Research Technical Paper

Modelling Inflation Dynamics:
A Critical Review of Recent Research

Jeremy Rudd   ### Karl Whelan
Federal Reserve Board   **CBFSAI**[*]

Abstract

In recent years, a broad academic consensus has arisen around the use of rational expectations sticky-price models to capture inflation dynamics. These models are seen as providing an empirically reasonable characterization of observed inflation behavior once suitable measures of the output gap are chosen; and, moreover, are perceived to be robust to the Lucas critique in a way that earlier econometric models of inflation are not. We review the principal conclusions of this literature concerning: 1) the ability of these models to fit the data; 2) the importance of rational forward-looking expectations in price setting; and 3) the appropriate measure of inflationary pressures. We argue that existing rational expectations sticky-price models fail to provide a useful empirical description of the inflation process, especially relative to traditional econometric Phillips curves of the sort commonly employed for policy analysis and forecasting.

1 Introduction

Robert Solow (1976) once observed that "any time seems to be the right time for reflections on the Phillips curve." However, right now seems to present a particularly appropriate moment to take stock of the empirical evidence on inflation dynamics. Recent years have seen an explosion in empirical research on inflation, with most of it related to the so-called "new-Keynesian" Phillips curve, which has provided a modern take on the traditional Phillips curve relationship by deriving it from an optimizing framework featuring rational expectations and nominal rigidities. That the current conference has taken its inspiration from the 1970 Federal Reserve conference on "The Econometrics of Price Determination" also seems appropriate because, like now, the 1970s witnessed an intense debate over the theoretical and empirical underpinnings of a popular econometric model of inflation. And, like now, these debates largely revolved around the merits of what appeared to be a new paradigm for understanding the behavior of inflation and the macroeconomy.[1]

In this paper, we offer a selective and critical review of recent developments in the theoretical and empirical modelling of U.S. inflation dynamics. We use the term "selective" because we are not attempting to provide a comprehensive summary of the huge amount of research devoted to this topic in recent years. Rather, we hope to shed light on a couple of key issues: first, how are inflation expectations formed; and second, what is an appropriate empirical measure of inflationary pressures.

We use the term "critical" because we have had some involvement in the recent debates that we review and our survey will reflect the answers to these questions that we have proposed in our earlier work.[2] In particular, our research has suggested a number of reasons to be skeptical about the new-Keynesian framework that is bidding to become the new benchmark model for inflation analysis. More generally, we discuss some reasons to doubt some of the stronger implications of the rational expectations hypothesis for the modelling of inflation. In that sense, our work connects back to many of the themes of the 1970 conference, and it is with those earlier debates that we begin.

We start by reviewing the origins of the Phillips curve and the debates over the accelerationist version of the model introduced by Friedman (1968) and Phelps (1967). We

---

[1]As Solow (1968, p. 3) also noted, ". . . the theory of inflation seems to make progress by way of a series of controversies. It is not uncommon for economics, or even for natural science, to proceed in this adversary manner, but I rather think it is especially characteristic of the analysis of inflation."

[2]See Rudd and Whelan (2003), (2005a), and (2005b).

discuss the important critiques of econometric Phillips curves made by Sargent (1971) and Lucas (1972a), and trace how this led to the development of the modern "new-Keynesian" Phillips curve. We outline how, despite their apparent similarity, the accelerationist and new-Keynesian models turn out to have very different implications for monetary policy and for econometric modelling and forecasting.

The paper next provides an empirical assessment of the new-Keynesian Phillips curve. This is a structural model, designed to be capable of explaining the behavior of inflation without being subject to the Lucas critique, but it is well known that it generates extremely counterfactual predictions when traditional output gaps (based on naïve detrending procedures) are used as a measure of inflationary pressures. However, in recent years it has become widely accepted that an alternative approach, which substitutes labor's share of income in place of detrended output, is theoretically superior and yields a good empirical model of inflation dynamics. We argue that the theoretical case for this approach—which was advocated in an influential paper by Galí and Gertler (1999)—is quite weak, and that the labor's share version of the new-Keynesian model actually provides a very poor description of observed inflation behavior. This failure of the model extends along two dimensions: first, labor's share fails to provide a good measure of inflationary pressures; and second, this version of the model cannot account for the important role played by lagged inflation in empirical inflation regressions.

We also review the evidence relating to the so-called "hybrid" class of new-Keynesian models, which add a dependence of inflation on its own lagged values to otherwise purely forward-looking models. These models are often viewed as striking a compromise between the need for rigorous microfoundations of the sort underlying the pure new-Keynesian model and the need for reasonable empirical fit; thus, they have commonly been adopted for use in applied monetary policy analysis. Galí and Gertler's conclusion that rational forward-looking behavior plays the dominant role in these models is widely cited as a stylized fact in this literature. We provide an alternative interpretation of the empirical estimates obtained from these models, and argue that the data actually provide very little evidence of an important role for rational forward-looking behavior of the sort implied by these models.

Finally, we end with a brief discussion of the properties of reduced-form econometric Phillips curves. The importance of lagged inflation in these models has led to their being criticized as being especially susceptible to the Lucas critique. As we show, however, this potential shortcoming of these models seems to be relatively unimportant in practice.

2 Econometric Inflation Equations: A Potted History

The correlation between nominal wage inflation and unemployment in United Kingdom data that was documented by A.W. Phillips (1958) quickly crossed the Atlantic and was adopted by Samuelson and Solow (1960) as a relationship between price inflation and unemployment that was relevant for the United States. Samuelson and Solow were aware that the empirical regularity documented in their paper was potentially fragile, and offered a number of clear caveats regarding its usefulness as a policy tool. Nevertheless, Keynesian policy discussions of the 1960s leaned heavily on the idea of the Phillips curve as presenting an exploitable tradeoff between the level of inflation and the level of unemployment.

Despite their popularity, it is clear in retrospect that the theoretical foundations of these early formulations were not completely sound, with a particular weak point being their treatment of how expectations would enter wage and price setting. This weakness was thoroughly critiqued in the seminal contributions of Phelps (1967) and Friedman (1968), with the latter paper, in particular, having a profound influence.

2.1 The Accelerationist Phillips Curve

Friedman started from the assumption that workers and firms could distinguish between real and nominal wages, and that it was real wages, or more to the point *expected* real wages, that were central to the wage bargain. By reducing the *ex post* real wage, a *surprise* increase in the price level could result in higher employment. But a fully anticipated increase in prices would be fully reflected in nominal wage demands, and would therefore have no effect on the real wage rate. As a result, a policymaker who aimed to engineer a permanent reduction in unemployment through a permanent one-time increase in the inflation rate would ultimately be disappointed: Once people came to expect it, the faster rate of inflation would provide the basis for bargaining over nominal wage increases and therefore stop having any real effects. When this happened, unemployment would return to a "natural" rate that was determined by the structure of labor and product markets.[3]

Hence, the correct formulation of the inflation-unemployment tradeoff was not a relation of the form $\pi_t = \gamma U_t$, but rather an "expectations-augmented" Phillips curve of the form:

$$\pi_t = \gamma(U_t - U^*) + \pi_t^e, \tag{1}$$

_____

[3]Both of these ideas—that changing expectations would shift the Phillips relation, and that eventually the economy would return to an equilibrium unemployment rate once actual and expected inflation were equal—are also present in Phelps (1967).

where inflation, π_t, is (negatively) correlated with deviations of the unemployment rate from its natural rate U^*, and where the entire curve is shifted up or down one-for-one with changes in π_t^e (the rate of inflation that agents had expected to prevail in time t).

In the Friedman-Phelps framework, then, there is no permanent tradeoff between the *level* of inflation and the unemployment rate—in the long run, $\pi_t^e = \pi_t$ and the Phillips curve is vertical at $U_t = U^*$. However, to the extent that agents' expectations were slow to catch up to reality, a policymaker could keep unemployment below the natural rate by constantly boosting the inflation rate. For this reason, the Friedman-Phelps characterization of the inflation process also came to be known as the "accelerationist hypothesis," since an acceleration in prices (an increase in the rate of inflation) would occur should policymakers attempt to permanently keep unemployment below its natural rate.

