governing the degree of price stickiness in Calvo-style models of price staggering. Consider, for instance, the case of indexation. In that case, even though firms change prices every period—implying the highest possible frequency of price changes—prices themselves may be highly sticky, for they may be only reoptimized at much lower frequencies.

Figure 1 displays with a solid line the relationship between the degree of price stickiness, α, and the optimal rate of inflation in percent per year, π, implied by the model studied in Schmitt-Grohé and Uribe (2007b). When α equals 0.5, the lower range of the available empirical evidence using macrodata, the optimal rate of inflation is -2.9 percent, which is the level called for by the Friedman rule. For a value of α of 0.8, which is near the upper range of the available empirical evidence using macrodata, the optimal level of inflation rises to -0.4 percent, which is close to price stability.

Besides the uncertainty surrounding the estimation of the degree of price stickiness, a second aspect of the apparent difficulty in establishing reliably the optimal long-run level of inflation has to do with the shape of the relationship linking the degree of price stickiness to the optimal level of inflation. The problem resides in the fact that, as is evident from Figure 1, this relationship becomes significantly steep precisely for that range of values of α that is empirically most compelling. It turns out that an important factor determining the shape of the function relating the optimal level of inflation to the degree of price stickiness is the underlying fiscal policy regime.

Fiscal considerations fundamentally change the long-run tradeoff between price stability and the Friedman rule. To see this, we now consider an economy where lump-sum taxes are unavailable ($\tau^L = 0$). Instead, the fiscal authority must finance government purchases by means of proportional capital and labor income taxes. The social planner jointly sets monetary and fiscal policy in a welfare-maximizing (i.e., Ramsey-optimal) fashion.[2] Figure 1 displays the relationship between the degree of price stickiness, α, and the optimal rate of inflation, π. The solid line corresponds to the case discussed earlier featuring lump-sum taxes. The dash-circled line corresponds to the economy with optimally chosen distortionary income taxes. In stark contrast to what happens under lump-sum taxation, under optimal distortionary taxation the function linking π and α is flat and very close to zero for the entire range of macrodata-based empirically plausible values of α, namely 0.5 to 0.8. In other words, when taxes are distortionary and optimally determined, price stability emerges as a prediction that is robust to the existing uncertainty about the exact degree of price stickiness. Even if one focuses on the evidence of price stickiness stemming from microdata, the model with distortionary Ramsey taxation predicts an optimal long-run level of inflation that is much closer to zero than to the level called for by the Friedman rule.

Our intuition for why price stability arises as a robust policy recommendation in the economy with optimally set distortionary taxation runs as follows. Consider the economy with lump-sum taxation. Deviating from the Friedman rule (by raising the inflation rate) has the benefit of reducing price adjustment costs. Consider next the economy with optimally chosen income taxation and no lump-sum taxes. In this economy, deviating from the Friedman rule still provides the benefit of reducing price adjustment costs. However, in this economy, increasing inflation has the additional benefit of increasing seigniorage revenue, thereby allowing the social planner to lower distortionary income tax rates. Therefore, the Friedman rule versus price stability tradeoff is tilted in favor of price stability.

[2] The details of this environment are contained in Schmitt-Grohé and Uribe (2006). The structure of this economy is identical to that studied in Schmitt-Grohé and Uribe (2007b), except for the inclusion of fiscal policy.

It follows from this intuition that what is essential in inducing the optimality of price stability is that, on the margin, the fiscal authority trades off the inflation tax for regular taxation. Indeed, it can be shown that if distortionary tax rates are fixed, even if they are fixed at the level that is optimal in a world without lump-sum taxes, and the fiscal authority has access to lump-sum taxes on the margin, the optimal rate of inflation is much closer to the Friedman rule than to zero. In this case, increasing inflation no longer has the benefit of reducing distortionary taxes. As a result, the Ramsey planner has less incentives to inflate.

We close this section by drawing attention to the fact that, quite independently of the precise degree of price stickiness, the optimal inflation target is below zero. In light of this robust result, it is puzzling that all countries that self-classify as inflation targeters set inflation targets that are positive. In effect, in the developed world inflation targets range between 2 and 4 percent per year. Somewhat higher targets are observed across developing countries. An argument often raised in defense of positive inflation targets is that negative inflation targets imply nominal interest rates that are dangerously close to the zero lower bound on nominal interest rates and, hence, may impair the central bank's ability to conduct stabilization policy. In Schmitt-Grohé and Uribe (2007b) we find, however, that this argument is of little relevance in the context of the medium-scale estimated model within which we conduct policy evaluation. The reason is that under the optimal policy regime, the mean of the nominal interest rate is about 4.5 percent per year with a standard deviation of only 0.4 percent. This means that for the zero lower bound to pose an obstacle to monetary stabilization policy, the economy must suffer from an adverse shock that forces the interest rate to be more than ten standard deviations below target. The likelihood of such an event is practically nil.

4. THE OPTIMAL VOLATILITY OF INFLATION

Two distinct branches of the existing literature on optimal monetary policy deliver diametrically opposed policy recommendations concerning the cyclical behavior of prices and interest rates. One branch follows the theoretical framework laid out in Lucas and Stokey (1983). It studies the joint determination of optimal fiscal and monetary policy in flexible-price environments with perfect competition in product and factor markets. In this strand of the literature, the government's problem consists of financing an exogenous stream of public spending by choosing the least disruptive combination of inflation and distortionary income taxes.

Calvo and Guidotti (1990, 1993) and Chari, Christiano, and Kehoe (1991) characterize optimal monetary and fiscal policy in stochastic environments with nominal nonstate-contingent government liabilities. A key result of these papers is that it is optimal for the government to make the inflation rate highly volatile and serially uncorrelated. For instance, Schmitt-Grohé and Uribe (2004b) show, in the context of a flexible-price model calibrated to the U.S. economy, that under the optimal policy the inflation rate has a standard deviation of 7 percent per year and a serial correlation of -0.03. The intuition for this result is that, under flexible prices, highly volatile and unforecastable inflation is nondistorting and at the same time carries the fiscal benefit of acting as a lump-sum tax on private holdings of government-issued nominal assets. The government is able to use surprise inflation as a nondistorting tax to the extent that it has nominal, nonstate-contingent liabilities outstanding. Thus, price changes play the role of a shock absorber of unexpected innovations in the fiscal deficit. This "front-loading" of government revenues via inflationary shocks allows the fiscal authority to keep income tax rates remarkably stable over the business cycle.