2.2   Tests of the Accelerationist Model and the Sargent-Lucas Critique

The Friedman-Phelps description of the Phillips curve was predicated on the assumption that expectations about inflation evolved over time as a result of actual past experience—that is, that expectations were formed *adaptively*.[4] In empirical evaluations of the accelerationist hypothesis, therefore, researchers used a distributed lag of past inflation rates to proxy for π_t^e, and then tested whether such proxies received a coefficient of unity. Indeed, for much of the profession, a test of whether ρ equalled one in regressions of the form

$$\pi_t = \alpha + \gamma U_t + \rho \sum_{i=1}^{N} \beta_i \pi_{t-i} + \epsilon_t \tag{2}$$

(where the weights β_i were constrained to sum to unity), was viewed interchangeably as both a test of the accelerationist hypothesis and of the existence of a natural rate of unemployment. A number of the contributions to the 1970 Federal Reserve conference on price determination focused on testing this hypothesis (and, when it was rejected, on quantifying the implied long-run tradeoff between inflation and unemployment), and, at first, many Keynesians concluded that the accelerationist view was not supported by the data.[5]

This empirical strategy was criticized in two remarkable papers by Sargent (1971) and Lucas (1972a) that can fairly be said to have ushered in the rational expectations revolution

---

[4]In Phelps (1967), the appeal to adaptive expectations is explicit. The term is not used by Friedman (1968), who provides an informal discussion of a gradual adjustment process.

[5]See Tobin (1972) for a summary of the 1970 conference, and Solow (1968, 1969) for early contributions that considered and rejected the accelerationist model.

in macroeconomics. For these economists, the treatment of expectations implicit in estimates of equations like (2) was deficient in that it was inconsistent with rational behavior. Sargent argued that (at the time) the U.S. inflation process appeared to be mean-stationary, so it was quite reasonable for the public to formulate inflation expectations in a manner that was consistent with mean-reversion. Thus, forcing the distributed lags on past inflation to sum to one—as was done in conventional tests of the accelerationist model—was inconsistent with the forecasts that a rational agent would make, and would lead to empirical estimates of ρ that were less than one even if the accelerationist hypothesis were correct.

Almost simultaneously, Lucas's important contribution to the 1970 Federal Reserve conference provided an explicit example of a rational expectations model in which the natural rate hypothesis (that the only relationship between activity and inflation involved *unanticipated* inflation) held, but under which standard econometric tests would point to rejection. In Lucas's setup, the central bank pursues a monetary policy in which money growth (and hence inflation) is mean stationary, so reduced-form regressions for inflation would yield values of ρ less than one even though agents had rational expectations and the economy *by construction* did not manifest a long-run tradeoff between output and inflation.

Although Lucas's more comprehensive critique of contemporary econometric models did not appear for several years (Lucas, 1976), the essence of this critique as it applies to models of inflation was clearly spelled out in this earlier paper. An economy characterized by rational expectations and a stable inflation objective could produce a reduced-form process for inflation with a value of ρ less than one. However, any suggestion that this implied a long-run tradeoff between inflation and output was simply incorrect: A change in policy that attempted to exploit any perceived tradeoff would result in higher average inflation and a different reduced-form representation of the inflation-output relationship.

2.3 Towards a Modern Econometric Phillips Curve

In the decade or so that followed the publication of the Phelps and Friedman papers, the notion that the accelerationist view of the inflation process was correct gained wider acceptance, even among those who did not fully subscribe to all elements of the rational expectations framework proposed by Sargent and Lucas.[6] Several factors contributed to this change in attitudes. The first, of course, was the strength of the theoretical arguments themselves. Second, it became apparent by the mid-1970s that the inflation-unemployment

---

[6]Indeed, by 1975 Arthur Okun could claim that "we are all accelerationists now" (Okun, 1975, p. 356).

tradeoff implied by the short-run Phillips curve had shifted. Finally, it became easier to find that the lags of inflation in empirical Phillips curves summed to one.[7] In addition, the important contribution of "supply shocks" to price acceleration in the early 1970s led to food, energy, and/or import prices receiving special treatment in empirical descriptions of inflation. What emerged in this period, therefore, was a benchmark econometric model of inflation of the form:

$$\pi_t = \alpha + \gamma U_t + B(L)\pi_{t-1} + \lambda z_t + \epsilon_t, \tag{3}$$

where $B(1)$ is constrained to equal one and z_t denotes a vector of supply shocks.[8] In this specification, then, inflation dynamics are determined by three sources: real activity (as summarized by the unemployment rate), supply shocks, and "inertia" (as captured by the lagged inflation terms); for this reason, it is sometimes called the "triangle model."

The basic framework of the triangle model has been widely adopted by applied macroeconomists, and has brought with it the important concept of the NAIRU, defined as the unemployment rate consistent with unchanging inflation in the absence of supply shocks. In the context of the simple version of the triangle model presented above (equation 3), the NAIRU is calculated from the estimated coefficients as $U^* = \frac{-\alpha}{\gamma}$. More refined analyses of the NAIRU include Staiger, Stock, and Watson's (1997a) improved approach for quantifying the precision with which the NAIRU can be measured, attempts by Staiger, Stock, and Watson (1997b), Gordon (1998), and Ball and Mankiw (2002) to model variations over time in the NAIRU, and work by Braun (1984) and Ball and Moffitt (2002) that seeks to assess the effect of changes in trend productivity on the NAIRU.

Taken literally, the characterization of inflation dynamics that the triangle model provides carries important implications for the conduct of macroeconomic policy. To the extent that lagged inflation captures true inertia in the price-setting process (resulting, for instance, from how expectations are formulated), the model implies that rapid reductions in inflation can only be produced at the cost of a substantial increase in unemployment. Hence, the model points to a gradualist approach as providing the best way to effect a large

---

[7]See Gordon (1976) for a contemporaneous account, and McCallum (1994) for a retrospective one. In addition, McCallum's (1976) more direct test of expectations in wage setting provided further empirical support for the accelerationist hypothesis.

[8]See Gordon (1998) for a typical implementation. We are, of course, presenting a simplified history of the development of this model. For example, many early papers featured wage-price Phillips curves combined with markup equations; over time, the system was collapsed into a single equation. See Gordon (1988) for a discussion.

reduction in inflation. In addition, policymakers must be mindful of the presence of long time lags between when macroeconomic shocks (including policy actions) are initially felt, and when they have their full effects on inflation. Thus, this framework provides a strong argument in support of preemptive action to cut off the full effect of an inflationary shock.

3   Rational Expectations and the New-Keynesian Phillips Curve

Lucas and Sargent's introduction of rational expectations into the modelling of economic dynamics had a significant influence on the profession and largely dictated the development of macroeconomic theory from the mid-1970s onwards. However, the alternative Phillips curve model advanced by Lucas (1972b)—the so-called "islands" model in which a short-run inflation-output tradeoff resulted from imperfect information on the part of price-setters—generally failed to convince as a model of economic fluctuations. In part, this initial resistance to the rational expectations approach likely reflected the Lucas model's implication that only purely random and transitory policy shocks could affect output.

The challenge, then, was to demonstrate that persistent effects of nominal disturbances—such as a shock to the money growth rate—could obtain in a rational expectations framework. This challenge was met by Stanley Fischer (1977) and John Taylor (1979), who constructed alternative models in which nominal rigidities interacted with rational expectations in such a way as to deliver conclusions about the effects of policy that were much more Keynesian in flavor. This in turn set the stage for the emergence of "new-Keynesian" macroeconomics, which accepted the idea of rational expectations, but which focused on the role played by various rigidities and market imperfections in yielding Keynesian-style prescriptions for fiscal and monetary policy.

Importantly, though, despite the significant influence of the rational expectations school of thought, and despite a general acceptance of the Lucas-Sargent critique of econometric Phillips curves on theoretical grounds, accelerationist econometric models of inflation have remained popular and highly influential in policy circles. For example, the concept of the NAIRU has featured prominently in many editions of the *Economic Report of the President* from the early 1980s to the Bush administration.[9] Also, the notion that inflation is sluggish or inertial is often invoked as a key consideration in the efficient design of monetary policy.[10]

---

[9]Romer and Romer (2002) provide specific examples.

[10]See the first chapter of Blinder (1998), for instance.

This division between theoretical Keynesian-style models, on the one hand, and empirical econometric modelling, on the other, has become more controversial of late. This is because many academic economists have become convinced that certain theoretical new-Keynesian models can provide a good description of the empirical inflation process. In part, this development stemmed from the realization that a number of popular new-Keynesian models of price-setting each implied a sort of Phillips curve relationship, a pattern that was noted in an important paper by John Roberts (1995). We now turn to a detailed description of this "new-Keynesian Phillips curve."

3.1   The New-Keynesian Phillips Curve: Derivation

By far the most popular formulation of the new-Keynesian Phillips curve is based on Gulliermo Calvo's (1983) model of random price adjustment. This model assumes that in each period a random fraction $(1 - \theta)$ of firms reset their price, while all other firms keep their prices unchanged. Thus, the evolution of the (log) price level is given by

$$p_t = \theta p_{t-1} + (1 - \theta) z_t, \tag{4}$$

where z_t is the price chosen by those who can reset their prices. Assuming an imperfectly competitive market structure such that, absent any frictions, firms would set their price as a fixed markup over marginal cost, a firm's optimal reset price is determined by

$$z_t = \mu + (1 - \theta\beta) \sum_{k=0}^{\infty} (\theta\beta)^k E_t \widetilde{mc}_{t+k}, \tag{5}$$

where μ is the frictionless optimal markup, β is the firm's discount factor, and $\widetilde{mc}_t$ is nominal marginal cost. In other words, firms take into account that their prices will likely be fixed over some period by setting their price equal to a weighted average of expected future nominal marginal costs. These two equations can be combined to yield a new-Keynesian Phillips curve (NKPC) of the form:

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1 - \theta)(1 - \theta\beta)}{\theta} (mc_t + \mu), \tag{6}$$

which relates inflation, π_t, to next period's expected inflation rate and the current deviation of real marginal cost ($mc_t = \widetilde{mc}_t - p_t$) from its frictionless optimal level. Under relatively general conditions, aggregate real marginal cost is proportional to the gap between output and its potential level. With this assumption, the NKPC becomes

$$\pi_t = \beta E_t \pi_{t+1} + \gamma y_t, \tag{7}$$

where y_t is the output gap.

This model of inflation has several appealing features, not least of which is that it provides a microfounded formulation of the inflation-output tradeoff that is consistent with rational expectations. The appeal of the model has gone beyond what was normally associated with the earlier theoretical demonstrations of the new-Keynesian literature. For instance, with many studies showing that traditional RBC models fail to explain the business cycle, and with solid empirical evidence now available to support the conjecture that prices are relatively sticky, Calvo-style nominal rigidities have been added to many dynamic stochastic general equilibrium models.[11]

3.2 Implications of the NKPC for Policy and Practice

Even for those economists who favor traditional econometric Phillips curves, the NKPC appears on the surface to provide a plausible theoretical model of inflation. While standard econometric models include lagged inflation, these lags are often understood to be proxying for expectations, so the NKPC bears some resemblance to the theoretical model that these analysts have in mind. However, while it is superficially similar to the accelerationist Phillips curve, the NKPC carries implications for policy and for macroeconomics more generally that are dramatically different from the conventional wisdom derived from the older class of econometric models.

No Inflation Inertia: Probably the most important implication of the NKPC model is that there is no "intrinsic" inertia in inflation, in the sense that there is no structural dependence of inflation on its own lagged values. Instead, inflation is determined in a completely forward-looking manner, as can be seen by solving forward equation (7) under the assumption of rational expectations to obtain the following closed-form solution for π_t:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k}. \tag{8}$$

This equation makes it clear that much of the conventional wisdom concerning the conduct of monetary policy is potentially incorrect. The idea that there is considerable inertia in inflation, and hence that it is difficult to reduce inflation quickly, does not hold in this framework—indeed, according to the NKPC, inflation can be costlessly controlled by a

---

[11]See Bils and Klenow (2004) for evidence on price rigidities for items in the U.S. consumer price index.

credible commitment to keep output close to its potential level now and in the future.[12]

Lucas-Sargent Revisited: For advocates of the NKPC, lagged inflation terms enter econometric Phillips curves merely because they are proxying for expectations of future values of the output gap, which are what truly determine current inflation (see equation 8). Thus, the supposed inertial nature of inflation is in reality only a sort of statistical mirage, and so should not have any bearing on the conduct of policy. More problematically, the strength of this statistical correlation is likely to vary across monetary policy regimes: In periods when the central bank has little credibility, the public may formulate its inflation expectations based on actual recent inflation performance, rather than on the public statements of the central bank. By contrast, if the central bank maintains a credible inflation target, then recent lagged values of inflation may play only a small role in the formulation of expectations.

For these reasons, as the NKPC has increased in popularity, the Lucas-Sargent critique of econometric Phillips curves has returned to center stage in discussions of empirical modelling of inflation. For example, the wisdom of persisting with an accelerationist Phillips curve specification (in which the coefficients on lagged inflation are constrained to sum to one) in the face of a credible and stable monetary policy regime has been questioned. Conversely, Taylor (1998) and Cogley and Sargent (2002) have warned about the possible danger of deviating from the accelerationist assumption and using inflation models whose own lags sum to less than one—even if these are justified empirically. In particular, if econometric estimates of the sum of the coefficients on lagged inflation fall below one, this may falsely signal the return of an exploitable long-run tradeoff between inflation and output.

Misleading Nature of NAIRU-Based Analyses: Another implication of the new-Keynesian Phillips curve is that it is probably inappropriate to use NAIRU estimates to guide macroeconomic policy. For instance, recent estimates of the time-varying NAIRU that are obtained by inverting an accelerationist inflation equation suggest that the surprisingly low inflation of recent years can be explained in terms of a reduction in the NAIRU. Implicitly, then, this indicates that the Federal Reserve could run the economy at a higher level of activity without pushing up inflation. If, however, recent low rates of inflation have

---

[12]In addition, as Ball (1994) showed, in certain new-Keynesian models it is possible to reduce inflation while generating an *increase* in output.

instead resulted from a centering of inflation around a credible low target level (in other words, if actual and expected inflation have become more mean-reverting in nature), then the measured reductions in the NAIRU are actually an artefact of preserving the accelerationist restriction in the face of a reduction in the importance of lagged inflation in the inflation process.

In fact, this suggests a more realistic example of how reliance on traditional econometric analysis could lead to incorrect policy advice than the Taylor-Cogley-Sargent concern that central bank economists may again start to believe in the existence of a long-run tradeoff between inflation and real activity. As far as we can tell, the natural rate hypothesis is still firmly established in the minds of central bank economists as a fundamental characteristic of the economy. Importantly, this widespread rejection of the presence of any long-run tradeoff between inflation and activity is based on more than statistical correlations; at least as important is the theoretical plausibility of the natural rate hypothesis. Against this background, it seems more likely (at least over the next several years) that policy errors could result from the inappropriate use of empirical models that impose the accelerationist property, rather than from the advocacy of a long-run inflation-unemployment tradeoff.

4   Problems with the New-Keynesian Phillips Curve

The implications of the new-Keynesian Phillips curve for policy and forecasting represent a dramatic departure from conventional wisdom. As such, we would hope that practitioners would be able to find persuasive evidence in favor of the NKPC before adopting it as a framework for understanding and predicting inflation. Unfortunately, as we document in this section and the next, the empirical case against the NKPC turns out to be quite strong.

We first consider the version of the model that is directly suggested by equation (7), which relates current inflation to expected inflation and a measure of the output gap. The first row of Table 1 reports results from GMM estimation of this equation over the sample 1960:Q1-2004:Q3, where inflation is defined as the log-difference of the price index for nonfarm business sector output, and where the output gap is constructed by detrending the log of real GDP with a quadratic polynomial in calendar time.[13] As is well established,

---

[13]In general, we prefer using the H-P filter to construct our measure of the output gap because it represents a relatively uncontroversial way to detrend GDP. Technically, however, it is not valid to use an H-P filtered series as an instrument for GMM estimation because the filter employs future information in computing the estimated trend. For all of our GMM models, therefore, we follow Galí and Gertler (1999) and measure

the results from this type of exercise provide bad news for this formulation of the model. The coefficient on the output gap is highly statistically significant but *negative*—a finding that is robust across a wide range of possible instrument sets, measures of inflation, and GDP detrending procedures.[14]

One way to understand the failure of the output gap version of the NKPC is in terms of the model's closed-form solution (equation 8), which states that inflation should be determined by a discounted sum of expected future output gaps. Thus, the model predicts that higher inflation should lead increases in output relative to trend. In fact, however, there is little evidence of such a pattern in the data: As can be seen from panel A of Table 2, in regressions of detrended GDP on its own lags and lags of inflation we find that higher inflation is a leading indicator for future *reductions* in the output gap.