However, as discussed in Section 2, the New Keynesian literature, aside from emphasizing the role of price rigidities and market power, differs from the earlier literature described above in two important ways. First, it assumes, either explicitly or implicitly, that the government has access to (endogenous) lump-sum taxes to finance its budget. An important implication of this assumption is that there is no need to use unanticipated inflation as a lump-sum tax; regular lump-sum taxes take on this role. Second, the government is assumed to be able to implement a production (or employment) subsidy to eliminate the distortion introduced by the presence of monopoly power in product and factor markets.

The key result of the New Keynesian literature, which we presented in Sections 2 and 3, is that the optimal monetary policy features an inflation rate that is zero or close to zero at all times (i.e., both the optimal mean and volatility of inflation are near zero). The reason price stability is optimal in environments of the type described there is that it minimizes (or completely eliminates) the costs introduced by inflation under nominal rigidities.

Together, these two strands of research on optimal monetary policy leave the monetary authority without a clear policy recommendation. Should the central bank pursue policies that imply high or low inflation volatility? In Schmitt-Grohé and Uribe (2004a), we analyze the resolution of this policy dilemma by incorporating in a unified framework the essential elements of the two approaches to optimal policy described above. Specifically, we build a model that shares two elements with the earlier literature: (a) The only source of regular taxation available to the government is distortionary income taxes. As a result, the government cannot implement production subsidies to undo distortions created by the presence of imperfect competition, and (b) the government issues only nominal, one-period, nonstate-contingent bonds. At the same time, the setup shares two important assumptions with the more recent body of work on optimal monetary policy: (a) Product markets are imperfectly competitive, and (b) product prices are assumed to be sticky and, hence, the model features a New Keynesian Phillips curve. Schmitt-Grohé and Uribe (2004a) introduce price stickiness as in the previous section by assuming that firms face a convex cost of price adjustment (Rotemberg 1982).

In this environment, the government faces a tradeoff in choosing the path of inflation. On the one hand, the government would like to use unexpected inflation as a nondistorting tax on nominal wealth. In this way, the fiscal authority could minimize variations in distortionary income taxes over the business cycle. On the other hand, changes in the rate of inflation come at a cost, for firms face nominal rigidities.

When price changes are brought about at a cost, it is natural to expect that a benevolent government will try to implement policies consistent with a more stable behavior of prices than when price changes are costless. However, the quantitative effect of an empirically plausible degree of price rigidity on optimal inflation volatility is not clear a priori. In Schmitt-Grohé and Uribe (2004a), we show that for the degree of price stickiness estimated for the U.S. economy, this tradeoff is overwhelmingly resolved in favor of price stability. The Ramsey allocation features a dramatic drop in the standard deviation of inflation from 7 percent per year under flexible prices to a mere 0.17 percent per year when prices adjust sluggishly.[3]

Indeed, the impact of price stickiness on the optimal degree of inflation volatility turns out to be much stronger than suggested by the numerical results reported in the previous paragraph. Figure 2, taken from Schmitt-Grohé and Uribe (2004a), shows that a minimum amount of price stickiness suffices to make price stability the central goal of optimal policy. Specifically, when the degree of price stickiness, embodied in the parameter θ (see equation [7]), is assumed to be ten times smaller than the estimated value for the U.S. economy, the optimal volatility of inflation is below 0.52 percent per year, 13 times smaller than under full price flexibility.

A natural question elicited by Figure 2 is why even a modest degree of price stickiness can turn undesirable the use of a seemingly powerful fiscal instrument, such as large revaluations or devaluations of private real financial wealth through surprise inflation. Our conjecture is that in the flexible-price economy, the welfare gains of surprise inflations or deflations are very small. Our intuition is as follows. Under flexible prices, it is optimal for the central bank to keep the nominal interest rate constant over the business cycle. This means that large surprise inflations must be as likely as large deflations, as variations in real interest rates are small. In other words, inflation must have a near-i.i.d. behavior. As a result, high inflation volatility cannot be used by the Ramsey planner to reduce the average amount of resources to be collected via distortionary income taxes, which would be a first-order effect. The volatility

[3] This price stability result is robust to augmenting the model to allow for nominal rigidities in wages and indexation in product or factor prices (Schmitt-Grohé and Uribe 2006, Table 6.5).

Figure 2 Degree of Price Stickiness and Optimal Inflation Volatility

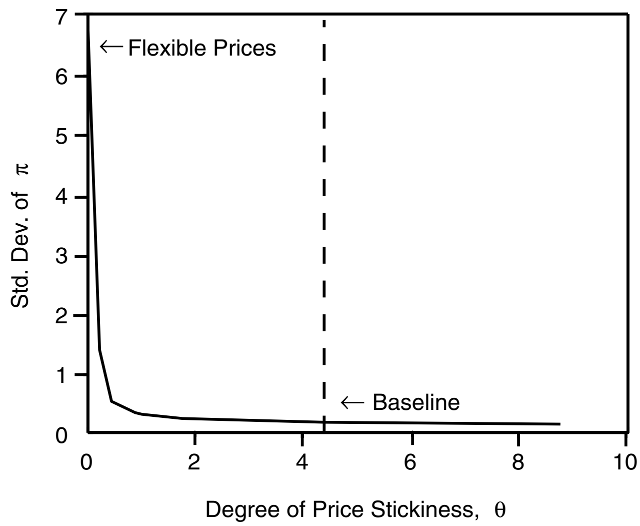

Notes: The parameter, θ, governs the cost of adjusting nominal prices as defined in equation (7). Its baseline value is 4.4, in line with available empirical estimates. The standard deviation of inflation is measured in percent per year.

of inflation primarily serves the purpose of smoothing the process of income tax distortions—a second-order source of welfare losses—without affecting their average level.