These poor results suggest two possible interpretations: Either the rational-expectations NKPC provides a bad description of inflation, or else this particular measure of the output gap is flawed. Perhaps unsurprisingly, the latter explanation has proven popular in recent years with proponents of the model. Typically, these researchers criticize traditional measures of the output gap on the grounds that naïve detrending procedures assume that potential GDP evolves smoothly over time. In theory, however, changes in potential output will be affected by any number of shocks, and so could fluctuate significantly (and stochastically) period to period. Moreover, even if potential GDP *could* be characterized by a relatively smooth trend, there may be little agreement over how, precisely, to estimate this trend component.

The recent literature on the NKPC has attempted to address this problem by replacing the GDP gap with a different measure of inflationary pressures. Before turning to a discussion of this alternative approach, we note a fact about traditional measures of the output gap that has tended to get lost in much of the recent literature. While the points just noted regarding the difficulties surrounding the measurement of potential output are clearly valid

the GDP gap as quadratically detrended log real GDP. We also re-did all of our GMM analyses using both H-P detrended GDP as well as an alternative (cubic) detrending procedure, and obtained results that were very similar to those that we report.

[14]The instrument set for this regression includes inflation, the labor income share, the output gap, an interest rate spread (defined as the difference between ten-year and three-month Treasury yields), hourly compensation growth, and commodity price inflation (measured by the PPI for crude materials). This is the same list of instruments employed by Galí and Gertler (1999); however, we use two lags of each variable instead of their four because the larger set fails Stock and Yogo's (2003) test for weak instruments.

and important, neither these points nor the failure of the new-Keynesian Phillips curve when a traditional output gap measure is used should, on their own, be taken as reasons to discount the usefulness of standard gap measures as indicators of inflationary pressures.

Indeed, despite the obvious difficulties that attend the measurement of potential GDP, traditional output gaps actually perform very well in conventional econometric inflation models: In a regression of nonfarm business price inflation on four of its own lags and one lag of H-P filtered real GDP, the lagged output gap receives a t-statistic of 5.1, which is significant at any conventional level.[15] And, of course, the fact that the output gap leads inflation positively while inflation leads the output gap negatively is perfectly consistent with a model of the economy consisting of an accelerationist Phillips curve coupled with a central-bank reaction function (such as a Taylor rule).

5   The New-Keynesian Phillips Curve with Labor's Share

As is evident from equation (6), the theoretical models that underpin the new-Keynesian Phillips curve predict that it is real marginal cost that drives inflation. Hence, recent attempts to salvage the NKPC as a workable model have sought to construct a convincing empirical proxy for mc_t. In particular, Sbordone (2002) and Galí and Gertler (1999) have proposed using average unit labor costs to measure nominal marginal cost, with the former concept defined as $\frac{wL}{Y}$ (where w denotes the nominal compensation rate, L equals total labor input, and Y equals real output). The resulting proxy for *real* marginal cost, $\frac{wL}{pY}$, is therefore labor's share of income.[16] In this section, we first discuss the theoretical case for using labor's share as a proxy for real marginal cost, and then assess the empirical evidence in favor of the version of the NKPC that employs this measure of inflationary pressures.

5.1   Theoretical Issues

Before focusing on labor's share itself, it is worth briefly reviewing the theoretical case for the relationship underlying equation (7), which allows the NKPC to be written as a function of the gap between output and its potential level. As Woodford (2003) has pointed out, increases in output that are not driven by increases in technological efficiency will tend to raise nominal marginal costs more than prices in a broad class of standard models, as

---

[15]Similar results obtain for measures of deterministically detrended GDP.

[16]Despite the fact that the labor income share is a well-established concept in macroeconomics, most papers in the recent literature on the NKPC have instead referred to this variable as real unit labor costs.

workers require a higher real wage in order to be induced to supply extra hours.[17] (In practice, labor-market institutions such as overtime premia and factors like upward-sloping supply curves for firm-specific inputs considerably strengthen the case for the procyclicality of real marginal cost.)

In light of the theoretical case for procyclicality, an examination of Figure 1 reveals a potential difficulty with using labor's share as a proxy for real marginal cost. The figure shows time-series plots of detrended GDP and the labor share, with the shaded bars displaying official NBER recession dates. Rather than moving procyclically, the labor share (lower panel) has typically displayed a pattern that would be considered *countercyclical*, with the series spiking upward during each postwar recession in the United States.

This negative correlation between the labor share and traditional measures of the output gap has been noted by those who advocate using labor's share to measure mc_t, and has even been cited as evidence that traditional GDP gaps might be negatively correlated with the "true" gap between output and its potential level.[18] In our opinion, however, this is not a satisfactory interpretation of the differential behavior of the two series shown in Figure 1. The measurement of potential output is clearly a difficult problem and simple trend-fitting procedures no doubt yield relatively crude approximations. That said, the upper panel of Figure 1—which plots the H-P filtered measure of the GDP gap—suggests to us that these simple methods still do a reasonable job of detecting periods when output deviates from potential.

Consider, for instance, the periods identified as recessions by the NBER. Despite the various conceptual difficulties surrounding the notion of potential output, it is generally agreed that these were times in which output was below potential, with this assessment being based on the behavior of numerous economic indicators, rather than on the deviation of output from a prevailing trend. Nevertheless, the business cycle implied by the movements in this measure of the output gap is consistent with the NBER concept, in that the gap manifests substantial peak-to-trough declines during each of these recessionary periods.[19]

In contrast, the labor share tends to jump upward and reach a local peak near the onset of the NBER recessions. For the labor share to be a good proxy for real marginal cost, and for real marginal cost to be positively correlated with the gap between output and

---

[17]Woodford (2003, p. 180) provides a formal derivation.

[18]See Woodford (2003, p. 206) for a version of this argument.

[19]The same is true for the quadratically or cubically detrended GDP gap measures (not shown).

its potential level, we would have to conclude from the behavior of the labor share that output was actually *above* potential during each postwar recession. While this pattern is at least theoretically possible, and might find some support among those who view negative technology shocks as the principal cause of recessions, we believe that the evidence to the contrary—as outlined, for instance, by Galí (1999)—is quite compelling: Technology shocks do not appear to be the dominant force behind business cycle fluctuations.[20]

How can the observed behavior of the labor share be reconciled with the theoretical prediction that real marginal cost should be procyclical? Most likely, the answer is that average unit labor costs are simply a poor proxy for nominal marginal cost. Indeed, as Rotemberg and Woodford (1999) discuss in detail, there are several reasons to believe that marginal and average cost manifest different cyclical patterns. Some of the phenomena that would cause real marginal cost to be procyclical, such as overtime premia or adjustment costs, relate to patterns that can have substantial effects at the margin, but that could have little impact on average costs. Conversely, average cost can be affected by a range of factors that will have no effect on marginal cost. For instance, cyclical utilization rates and the existence of overhead labor can cause unit labor costs (and the labor share) to increase substantially during recessions, even as marginal cost may be declining. All told, it seems likely that real marginal cost is procyclical and is related to the gap between output and potential, but that the labor share proxy is simply not capturing these cyclical movements correctly.

By and large, the important measurement issues that surround the use of average unit labor costs as a proxy for marginal cost have been ignored in the recent literature on the NKPC. Indeed, proponents of the labor share approach such as Galí, Gertler, and Lopez-Salido (2005) generally refer to the labor share series simply as "real marginal cost," as though these two concepts are identical. In some cases, advocates of this approach have conceded that labor's share provides only a crude proxy for mc_t, but have then suggested that the performance of the NKPC model (which they argue is already good when labor's share is used) can only be improved with a more sophisticated measure of real marginal

---

[20]While potential output is often viewed as depending only on the level of technology, Woodford (2003) defines it as that level of output consistent with fully flexible prices, and notes that it can vary with labor-leisure preferences and with variations in exogenous government spending. That said, it hardly seems more reasonable to view the NBER recession periods as resulting from preference shocks that led people to demand more leisure, or from changes in government spending. Thus, even under this definition these recessions seem likely to have been periods during which output was below potential.

cost. In our opinion, this line of reasoning is suspect inasmuch as a good proxy for real marginal cost is likely to display a cyclical pattern that is the opposite of that manifested by labor's share. Thus, if the labor share version of the model were to fit well, then it would be unlikely that a model based on a marginal cost proxy manifesting completely different cyclical behavior would also fit well.

5.2 Empirical Performance

Despite our theoretical reservations, any assessment of the labor's share version of the new-Keynesian Phillips curve must ultimately rest upon the model's ability to match the data. If this variant of the NKPC works well empirically, then perhaps the theoretical problems just mentioned are not critical. It turns out, however, that the empirical success of this version of the NKPC is very limited.