Another way to gain intuition for the dramatic decline in optimal inflation volatility that occurs even at very modest levels of price stickiness is to interpret price volatility as a way for the government to introduce real state-contingent public debt. Under flexible prices, the government uses state-contingent changes in the price level as a nondistorting tax or transfer on private holdings of government assets. In this way, nonstate-contingent nominal public debt becomes state-contingent in real terms. So, for example, in response to an unexpected increase in government spending, the Ramsey planner does not need to increase tax rates by much because by inflating away part of the public debt he can ensure intertemporal budget balance. It is, therefore, clear that introducing costly price adjustment is the same as if the government were limited in its ability to issue real state-contingent debt. It follows that the larger the welfare gain associated with the ability to issue real state-contingent public debt—as opposed to nonstate-contingent debt—the larger the amount of price stickiness required to reduce the optimal degree of inflation volatility. Aiyagari et al. (2002) show that indeed the level of welfare under the Ramsey policy in an economy without real state-contingent public debt is virtually the same as in an economy with state-contingent debt. Our finding that a small amount of price stickiness is all it takes to bring the optimal volatility of inflation from a very high level to near zero is thus perfectly in line with the finding of Aiyagari et al. (2002).

If this intuition is correct, then the behavior of tax rates and public debt under sticky prices should resemble that implied by the Ramsey allocation in economies without real state-contingent debt. Indeed, in financing the budget, the Ramsey planner replaces front-loading with standard debt and tax instruments. For example, in response to an unexpected increase in government spending, the planner does not generate a surprise increase in the price level. Instead, he chooses to finance the increase in government purchases partly through an increase in income tax rates and partly through an increase in public debt. The planner minimizes the tax distortion by spreading the required tax increase over many periods. This tax-smoothing behavior induces near-random walk dynamics into the tax rate and public debt. By contrast, under full price flexibility (i.e., when the government can create real-state-contingent debt), tax rates and public debt inherit the stochastic process of the underlying shocks.

An important conclusion of this analysis is, thus, that the Aiyagari et al. (2002) result, namely, that optimal policy imposes a near-random walk behavior on taxes and debt, does not require the unrealistic assumption that the government can issue only nonstate-contingent real debt. This result emerges naturally in economies with nominally nonstate-contingent debt—clearly the case of greatest empirical relevance—and a minimum amount of price rigidity.[4] However, if government debt is assumed to be state contingent, the presence of sticky prices may introduce no difference in the Ramsey real allocation, depending on the precise specification of the demand for money (see Correia, Nicolini, and Teles 2008). The reason for this result is that, as shown in Lucas and Stokey (1983), if government debt is state-contingent and prices are fully flexible, the Ramsey allocation does not pin down the price level uniquely. In this case, there is an infinite number of price-level processes (and thus of money supply processes) that can be supported as Ramsey outcomes.

[4] It is of interest to relate the near-random walk in taxes and debt that emerges as the optimal policy outcome in a model featuring a New Keynesian Phillips curve with the celebrated tax-smoothing result of Barro (1979). In Barro's formulation, the objective function of the government is the expected present discounted value of squared deviations of tax rates from a target or desired level. The government minimizes this objective function subject to a sequential budget constraint, which is linear in debt and tax rates. The resulting solution resembles the random walk model of consumption with taxes taking the place of consumption and public debt taking the place of private debt. The analysis in Schmitt-Grohé and Uribe (2004a) departs from Barro's ad hoc loss function and replaces it with the utility function of the representative optimizing household inhabiting a fully-articulated, dynamic, stochastic, general-equilibrium economy. In this environment, the random walk result obtains from a more subtle channel, namely, the introduction of a miniscule amount of nominal rigidity in product prices.

Loosely speaking, the introduction of price stickiness simply "uses this degree of freedom" to pin down the equilibrium process of the price level without altering other aspects of the Ramsey solution.

5. IMPLEMENTATION OF OPTIMAL POLICY

We established that in the simple New Keynesian model presented in Section 2, the optimal policy consists of setting the inflation rate equal to zero at all times ($\pi_t = 1$) and imposing a constant output subsidy ($\tau_t^D = 1/[1 - \eta]$). The question we pursue in this section is how to implement the optimal policy. Because central banks in the United States and elsewhere use the short-term nominal interest rate as the monetary policy instrument, it is of empirical interest to search for interest rate rules that implement the optimal allocation.

Using the Ramsey-Optimal Interest Rate Process as a Feedback Rule

One might be tempted to believe that implementation of optimal policy is trivial once the interest rate associated with the Ramsey equilibrium has been found. Specifically, in the Ramsey equilibrium, the nominal interest rate can be expressed as a function of the current state of the economy. Then, the prescription would be simply to use this function as a policy rule in setting the nominal interest rate at all dates and under all circumstances. It turns out that conducting policy in this fashion would, in general, not deliver the intended results. The reason is that although such a policy would be consistent with the optimal equilibrium, it would at the same time open the door to other (suboptimal) equilibria. It follows that the solution to the optimal policy problem is mute with respect to the issue of implementation of such policy. To see this, it is convenient to consider as an example a log-linear approximation to the equilibrium conditions associated with the cashless, sticky-price model presented in Section 2. It can be shown that the resulting linear system is given by[5]

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa \hat{c}_t - \gamma \hat{z}_t \tag{11}$$

and

$$-\sigma \hat{c}_t = \hat{R}_t - \sigma E_t \hat{c}_{t+1} - E_t \hat{\pi}_{t+1}, \tag{12}$$

where σ, κ, and $\gamma > 0$ are parameters. Hatted variables denote percent deviations of the corresponding nonhatted variables from their respective values

[5] The log-linearization is performed around the nonstochastic steady state of the Ramsey equilibrium. In performing the linearization, we assume that the period utility function is separable in consumption and hours and that the production function is linear in labor.

in the deterministic steady state of the Ramsey equilibrium. Equation (11) results from combining equations (8), (9), and (10) and is typically referred to as the New Keynesian Phillips curve.[6] Equation (12) is a linearized version of an Euler equation that prices nominally risk-free bonds, where R_t denotes the gross nominal risk-free interest rate between periods t and $t+1$. This equation is typically referred to as the intertemporal IS equation.