One result that is often cited in favor of the model is that GMM estimation of

$$\pi_t = \beta E_t \pi_{t+1} + \gamma s_t, \tag{9}$$

where s_t denotes the log of labor's share of income, produces a correctly signed positive estimate for $\hat{\gamma}$ (as opposed to the negative sign that obtains from using detrended output). Line 2 of Table 1 confirms this result for our updated sample. Importantly, however, the results also indicate that a key result in Galí and Gertler (1999) is overturned: While positive, the coefficient on s_t is statistically indistinguishable from zero, so that one cannot reject the hypothesis that inflation and the labor share are completely unrelated. Experimentation revealed that this result is robust to changes in the instrument set. Moreover, the result also holds when we fit the model over the same sample period (1960:Q1 to 1997:Q4) that Galí and Gertler employed in their work, which suggests that their initial finding of a significant coefficient on labor's share can no longer be obtained in newer vintages of the data.[21]

Perhaps a more concrete way to illustrate the weakness of this model is to construct explicit measures of $E_t s_{t+k}$ using linear forecasting models.[22] For instance, one can fit the model

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t s_{t+k} \tag{10}$$

---

[21] While revisions to real GDP are often cited as problematic for the estimation of potential output and the GDP gap, it is worth noting that compensation and labor's share also commonly undergo substantial revisions in different vintages of the U.S. national accounts (see Koenig, 2003, for a discussion).

[22] The remaining results in this section update work discussed in greater detail in Rudd and Whelan (2005a).

using a VAR to forecast the values of s_{t+k} in a manner similar to Campbell and Shiller's (1987) methodology. Figure 2 shows the fit of the labor share variant of the NKPC that obtains from using a bivariate VAR in labor's share and H-P detrended output.[23] Both equations in the VAR fit well; in particular, detrended GDP enters as a highly significant predictor of future labor shares. Figure 2 reveals, however, that the discounted sum of expected labor shares (generated using a value of β equal to 0.99) provides a very poor model of inflation: This version of the NKPC model has an R^2 of only 0.13.

The fit for this model is, of course, well below what is obtained from simple autoregressive models of inflation: For example, an $AR(4)$ model receives an $\bar{R}^2$ of 0.75. The ability of an AR model to provide such a good characterization of the inflation process raises an obvious question: Why not use inflation to help construct forecasts of the future values of the driving term? In theory, if the NKPC model were correct, then lagged inflation would embody preceding periods' expectations about future labor shares, and thus should probably add useful forecasting power to the VAR.[24] However, in practice the data reject the hypothesis that inflation Granger causes the labor income share (see panel B of Table 2). In itself, this suggests that the essential story behind this model—that inflation reflects expectations of future labor shares—is not at all evident in the data.

One obvious implication of the absence of Granger causality is that one should not attempt to assess the fit of the NKPC model by generating $E_t s_{t+k}$ values from VAR systems that include inflation itself. However, even if one does so—which implies that the linear combination of variables defining the expected discounted sum will contain *contemporaneous* inflation—the resulting fit is still poor compared to what obtains from a reduced-form model. For instance, if one generates the discounted sum using a VAR in labor's share and inflation, the resulting fitted values for π_t have an R^2 of only 0.41, again well below what is obtained by simple reduced-form models. Together with the Granger causality results, this finding squarely contradicts the notion that the presence of lagged inflation in conventional empirical Phillips curves is attributable to its role as a proxy for future values of the labor share.

---

[23]We use two lags in the VAR based on the Schwartz criterion, but the result reported here is robust to using systems with longer lags.

[24]We say "probably" because it is possible that the labor share could follow a pure AR model. So long as it does not, however—and so long as agents have larger information sets than we do as econometricians—inflation should contain additional information for s_t if this version of the NKPC model is correct.

5.3 Assessment

The new-Keynesian Phillips curve clearly fails as an empirical description of inflation when detrended output is used as its driving variable. Recently, the idea that labor's share provides a theoretically preferable measure of inflationary pressure that permits the NKPC to fit the data well has gained substantial acceptance among a significant fraction of the profession. For example, one prominent researcher has offered his opinion that "the fact that simple optimizing models of pricing dynamics fit well when real unit labor costs is used as a proxy for the output gap, but not at all when detrended output is used, suggests that the real unit labor cost series is in fact the better measure of the output-gap concept that is relevant to inflation determination."[25]

By contrast, we would argue that neither the theoretical nor the empirical arguments in favor of the labor share version of the NKPC are anywhere as strong as its proponents have suggested. On theoretical grounds, much of the recent literature has been dismissive of the *ad hoc* nature of output gap measures based on the naïve detrending of real output, with the labor share approach promoted as being better founded in theory. However, proxying for marginal cost with average cost is also somewhat *ad hoc*, and a relatively simple examination of the data suggests to us that this measure provides a far less convincing proxy for the output gap than does detrended GDP. At a minimum, we think these considerations suggest that the measurement issues associated with proxying for real marginal cost are at least as serious as those surrounding the measurement of output gaps, and that claims for the theoretical superiority of labor's share over detrended GDP are largely rhetorical.

On empirical grounds, the idea that labor's share is a statistically significant determinant of inflation and, more specifically, that inflation is proportional to a discounted sum of expected future labor shares, receives very little support from the data. In particular, a comparison of the labor share version of the NKPC to traditional empirical Phillips curves, which fit the data well and incorporate a highly statistically significant relationship between detrended output and inflation, clearly reveals that the labor share approach suffers from substantial shortcomings.

---

[25]Woodford (2003, p. 616).

6 "Hybrid" New-Keynesian Models

The debate over the adequacy of the new-Keynesian Phillips curve is ongoing: Galí and Gertler (1999) characterize the labor share version of the model as representing "a good first approximation" to empirical inflation dynamics, while we have argued that the NKPC almost completely fails as an empirical model. However, even some of the NKPC's more enthusiastic supporters, including Galí and Gertler, concede that the model fails to fully capture the empirical dependence of inflation on its own lagged values. This has resulted in various proposals for so-called "hybrid" variants of the NKPC, which take the form

$$\pi_t = \gamma_f E_t \pi_{t+1} + \gamma_b \pi_{t-1} + \kappa x_t. \tag{11}$$

Hence, these models combine a rational forward-looking element with some dependence on lagged inflation, together with a measure x_t of inflationary pressures.

For many, this class of models represents a sort of common-sense middle ground that preserves the insights of standard rational expectations sticky-price models while allowing for better empirical fit by dealing directly with a well-known empirical deficiency of the pure forward-looking model. As a result, this class of models has been widely used in applied monetary policy analysis, with the policy implications depending importantly on the values of γ_b and γ_f.[26] Viewed from this perspective, then, the relative size of the γ_b and γ_f coefficients is a key issue; accordingly, Galí and Gertler (1999) and Galí, Gertler, and Lopez-Salido (2005) have reported influential and widely cited estimates of this equation from which they conclude that "forward-looking behavior is dominant" in empirical inflation dynamics. As a result, there is now a near-consensus among new-Keynesian economists that a hybrid model featuring a large γ_f and a small γ_b provides a sensible microfounded theoretical framework for analyzing the behavior of inflation.

Here, we want to offer some counterarguments to this consensus position. First, we argue that this class of models actually has very weak microfoundations and that these models are, most likely, as open to the Lucas critique as the traditional models that they seek to replace. Second, despite Galí and Gertler's widely cited estimates, we argue that the evidence for an important role for forward-looking behavior from empirical implementations of these models is, in fact, substantially weaker than is commonly believed.

---

[26]See, for instance, Clarida, Galí, and Gertler (1999).

6.1 Theoretical Issues

There are three well-known ways to motivate the hybrid model theoretically. In each case, the resulting models are arguably more *ad hoc* than microfounded.

Relative Wage Model: In an important and influential paper, Jeffrey Fuhrer and George Moore (1995) critiqued standard new-Keynesian pricing models as being incapable of explaining the persistence of the empirical inflation process, and suggested an alternative contracting model in which each cohort of workers bargains over its real wage with reference to the real wages earned by other cohorts of workers. This results in a hybrid model of the form

$$\pi_t = 0.5 E_t \pi_{t+1} + 0.5 \pi_{t-1} + \kappa y_t, \tag{12}$$

in which inflation depends equally on its lagged and expected future values, as well as on the output gap. While this model offers one potential explanation for the empirical role of lagged dependent variables in inflation regressions, it is not particularly obvious that the model's underlying contracting mechanism provides a reasonable characterization of how wages are actually set. In addition, Driscoll and Holden (2003) have demonstrated that the presence of a lagged inflation term in the model depends on somewhat arbitrary assumptions regarding which other cohorts a given cohort will compare its real wage against.

Rule-of-Thumb Price-Setters: Galí and Gertler (1999) provide an alternative motivation for a hybrid model in which nominal rigidities take the standard Calvo form for a subset of firms, while a remaining group of firms instead sets prices according to a "rule-of-thumb" that depends on the lagged inflation rate. Thus, their hybrid model takes the form

$$\pi_t = \gamma_f E_t \pi_{t+1} + \gamma_b \pi_{t-1} + \kappa s_t, \tag{13}$$

where s_t is again labor's share of income. Galí and Gertler motivate this model's departure from the pure rational expectations framework by appealing to the rule-of-thumb consumers model of Campbell and Mankiw (1989). However, it is much less clear in this case why part of the population behaves so differently from the rest. In the consumption example, the presence of rule-of-thumb consumers is often seen as resulting from the existence of liquidity constraints, which prevents the adoption of an optimal consumption path. In the pricing example, however, it is unclear why some firms would set prices in a optimal manner, while others would not.

Indexation: An alternative approach developed by Christiano, Eichenbaum, and Evans (2005) introduces a role for lagged inflation by changing the nature of price rigidities in the Calvo model. Indeed, in their model there is *no* price rigidity: All prices change in every period. However, at any point in time, only a random fraction of firms can reset their prices in a manner consistent with profit maximization; the remainder simply updates last period's price by "indexing" to last period's inflation rate. This yields a hybrid equation of the form

$$\pi_t = \frac{\beta}{1+\beta} E_t \pi_{t+1} + \frac{1}{1+\beta} \pi_{t-1} + \kappa x_t,$$

where β is the discount rate and x_t is real marginal cost. Relatedly, another well-known paper that invokes indexation to motivate a hybrid specification is Smets and Wouters (2003), in which the extent to which nonoptimizing firms index to the lagged inflation rate is a free parameter.

The indexation approach strikes us as providing an even less satisfactory theoretical motivation than the assumption of rule-of-thumb price-setters. The assumption that all prices change each period directly contradicts a substantial body of evidence regarding the nature and existence of nominal rigidities; indeed, such evidence provides the principal motivation for the new-Keynesian enterprise in the first place. The idea that firms find it easy to change prices each period, but somehow have difficulty doing so in an optimizing manner, also seems hard to justify. Finally, the Smets-Wouters approach raises the additional question of why "indexing" firms who change their price every period would then choose not to fully index to the lagged inflation rate, thereby in effect deliberately allowing their relative price to change systematically with inflation.

Beyond the specific questions raised by each of these theoretical motivations, it is worth noting that these models have been promoted as providing microfounded alternatives to reduced-form inflation equations, with estimated coefficients that are viewed as structural parameters that are immune to the Lucas critique. However, even if one is willing to accept theoretical devices such as indexation or the assumption of a fixed fraction of forward- and backward-looking price-setters, it is hard to imagine that parameters such as the fraction of backward-looking agents or the probability of setting an optimal non-indexed price would be "structural" parameters that are fixed across monetary policy regimes. Thus, it seems likely that models of this sort are just as vulnerable to the Lucas critique as are traditional econometric equations.

6.2 Importance of Forward-Looking Terms

We now discuss some of the evidence that has been presented on the relative importance of forward- and backward-looking terms in hybrid inflation equations.

6.2.1 Direct GMM Estimation

As before, one can estimate the hybrid specification (11) directly with GMM using variables dated time t or earlier as instruments. In this case, because equation (11) is a linear model, such a procedure is equivalent to a two-stage least squares approach in which fitted values from a first-stage regression of π_{t+1} on the instrument set are used to proxy for $E_t\pi_{t+1}$. Using this methodology (with x_t defined as labor's share), Galí and Gertler (1999) obtained large values for γ_f along with values for γ_b that were significantly smaller than those typically found in reduced-form inflation regressions; they concluded from these estimates that price setting is dominated by forward-looking behavior.

It turns out, however, that such estimates do not really allow us to distinguish between forward- and backward-looking models of inflation. This is because this methodology can signal the presence of an important role for forward-looking behavior—in the form of a large coefficient on $E_t\pi_{t+1}$—even when such behavior is actually completely absent. To understand why, note that the "second-stage" regression (the hybrid equation itself) is quite likely to be misspecified, since it almost certainly omits a number of variables that belong in the true model for inflation. Moreover, it is also highly likely that some of these omitted variables—or variables that are correlated with them—will be included in the instrument set that we use to implement this method (intuitively, anything that is correlated with π_t but not included in the hybrid model directly will serve as a good instrument for π_{t+1} in the first-stage regression). For example, Galí and Gertler included additional lags of inflation, commodity prices, and detrended output in their instrument set, all of which are typically used in empirical inflation equations. Hence, the constructed proxy for $E_t\pi_{t+1}$ will capture the influence of these omitted variables and receive a large coefficient even if $E_t\pi_{t+1}$ itself has no independent influence whatsoever on inflation. Furthermore, whatever role lagged inflation plays in the true inflation process will also be partially captured by the $E_t\pi_{t+1}$ proxy if lags of π_t are included in the instrument set, and so this method will necessarily tend to yield a small coefficient on lagged inflation.[27]

---

[27] See Rudd and Whelan (2005b) for a formal discussion of these points. In principle, we would be able to detect the presence of this misspecification with a test of the model's overidentifying restrictions (an OID

This problem motivates Rudd and Whelan's (2005b) demonstration that the estimates reported by Galí and Gertler are also consistent with the true model's being a completely backward-looking formulation containing some of the additional variables employed as (or correlated with) instruments for the GMM estimation. This example has been criticized by Galí, Gertler, and Lopez-Salido (2005) as being based on an "extreme scenario" in which there is no forward-looking behavior in price setting. However, the extreme nature of this scenario helps to make a general point: To the extent that Galí and Gertler's estimates could have been obtained in a world where forward-looking behavior is completely absent, they can hardly be viewed as compelling evidence that such behavior is dominant.

6.2.2 Closed-Form Solutions

Given that direct GMM estimation of equation (11) sheds little light on the relative importance of forward- and backward-looking terms, it is useful to consider other ways to examine this question. A particularly useful approach can be derived by noting that one can go somewhat further than the simple hybrid equation in testing for the presence of rational forward-looking behavior. This is because the assumption that the forward-looking term in this model represents a rational expectation exactly pins down the fundamental determinants of $E_t \pi_{t+1}$. Once we have solved for this fundamental solution, we obtain an exact closed-form expression for the labor's share version of the model of the form

$$\pi_t = \delta_1 \pi_{t-1} + \mu \sum_{k=0}^{\infty} \delta_2^{-k} E_t s_{t+k}, \tag{14}$$

where δ_1 and δ_2 are related to the roots of the polynomial equation

$$\gamma_b z^2 - z + \gamma_f = 0. \tag{15}$$

Equation (14) provides a more clear-cut way of distinguishing the hybrid model from a traditional Phillips curve. Both types of model allow lagged inflation to matter. However,

---

test). However, one can show that such tests will have essentially no power in this context, which suggests that it is not worthwhile to use OID tests in order to assess the empirical performance of the Calvo model (as Eichenbaum and Fisher, 2004, do).

It is worth noting that a similar argument implies that we cannot use survey measures of inflation expectations in order to estimate the new-Keynesian Phillips curve (as is done, for example, by Roberts, 1995). If private agents use variables to forecast inflation that are not included in the econometrician's test equation, then survey measures of expected inflation will tend to receive a large weight even if they play no direct role in determining current inflation.

while standard Phillips curve specifications use a conventional output gap measure to capture inflationary pressures, this version of the hybrid model sees the relevant driving term as being an expected discounted sum of future labor shares. It is therefore these current and expected labor share terms that are the distinguishing feature of this variant of the hybrid model. Hence, in our opinion, the key question is whether these terms are an important determinant of empirical inflation dynamics. The evidence suggests that they are not.

One way to reveal both the strengths and weaknesses of the labor's share version of the hybrid model is to again construct proxies for $E_t s_{t+k}$ using a VAR. To do this, we again used our two-lag VAR in labor's share and H-P filtered GDP. Figure 3 plots the predicted values of π_t from the best-fitting parameterization of the model. Specifically, we constructed separate discounted sum terms over a grid of values for the parameter δ_2^{-1} ranging between zero and one, and then used the one that gave the best fit for inflation (this turned out to obtain when $\delta_2^{-1} = 1$).