Substituting the welfare-maximizing rate of inflation, $\hat{\pi}_t = 0$, into the intertemporal IS curve (12) implies that the nominal interest rate is given by $\hat{R}_t = \sigma E_t(\hat{c}_{t+1} - \hat{c}_t)$, which states that under the optimal policy the nominal and real interest rates coincide. Suppose that the central bank adopted this expression as a policy feedback rule for the nominal interest rate. The question is whether this proposed rule implements the Ramsey equilibrium uniquely. The answer to this question is no. To see why, consider a solution of the form $\hat{\pi}_t = \epsilon_t$, where ϵ_t is i.i.d. normal with mean zero and an arbitrary standard deviation $\sigma_\epsilon \geq 0$. Notice that for all positive values of σ_ϵ, the proposed solution for inflation is different from the optimal one. It is straightforward to see that the proposed solution satisfies the intertemporal IS equation (12). The solution for consumption can be read off the NKPC as being $\hat{c}_t = \gamma/\kappa \hat{z}_t + (1/\kappa)\epsilon_t$. We have, therefore, constructed a competitive equilibrium in which a nonfundamental source of uncertainty, embodied in the random variable ϵ_t, causes stochastic deviations of consumption and inflation from their optimal paths. Notice that in this example there exists an infinite number of different equilibria indexed by the parameter, σ_ϵ, governing the volatility of the nonfundamental shock, ϵ_t.

One possible objection against the interest rate feedback rule proposed in the previous paragraph is that it is cast in terms of the endogenous variable, c_t. In particular, one may wonder whether this endogeneity is responsible for the inability of the proposed rule to implement the Ramsey equilibrium uniquely. This concern is indeed unfounded. For, even if the interest rate feedback rule were cast in terms of exogenous fundamental variables, the failure of the strategy of using the Ramsey solution for R_t as an interest rate feedback rule remains. Specifically, substituting the optimal rate of inflation, $\hat{\pi}_t = 0$, into the New Keynesian Phillips curve (11) yields $\hat{c}_t = \gamma\kappa^{-1}\hat{z}_t$. In turn, substituting this expression into the intertemporal IS curve (12) implies that in the optimal equilibrium the nominal interest rate is given by $\hat{R}_t = \hat{r}_t^n \equiv \sigma\gamma\kappa^{-1}E_t(\hat{z}_{t+1} - \hat{z}_t)$. The variable r_t^n denotes the risk-free real (as well as nominal) interest rate that prevails in the Ramsey optimal equilibrium and is referred to as the natural rate of interest. Using this expression, equations (11)–(12) become a system of two linear stochastic difference equations

[6] For detailed derivations of this expression, see, for instance, Woodford (2003). This linear expression is the NKPC studied in the papers by Nason and Smith (2008) and Schorfheide (2008) that appear in this issue.

in the endogenous variables $\hat{\pi}_t$ and $\hat{c}_t$. This system possesses one eigenvalue inside the unit circle and one outside. It follows that the solution for inflation and consumption is of the form $\hat{\pi}_{t+1} = \zeta_{\pi\pi}\hat{\pi}_t + \zeta_{\pi z}\hat{z}_t + \epsilon_{t+1}$ and $\hat{c}_t = \zeta_{c\pi}\hat{\pi}_t + \zeta_{cz}\hat{z}_t$, where $\hat{\pi}_0$ is arbitrary, ϵ_t is a nonfundamental shock such as the one introduced above, and the parameter $\zeta_{\pi\pi}$ is less than unity in absolute value. Clearly, the competitive equilibrium that the proposed rule implements displays persistent and stochastic deviation from the optimal solution. We conclude that the use of the Ramsey-optimal interest rate process as a policy feedback rule fails to implement the desired competitive equilibrium.

Can the Taylor Rule Implement Optimal Policy?

A Taylor-type rule is an interest rate feedback rule whereby the nominal interest rate is set as an increasing linear function of inflation and deviations of real output from trend, with an inflation coefficient greater than unity and an output coefficient greater than zero. Formally, a Taylor-type interest rate rule can be written as $\hat{R}_t = \alpha_\pi\hat{\pi}_t + \alpha_y\hat{y}_t$, where $\alpha_\pi > 1$ and $\alpha_y > 0$ are parameters and $\hat{y}_t$ represents the percent deviation of real output from trend. Taylor's rule has been widely studied in monetary economics since the publication of Taylor's (1993) seminal article. It has been advocated as a desirable policy specification and is considered by some to be a reasonable approximation of actual monetary policy in the United States and many other developed countries. For this reason, we now consider the question of whether a Taylor rule can implement the optimal allocation. The answer is that, in general, an interest rate feedback rule of the type proposed by Taylor is unable to support the Ramsey optimal equilibrium. This issue was first analyzed by Woodford (2001).

To establish whether the Taylor rule presented above can implement the optimal allocation, we set the inflation rate at its optimal value of zero ($\hat{\pi}_t = 0$) and combine the intertemporal IS equation (12) with the Taylor rule. This yields $\alpha_y\hat{c}_t = \sigma(E_t\hat{c}_{t+1} - \hat{c}_t)$. Now, replacing $\hat{c}_t$ with its optimal value of $\gamma/\kappa\hat{z}_t$, we obtain $E_t\hat{z}_{t+1} = (1 + \alpha_y/\sigma)\hat{z}_t$. This expression represents a contradiction because the productivity shock, $\hat{z}_t$, is assumed to be an exogenous stationary process with a law of motion independent of the policy parameter, α_y, and the preference parameter, σ. We have established that the proposed Taylor rule fails to implement the optimal equilibrium. One can show that this result also obtains when the nominal interest rate is assumed to respond to deviations of output from its natural level, which is defined as the level of output associated with the optimal equilibrium and is given by $y_t^n \equiv \gamma/\kappa\hat{z}_t$.

However, the optimal allocation can indeed be implemented by a modified Taylor rule of the form $\hat{R}_t = \hat{r}_t^n + \alpha_\pi\hat{\pi}_t$ as long as $\alpha_\pi > 1$. In this rule, $\hat{r}_t^n$ denotes the natural rate of interest defined earlier. The first term in the modified Taylor rule makes the Ramsey allocation feasible as an equilibrium outcome.

The second term makes it unique. The key difference between the standard and modified Taylor rules is that the latter features a time-varying intercept that allows the nominal interest rate to accommodate movements in the real interest rate one-for-one without requiring changes in the price level. More generally, the optimal competitive equilibrium can be implemented via rules of the form $\hat{R}_t = \hat{r}_t^n + \alpha_\pi \hat{\pi}_t + \alpha_y(\hat{y}_t - \hat{y}_t^n)$, with policy parameters α_π and α_y satisfying the restrictions imposed by the definition of a Taylor-type rule given above.

Applying this type of rule can be quite impractical, for it would require knowledge on the part of the central bank of current and expected future values taken by all of the shocks that affect the real interest rate, as well as of the function mapping such values to the natural rate of interest. This difficulty raises the question of how close a less sophisticated interest rate rule would get to implementing the optimal equilibrium. We turn to this issue next.

Optimal, Simple, and Implementable Rules

In this subsection, we analyze the ability of simple, implementable interest rate rules to approximate the outcome of optimal policy. We draw from our previous work (Schmitt-Grohé and Uribe 2007a), where we evaluate policy in the context of a calibrated model of the U.S. business cycle featuring monopolistic competition, sticky prices in product markets, capital accumulation, government purchases financed by lump-sum or distortionary taxes, and with or without a transactional demand for money.[7] In the model, business cycles are driven by stochastic variations in the level of total factor productivity and government consumption. We impose two requirements for an interest rate rule to be implementable. First, the rule must deliver a unique rational expectations equilibrium. Second, it must induce nonnegative equilibrium dynamics for the nominal interest rate. For an interest rule to be simple, we require that the interest rate be set as a function of a small number of easily observable macroeconomic indicators. Specifically, we study interest rate feedback rules that respond to measures of inflation, output, and lagged values of the nominal interest rate. The family of rules we consider is of the form

$$\begin{aligned} \ln(R_t/R^*) &= \alpha_R \ln(R_{t-1}/R^*) + \alpha_\pi E_t \ln(\pi_{t-i}/\pi^*) + \alpha_y E_t \ln(y_{t-i}/y^*); \\ i &= -1, 0, 1, \end{aligned} \tag{13}$$

where y^* denotes the nonstochastic Ramsey steady-state level of aggregate demand, and R^*, π^*, α_R, α_π, and α_y are parameters. The index, i, can take three values: 1, 0, and -1. When $i = 1$, we refer to the interest rate rule as backward-looking, when $i = 0$ as contemporaneous, and when $i = -1$ as

[7] See also Rotemberg and Woodford (1997).

Table 1 Evaluating Interest Rate Rules

	α_π	α_y	α_R	**Welfare Cost**	σ_π	σ_R
Ramsey Policy	—	—	—	0.000	0.01	0.27
Optimized Rules						
Contemporaneous ($i = 0$)						
Smoothing	3	0.01	0.84	0.000	0.04	0.29
No Smoothing	3	0.00	—	0.001	0.14	0.42
Backward ($i = 1$)	3	0.03	1.71	0.001	0.10	0.23
Forward ($i = -1$)	3	0.07	1.58	0.003	0.19	0.27
Nonoptimized Rules						
Taylor Rule ($i = 0$)	1.5	0.5	—	0.522	3.19	3.08
Simple Taylor Rule	1.5	—	—	0.019	0.58	0.87
Inflation Targeting	—	—	—	0.000	0.00	0.27

Notes: The interest rate rule is given by $\ln(R_t/R^*) = \alpha_R \ln(R_{t-1}/R^*) + \alpha_\pi E_t \ln(\pi_{t-i}/\pi^*) + \alpha_y E_t \ln(y_{t-i}/y^*)$; $i = -1, 0, 1$. In the optimized rules, the policy parameters α_π, α_y,and α_R are restricted to lie in the interval $[0, 3]$. The welfare cost is defined as the percentage decrease in the Ramsey-optimal consumption process necessary to make the level of welfare under the Ramsey policy identical to that under the evaluated policy. Thus, a positive figure indicates that welfare is higher under the Ramsey policy than under the alternative policy. The standard deviation of inflation and the nominal interest rate is measured in percent per year.

forward-looking. The optimal simple and implementable rule is the simple and implementable rule that maximizes welfare of the representative agent. Specifically, we characterize values of α_π, α_y, and α_R that are associated with the highest level of welfare of the representative agent within the family of simple and implementable interest rate feedback rules defined by equation (13). As a point of comparison for policy evaluation, we also compute the real allocation associated with the Ramsey optimal policy.

The first row of Table 1 shows that under the Ramsey policy inflation is virtually equal to zero at all times.[8] The remaining rows of Table 1 report policy evaluations. The welfare associated with each interest rate feedback rule is compared to the level of welfare associated with the Ramsey-optimal policy. Specifically, the welfare cost is defined as the fraction, in percentage points, of the consumption stream an agent living in the Ramsey economy would be willing to give up to be as well off as in an economy in which

[8] In the deterministic steady state of the Ramsey economy, the inflation rate is zero. One may wonder why, in an economy featuring sticky prices as the single nominal friction, the volatility of inflation is not exactly equal to zero at all times under the Ramsey policy. The reason is that we do not follow the standard practice of subsidizing factor inputs to eliminate the distortion introduced by monopolistic competition in product markets. Introducing such a subsidy would result in a constant Ramsey-optimal rate of inflation equal to zero.

monetary policy takes the form of the respective interest rate feedback rule shown in the table.