The positive news for the model from this chart is that the model's ability to fit the inflation process is relatively good. The model has an R^2 of 0.713, which is considerably higher than that received by the pure forward-looking NKPC. Indeed, Galí and Gertler (1999) also reported obtaining a good fit from this type of exercise—the fitted values shown here correspond to their "fundamental inflation" series from Figure 2 of their paper—and portrayed this as an important endorsement of the hybrid model.

There is bad news for the model, however, when one focuses on the question that we are examining; namely, how much of this good fit comes from including expectations of future labor shares? The answer turns out to be "almost none." Indeed, the R^2 from a pure $AR(1)$ model of inflation is 0.699, while the R^2 from a regression of inflation on its own lag and contemporaneous s_t is 0.704.[28] So, the good fit of this model can hardly be attributed to the contribution made by rational expectations of future labor shares.[29]

To be more precise about whether expected future labor shares play a statistically

---

[28]The fit of the hybrid model is slightly better ($R^2 = 0.730$) when one uses the VAR in s_t and π_t that Galí and Gertler employ. However, as we discussed in Section 5, the use of this VAR is suspect in this case because it allows for contemporaneous inflation to be one of the variables determining the expectational proxy, even though inflation does not Granger cause the labor share. In any case, even when this VAR is used, the gain in fit over a purely backward-looking model remains very small.

[29]A similar calculation is reported in Rudd and Whelan (2005b), which discusses the improvement in fit that results from including a discounted sum obtained from a first-stage GMM regression. It is reported there as well that the discounted sum of labor's shares contributes little to model fit.

significant role in explaining inflation, we also applied a GMM estimation methodology to the closed-form solution of the hybrid model, equation (14). To deal with the infinite sum in equation (14), we can *partially* solve forward equation (11) to obtain an *exact* expression in a finite number of terms,

$$\pi_t = \delta_1 \pi_{t-1} + \mu \sum_{k=0}^{N} \delta_2^{-k} E_t s_{t+k} + E_t \left[\delta_2^{-(N+1)} \left(\pi_{t+N+1} - \delta_1 \pi_{t+N} \right) \right], \qquad (16)$$

which can be directly estimated with GMM. Table 3 reports the estimated parameters obtained for this model under three different values of N (four, eight, and twelve quarters). We also consider three different instrument sets. Specifically, the first instrument set is as described earlier in footnote 14, and consists of two lags of each of the instruments employed by Galí and Gertler (1999). The second instrument set is the one used by Galí, Gertler, and Lopez-Salido (2005), which consists of four lags of inflation and two lags each of detrended output, wage inflation, and the labor income share. The final instrument set was used in a similar exercise in Rudd and Whelan (2003), and is the same as the Galí, *et al.* set but with two lags of inflation instead of four. For every one of these cases, the estimated μ coefficient, which describes the influence of the expected discounted sum of labor shares on inflation, is statistically insignificant.

These findings reinforce the principal conclusion drawn from the results of the previous section. There appears to be little empirical relationship between inflation and expectations of future values of the labor income share, and this conclusion is not much affected by the inclusion of a lagged dependent variable in the inflation specification. Similarly, while Figure 3 shows that the labor's share version of the hybrid model can provide a reasonably good fit for inflation, this fit turns out to have little to do with either the presence of rational forward-looking agents or with the use of labor's share as a proxy for real marginal cost.

6.2.3 Comparison with Galí, Gertler, and Lopez-Salido (2005)

Of course, there is always more than one way to interpret any set of empirical results. Galí, Gertler, and Lopez-Salido (2005) have also estimated equations similar to these, and have arrived at conclusions that appear to be the exact opposite of our own. For instance, Galí, *et al.* conclude that backward-looking behavior, "while statistically significant, is nonetheless quantitatively modest," and thus that "forward-looking behavior is dominant." Such an extreme range of opinions from researchers examining the same equations with

essentially the same data is, on the face of it, very odd. As such, we think it is worth trying to explain the source of this difference as best we can.

To start with, we should note that this difference in conclusions does not stem from the use of different data or a different estimation methodology; instead, it comes from Galí, *et al.*'s use of a very different metric for judging the importance of forward-looking behavior. Rather than focus on the role played by expected future labor shares, they focus on the values of γ_f and γ_b implied by the estimated parameters δ_1 and δ_2. Indeed, they argue that "the only way to obtain a proper sense of the relative importance of forward versus backward-looking behavior is to . . . obtain direct estimates of γ_f and γ_b." However, the relationship between the parameters δ_1 and δ_2 on the one hand, and γ_f and γ_b, on the other, is very complex (recall that they are related through the second-order polynomial equation 15). And it turns out that the parameters that Galí, *et al.* have focused on are in fact almost completely unrelated to the question that we have been asking, which is whether there is a statistically significant role for expected future labor shares.

To see why this is so, consider an example in which $\gamma_f + \gamma_b = 1$.[30] In this case, if $\gamma_f \leq 0.5$, then the closed-form solution is given by

$$\pi_t = \pi_{t-1} + \mu \sum_{k=0}^{\infty} \left(\frac{\gamma_f}{1 - \gamma_f} \right)^k E_t s_{t+k}, \tag{17}$$

while if $\gamma_f \geq 0.5$, then the closed-form solution is

$$\pi_t = \left(\frac{1 - \gamma_f}{\gamma_f} \right) \pi_{t-1} + \mu \sum_{k=0}^{\infty} E_t s_{t+k}. \tag{18}$$

Now consider the properties of the estimates of γ_f that will be drawn from the closed-form solution. First, note that as long as the point estimate of the lagged inflation coefficient, δ_1, is less than one (as is likely to be the case in practice), then Galí, *et al.*'s procedure will estimate γ_f to be greater than one-half.[31] Second, this estimate will be obtained even if the only term that has explanatory power is lagged inflation. This latter point is particularly important. Even if the essential implication of rational expectations in this model—a dependence of inflation on expected future values of the labor share—is completely absent,

---

[30]See Rudd and Whelan (2003) for an extensive discussion of this class of models.

[31]Note that the estimation procedures being discussed here do not attempt to obtain median-unbiased estimates of δ_1, so even if the true value of this parameter were one, small-sample biases would ensure that the point estimate obtained would be less than one.

Galí, *et al.*'s metric will still signal that rational forward-looking behavior is "dominant." Hence, even when the true model is completely backward-looking, this procedure would imply the opposite. Indeed, as Table 3 shows, the values of δ_1 and δ_2 that we obtained imply values of γ_f that are almost always greater than γ_b, even though the μ coefficients were invariably insignificant.[32] In addition, we note that the implied values of γ_b and γ_f always sum to a value close to one, so the analytical example just discussed provides a useful intuition for how to interpret these estimates.

For these reasons, we do not believe that the estimates of γ_f and γ_b that are generated with this method can provide a useful way of assessing the importance of the rational forward-looking component of a hybrid inflation equation. In the absence of any significant role for expected future labor shares, the idea that the data reflect a mixture of forward- and backward-looking behavior can be rejected, and so any implied estimates of γ_f and γ_b are completely irrelevant. We should also note that this critique is not restricted to GMM-based estimation, but also applies to so-called maximum-likelihood estimation of rational-expectations pricing models, in which values of the "structural parameters" implied by model-consistent solutions such as these closed-form equations are computed.[33]

Beyond the estimation issues that are involved, the literature's focus on the specific point estimates of γ_f and γ_b obtained from hybrid inflation equations has, in our opinion, led many researchers to misapprehend both how badly the pure NKPC model fits the data, as well as how poorly it explains the presence and importance of lagged inflation in reduced-form inflation equations. For instance, the closed-form solution in the case $\gamma_b = \gamma_f = 0.5$ is often treated as providing a middle ground between the accelerationist Phillips curve and the "pure" NKPC. However, in practice, these γ_f and γ_b values will yield a closed-form solution such that $\delta_1 = 1$, and thus the dependence of inflation on its own lagged values will be exactly as predicted by the accelerationist model.

In addition, even if one wishes to accept Galí, *et al.*'s parameter estimates (which imply

---

[32]For completeness, the table also reports the values of κ (the coefficient on the driving variable in the hybrid inflation equation 11) that are implied by the estimated parameters μ, δ_1, and δ_2^{-1}. In line with our estimates of μ, we invariably find that κ is statistically indistinguishable from zero; this contrasts with Galí, *et al.*, who report estimates of this coefficient that are highly significant. Once again, the source of this difference appears to be our use of a more recent vintage of the labor's share series (as opposed to our slightly longer estimation period): If we use our data to re-estimate the model over the 1960:Q1 to 1997:Q4 sample that Galí, *et al.* employ, we continue to find estimates of κ that are statistically insignificant.