We consider seven different monetary policies: Four constrained-optimal interest rate feedback rules and three nonoptimized rules. In the constrained-optimal rule labeled no-smoothing, we search over the policy coefficients, α_π and α_y, keeping α_R fixed at zero. The second constrained-optimal rule, labeled smoothing in the table, allows for interest rate inertia by setting optimally all three coefficients, α_π, α_y, and α_R.

We find that the best no-smoothing interest rate rule calls for an aggressive response to inflation and a mute response to output. The inflation coefficient of the optimized rule takes the largest value allowed in our search, namely 3. The optimized rule is quite effective as it delivers welfare levels remarkably close to those achieved under the Ramsey policy. At the same time, the rule induces a stable rate of inflation, a feature that also characterizes the Ramsey policy. Taking together this finding and those obtained in the previous subsection, we conclude that although a Taylor rule cannot exactly implement the Ramsey allocation, it delivers outcomes that are so close to the optimum in welfare terms that, for practical purposes, it can be regarded as implementing the Ramsey allocation.

We next study a case in which the central bank can smooth interest rates over time. Our numerical search yields that the optimal policy coefficients are $\alpha_\pi = 3$, $\alpha_y = 0.01$, and $\alpha_R = 0.84$. The fact that the optimized rule features substantial interest rate inertia means that the monetary authority reacts to inflation much more aggressively in the long run than in the short run. The finding that the interest rule is not superinertial (i.e., α_R does not exceed unity) means that the monetary authority is backward-looking. So, again, as in the case without smoothing, optimal policy calls for a large response to inflation deviations in order to stabilize the inflation rate and for no response to deviations of output from the steady state. The welfare gain of allowing for interest rate smoothing is insignificant. Taking the difference between the welfare costs associated with the optimized rules with and without interest rate smoothing reveals that agents would be willing to give up less than 0.001 percent of their consumption stream under the optimized rule with smoothing to be as well off as under the optimized policy without smoothing.

The finding that allowing for optimal smoothing yields only negligible welfare gains spurs us to investigate whether rules featuring suboptimal degrees of inertia or responsiveness to inflation can produce nonnegligible welfare losses at all. Panel A of Figure 3 shows that, provided the central bank does not respond to output, $\alpha_y = 0$, varying α_π and α_R between 0 and 3 typically leads to economically negligible welfare losses of less than 0.05 percent of consumption. In the graph, crosses represent combinations of α_π and α_R that are implementable and circles represent combinations that are implementable

Figure 3 The Cashless Economy

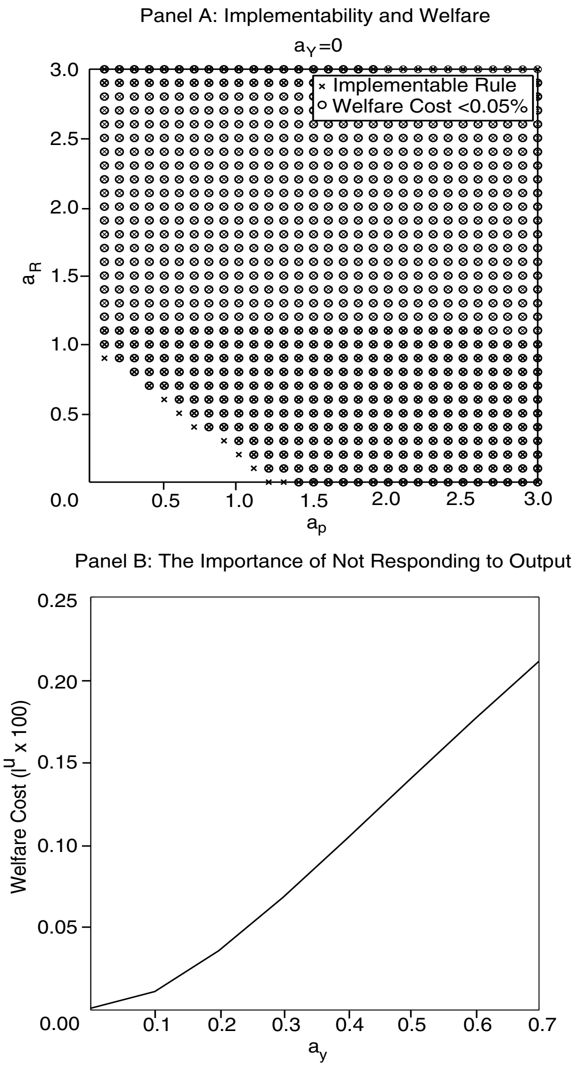

and that yield welfare costs less than 0.05 percent of consumption relative to the Ramsey policy.

The blank area in the figure identifies α_π and α_R combinations that are not implementable either because the equilibrium fails to be locally unique or because the implied volatility of interest rates is too high. This is the case for values of α_π and α_R such that the policy stance is passive in the long run, that is, $\frac{\alpha_\pi}{1-\alpha_R} < 1$. For these parameter combinations the equilibrium is not locally unique. This finding is a generalization of the result that, when the inflation coefficient is less than unity ($\alpha_\pi < 1$), the equilibrium is indeterminate, which obtains in the absence of interest rate smoothing ($\alpha_R = 0$). We also note that the result that passive interest rate rules render the equilibrium indeterminate is typically derived in the context of models that abstract from capital accumulation. It is, therefore, reassuring that this particular abstraction appears to be of no consequence for the finding that (long run) passive policy is inconsistent with local uniqueness of the rational expectations equilibrium. Similarly, we find that determinacy obtains for policies that are active in the long run, $\frac{\alpha_\pi}{1-\alpha_R} > 1$.

More importantly, Panel A of Figure 3 shows that virtually all parameterizations of the interest rate feedback rule that are implementable yield about the same level of welfare as the Ramsey equilibrium. This finding suggests a simple policy prescription, namely, that any policy parameter combination that is irresponsive to output and active in the long run, is equally desirable from a welfare point of view.

One possible reaction to the finding that implementability-preserving variations in α_π and α_R have little welfare consequences may be that in the class of models we consider, welfare is flat in a large neighborhood around the optimum parameter configuration, so that it does not really matter what the government does. This turns out not to be the case. Recall that in the welfare calculations underlying Panel A of Figure 3, the response coefficient on output, α_y, was kept constant and equal to zero. Indeed, interest rate rules that lean against the wind by raising the nominal interest rate when output is above trend can be associated with sizable welfare costs. Panel B of Figure 3 illustrates the consequences of introducing a cyclical component to the interest rate rule. It shows that the welfare costs of varying α_y can be large, thereby underlining the importance of not responding to output. The figure shows the welfare cost of deviating from the optimal output coefficient ($\alpha_y \approx 0$) while keeping the inflation coefficient of the interest rate rule at its optimal value ($\alpha_\pi = 3$) and not allowing for interest rate smoothing ($\alpha_R = 0$). Welfare costs are monotonically increasing in α_y. When $\alpha_y = 0.7$, the welfare cost is over 0.2 percent of the consumption stream associated with the Ramsey policy. This is a significant figure in the realm of policy evaluation at business-cycle frequency.[9] This finding suggests that bad policy can have significant welfare costs in our model and that policy mistakes are committed when policymakers are unable to resist the temptation to respond to output fluctuations.

[9] A similar result obtains if one allows for interest rate smoothing with α_R taking its optimized value of 0.84.

It follows that sound monetary policy calls for sticking to the basics of responding to inflation alone.[10] This point is conveyed with remarkable simplicity by comparing the welfare consequences of a simple interest rate rule that responds only to inflation with a coefficient of 1.5 to those of a standard Taylor rule that responds to inflation as well as output with coefficients 1.5 and 0.5, respectively. Table 1 shows that the Taylor rule that responds to output is significantly welfare inferior to the simple interest rate rule that responds solely to inflation. Specifically, the welfare cost of responding to output is about half a percentage point of consumption.[11]

The Ramsey-optimal monetary policy implies near complete inflation stabilization (see Table 1). It is reasonable to conjecture, therefore, that inflation targeting, interpreted to be any monetary policy capable of bringing about zero inflation at all times ($\pi_t = 1$ for all t), would induce business cycles virtually identical to those associated with the Ramsey policy. We confirm this conjecture by computing the welfare cost associated with inflation targeting. The welfare cost of targeting inflation relative to the Ramsey policy is virtually nil.

An important issue in monetary policy is determining what measures of inflation and aggregate activity the central bank should respond to. In particular, a question that has received considerable attention among academic economists and policymakers is whether the monetary authority should respond to past, current, or expected future values of output and inflation. Here we address this question by computing optimal backward- and forward-looking interest rate rules. That is, in equation (13) we let i take the values -1 and $+1$. Table 1 shows that there are no welfare gains from targeting expected future values of inflation and output as opposed to current or lagged values of these macroeconomic indicators. Also, a muted response to output continues to be optimal under backward- or forward-looking rules.

Under a forward-looking rule without smoothing ($\alpha_R = 0$), the rational expectations equilibrium is indeterminate for all values of the inflation and output coefficients in the interval [0,3]. This result is in line with that obtained by Carlstrom and Fuerst (2005). These authors consider an environment similar to ours and characterize determinacy of equilibrium for interest rate rules that depend only on the rate of inflation. Our results extend the findings of Carlstrom and Fuerst to the case in which output enters in the feedback rule.

We close this section by noting that most of the results presented here, extend to a model economy with a much richer battery of nominal and real rigidities. In Schmitt-Grohé and Uribe (2007b), we consider an economy featuring four real rigidities: habit formation, variable capacity utilization,

[10] Other authors have also argued that countercyclical interest rate policy may be undesirable (e.g., Ireland 1997 and Rotemberg and Woodford 1997).

[11] The simple interest rate rule that responds solely to inflation is implementable, whereas the standard Taylor rule is not, because it implies too high a volatility of nominal interest rates.

investment adjustment costs, and monopolistic competition in product and labor markets. The economy in that study also includes four nominal frictions, namely, sticky prices, sticky wages, money demand by households, and money demand by firms. Finally, the model features a more realistic shock structure that includes permanent stochastic variations in total factor productivity, permanent stochastic variations in the relative price of investment, and stationary stochastic variations in government spending. The values assigned to the structural parameters are based on existing econometric estimations of the model. These studies, in turn, argue that the model explains satisfactorily observed short-term fluctuations in the postwar United States. We find that the Ramsey policy calls for stabilizing price inflation. More importantly, a simple interest rate rule that responds only to inflation (with mute responses to wage inflation or output) attains a level of welfare remarkably close to that associated with the Ramsey optimal equilibrium.

6. CONCLUSION

In this article, we present a selective account of recent developments on the policy implications of the New Keynesian Phillips curve. The main lesson derived from our analysis is that price stability emerges as a robust policy prescription in models with product price rigidities. In fact, a minimum amount of price stickiness suffices to make inflation stabilization the overriding goal of monetary policy.

The desirability of price stability obtains in several variations of the standard New Keynesian framework that include expanding the set of nominal and real rigidities to allow for government spending financed by distortionary taxes, a transactional demand for money by households and firms, nominal wage rigidity, habit formation, variable capacity utilization, and investment adjustment costs.

A second important message that emerges is that a simple interest rate feedback rule that responds aggressively only to a measure of consumer price inflation delivers outcomes that are remarkably close to the Ramsey optimal equilibrium. In particular, to emulate optimal monetary policy it is not necessary that in setting the nominal rate the monetary authority respond to deviations of output from trend or past values of the interest rate itself. In this sense, the policy implications of the NKPC identified in this survey are consistent with a pure inflation targeting objective.

We have left out a number of important issues in the theory of inflation stabilization. For example, we limit attention to monetary policy under commitment. There is an active literature exploring the policy implications of the NKPC when the government is unable to commit to future actions. A central theme in this literature is to ascertain whether lack of commitment gives rise to an optimal inflation bias. A second omission in the present analysis concerns models with asymmetric costs of price adjustment. Here again, the central question is whether in the presence of downwardly rigid prices or wages, the policymaker should pursue a positive inflation target. Finally, our article does not discuss the recent literature on optimal monetary policy in models with credit constraints. An important focus of this literature is whether this type of friction introduces reasons for the central bank to respond to financial variables in setting the short-term interest rate.

REFERENCES

Aiyagari, S. Rao, Albert Marcet, Thomas J. Sargent, and Juha Seppälä. 2002. "Optimal Taxation Without State-Contingent Debt." *Journal of Political Economy* 110 (December): 1220–54.

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Lindé. 2005. "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle." Working Paper 11034. Cambridge, Mass.: National Bureau of Economic Research. (January).

Barro, Robert J. 1979. "On the Determination of Public Debt." *Journal of Political Economy* 87 (October): 940–71.

Bils, Mark, and Pete Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (October): 947–85.

Calvo, Guillermo. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (September): 383–98.

Calvo, Guillermo A., and Pablo E. Guidotti. 1990. "Indexation and Maturity of Government Bonds: An Exploratory Model." In *Capital Markets and Debt Management*, edited by R. Dornbusch and M. Draghi. New York: Cambridge University Press, 52–82.

Calvo, Guillermo A., and Pablo E. Guidotti. 1993. "On the Flexibility of Monetary Policy: The Case of the Optimal Inflation Tax." *Review of Economic Studies* 60 (July): 667–87.

Carlstrom, Charles T., and Timothy S. Fuerst. 2005. "Investment and Interest Rate Policy: A Discrete Time Analysis." *Journal of Economic Theory* 123 (July): 4–20.

Chari, Varadarajan V., Lawrence J. Christiano, and Patrick J. Kehoe. 1991. "Optimal Fiscal and Monetary Policy: Some Recent Results." *Journal of Money Credit and Banking* 23 (August): 519–39.

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

Correia, Isabel, Juan Pablo Nicolini, and Pedro Teles. 2008. "Optimal Fiscal and Monetary Policy: Equivalence Results." *Journal of Political Economy* 168: 141–70.

Del Negro, Marco, Frank Schorfheide, Frank Smets, and Raf Wouters. 2004. "On the Fit and Forecasting Performance of New-Keynesian Models." Manuscript (November).

Feenstra, Robert C. 1986. "Functional Equivalence Between Liquidity Costs and the Utility of Money." *Journal of Monetary Economics* 17 (March): 271–91.

Goodfriend, Marvin, and Robert G. King. 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio J. Rotemberg. Cambridge, Mass.: MIT Press, 231–83.

Ireland, Peter. 1997. "A Small, Structural, Quarterly Model for Monetary Policy Evaluation." *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 83–108.

Khan, Aubhik, Robert G. King, and Alexander Wolman. 2003. "Optimal Monetary Policy." *Review of Economic Studies* 70 (October): 825–60.

Kimbrough, Kent P. 1986. "The Optimum Quantity of Money Rule in the Theory of Public Finance." *Journal of Monetary Economics* 18 (November): 277–84.

Klenow, Pete, and Oleksiy Kryvtsov. 2005. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" Mimeo, Stanford University (July).

Levin, Andrew, Alexei Onatski, John Williams, and Noah Williams. 2006. "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models." In *NBER Macroeconomics Annual 2005*, vol. 20, edited by Mark Gertler and Kenneth Rogoff. Cambridge, Mass.: MIT Press.

Lucas, Robert E., Jr. 1982. "Interest Rates and Currency Prices in a Two-Country World." *Journal of Monetary Economics* 10: 335–60.

Lucas, Robert E., Jr., and Nancy L. Stokey. 1983. "Optimal Fiscal and Monetary Policy in an Economy Without Capital." *Journal Monetary Economics* 12: 55–93.

Nakamura, Emi, and Jón Steinsson. 2007. "Five Facts About Prices: A Reevaluation of Menu Cost Models." Mimeo, Harvard University (May).

Nason, James M., and Gregory W. Smith. 2008. "The New Keynesian Phillips Curve: Lessons from Single-Equation Econometric Estimation." Federal Reserve Bank of Richmond *Economic Quarterly* 94 (Fall): 361–96.

Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90 (December): 1187–211.

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

Schmitt-Grohé, Stephanie, and Martín Uribe. 2004a. "Optimal Fiscal and Monetary Policy under Sticky Prices." *Journal of Economic Theory* 114 (February): 198–230.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2004b. "Optimal Fiscal and Monetary Policy under Imperfect Competition." *Journal of Macroeconomics* 26 (June): 183–209.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2006. "Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model." In *NBER Macroeconomics Annual 2005*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, Mass. and London: MIT Press, 383–425.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2007a. "Optimal Simple and Implementable Monetary and Fiscal Rules." *Journal of Monetary Economics* 54 (September): 1702–25.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2007b. "Optimal Inflation Stabilization in a Medium-Scale Macroeconomic Model." In *Monetary Policy Under Inflation Targeting*, edited by Klaus Schmidt-Hebbel and Rick Mishkin. Santiago, Chile: Central Bank of Chile, 125–86.

Schorfheide, Frank. 2008. "DSGE Model-Based Estimation of the New Keynesian Phillips Curve." Federal Reserve Bank of Richmond *Economic Quarterly* 94 (Fall): 397–434.

Sidrauski, Miguel. 1967. "Rational Choice and Patterns of Growth in a Monetary Economy." *American Economic Review Papers and Proceedings* 57: 534–44.

Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (June): 586–606.

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

Woodford, Michael. 1996. "Control of the Public Debt: A Requirement for Price Stability?" Working Paper 5684. Cambridge, Mass.: National Bureau of Economic Research. (July).

Woodford, Michael. 2001. "The Taylor Rule and Optimal Monetary Policy." *American Economic Review Papers and Proceedings* 91 (May): 232–7.

Woodford, Michael. 2003. *Interest and Prices*. Princeton, N.J.: Princeton University Press.

Yun, Tack. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37 (July): 345–70.