[33]See Fuhrer (1997) for an example.

that $\gamma_f > \gamma_b$) and thus their assessment that backward-looking behavior (as they define it) is "quantitatively modest," it remains the case that the lagged inflation term provides essentially all of the explanatory power in equation (14), in stark contradiction to the predictions of the pure NKPC model. Thus, the commonly drawn conclusion that a finding of $\gamma_f > \gamma_b$ suggests that the NKPC is more right than wrong—that it provides "a good first approximation" to the data—is incorrect.

6.2.4 Sbordone's (2005) Calculations

Before concluding our discussion of this topic, we would note some recent results reported by Argia Sbordone (2005) that also appear to contradict several of the conclusions reported above. Sbordone provides estimates suggesting that the performance of a model of the same form as equation (14) worsens substantially when expectations of future values of the labor share are excluded, while almost no deterioration in fit obtains when the lagged inflation term is omitted. Again, the reason for this difference in conclusions stems from the use of a metric to assess the relative success of forward- and backward-looking inflation models that differs substantially from our own.

Sbordone assesses the model's performance using dynamic out-of-sample forecasts. Thus, her assessment of the model "without forward-looking terms" is based on finding the best-fitting series from a dynamic simulation of

$$\pi_t = \delta_1 \pi_{t-1} + \mu s_t. \tag{19}$$

This means that, for most of the historical sample, Sbordone's fitted values for her "backward-looking model" are essentially a geometrically weighted moving average of current and past values of the labor share of income. Because labor's share provides a very poor measure of inflationary pressures, we would be highly surprised if such a series were in fact to predict inflation well, and indeed Sbordone demonstrates that it does not. However, this hardly constitutes a useful piece of evidence against traditional Phillips curve specifications. Such models imply that current inflation depends on lagged values of *actual* inflation (and most definitely not a weighted average of past labor shares), and models based on this assumption clearly fit the data quite well.

Conversely, Sbordone's demonstration of the ability of a hybrid NKPC to fit the data is based on a dynamic simulation of

$$\pi_t = \delta_1 \pi_{t-1} + \mu (I - \beta A)^{-1} Z_t, \tag{20}$$

28

where $(I - \beta A)^{-1} Z_t$ is the expected discounted sum of future labor shares generated by a VAR in the companion form $Z_t = A Z_{t-1} + \epsilon_t$. In the results she reports, Sbordone uses a bivariate VAR in labor's share and inflation. Thus, the linear combination $(I - \beta A)^{-1} Z_t$ contains π_t (as well as its lags). Sbordone's finding of success for this model therefore stems in part from her use of the *actual value of π_{t+k} itself* in the information set that is used to generate the simulated values of π_{t+k}. More generally, as discussed above, this particular forecasting VAR is only valid for this purpose when there is evidence that inflation Granger causes labor's share, a result that we have failed to find here or elsewhere (*c.f.* Rudd and Whelan, 2005a).

Clearly, then, we do not view Sbordone's methodology as providing a useful way to assess the relative merits of backward- and forward-looking models. However, it does raise one interesting question, which is whether reduced-form econometric Phillips curves can perform well in a dynamic simulation. Figure 4 reports results from one such exercise for a "triangle-style" econometric Phillips curve for nonfarm business price inflation. The model features four lags of inflation, two lags of the H-P filtered output gap, and one lag of a supply shock term (defined as goods import price inflation *minus* NFB price inflation). The model was estimated over the full sample, 1960:Q1 to 2004:Q3, and then simulated jumping off of 1959:Q4.[34]

The chart shows that the simulated series tracks the historical behavior of inflation remarkably well. An economist living in 1959, and armed with full information about the future values of the output gap and import prices, could have constructed an exceptionally accurate out-of-sample forecast of U.S. inflation over the period 1960-2004 with this model. Of course, in practice, the accuracy of forecasts constructed from this type of model will depend upon how well one can forecast its driving terms (in this case, import prices and the output gap). But this is true of any forecasting model—for example, the NKPC would require accurate forecasts of real marginal cost to be useful as a forecasting tool. From a practical perspective, the relevant question is whether a reduced-form econometric model based on an assumed propagation mechanism involving substantial inertia, can generate sensible inflation projections. Figure 4 suggests that it can.[35]

---

[34]Note that this is a true dynamic simulation, in the sense that the measure of NFB price inflation used to construct the supply shock in each period is the simulated inflation rate, as opposed to its actual value.

[35]This last point is important in light of Sbordone's claim that lagged inflation is basically immaterial to inflation dynamics. Sbordone arrives at this conclusion because her dynamic simulation of $\pi_t = \delta_1 \pi_{t-1} + \mu s_t$ (equation 19) fits badly. But this tells us little more than that labor's share is a poor choice of driving variable

7 Other Models of Sticky Prices

The hybrid specifications discussed above are typically derived from sticky-price models with Calvo-style pricing rigidities. Because the Calvo framework carries with it several unappealing features, a number of researchers have suggested that estimates of a positive role for backward-looking behavior from these models may simply reflect a failure to correctly model the nature of the price-setting mechanism. For instance, Galí, Gertler, and Lopez-Salido (2005) suggest that models with Taylor-style price rigidities and purely rational agents may explain their finding of a positive role for lagged inflation. Similarly, Dotsey (2002) and Bakhshi, *et al.* (2003) show that "truncated" Calvo models, which rule out the possibility of price durations of arbitrarily long length, can generate data consistent with positive estimates of γ_b in hybrid models. Here, we briefly outline why we do not view these alternative models as providing an adequate explanation for the presence of lagged inflation in empirical regressions.[36]

Consider the general class of sticky-price models described by Goodfriend and King (1997). The structure of the price rigidity in these models is summarized by a hazard function $(a_1, a_2, \ldots, a_{n-1}, a_n)$ with a_k defined as the probability that a price that has remained in place for k periods will be changed. It is assumed that $a_n = 1$, so there is a maximum length of time that any price can remain fixed. Defining $f_0 = 1$ and

$$f_k = \prod_{r=1}^{k} (1 - a_r), \tag{21}$$

then this type of price rigidity implies convergence towards a stable distribution of uncompleted price spells, such that the share of firms with prices that are k periods old equals

$$s_k = \frac{f_k}{\sum_{i=0}^{n-1} f_i}. \tag{22}$$

Assuming the aggregate price index that obtains under Dixit-Stiglitz preferences, the price level can be log-linearized as

$$p_t = \sum_{k=0}^{n-1} s_k x_{t-k}, \tag{23}$$

where x_t is the log of the new contract price set at time t. Given this structure of price

---

for an inflation regression.

[36]This section presents results from a revised draft of a paper previously circulated as Whelan (2004).

rigidity, the log-linearized optimal contract price is

$$x_t = \frac{E_t\left[\sum_{k=0}^{n-1} \beta^k s_k \left(p_{t+k} + \gamma y_{t+k}\right)\right]}{\sum_{k=0}^{n-1} \beta^k s_k},$$

(24)

where we have again assumed a relationship between real marginal cost and the output gap, y_t.

Together, equations (23) and (24) determine the dynamics of inflation in this framework. In an appendix, we show that they imply an inflation equation of the form

$$\pi_t = \psi(L)\pi_{t-1} + \frac{\gamma\alpha(L)}{s_0 s_{n-1}}\left[\sum_{k=0}^{\infty} \kappa_k E_t Z_{t+k}\right],$$

(25)

where

$$\psi(L) = \sum_{k=1}^{n-2} \psi_k L^k,$$

(26)

$$\alpha(L) = \left(\sum_{j=0}^{n-1} s_j L^j\right), \text{ and}$$

(27)

$$Z_t = \sum_{k=0}^{n-1} \beta^k s_k y_{t+k}.$$

(28)

This result shows that this class of models does indeed allow for a role for lagged inflation, even after one has controlled for the role played by economic fundamentals (meaning output or real marginal cost). However, it turns out that the nature of this role is dramatically different from that implied by empirical evidence: The models predict that these lagged inflation coefficients will be *negative*.

To illustrate this pattern, Table 4 reports the coefficients of the lag polynomial $\psi(L)$ for various sticky-price models. First, the upper panel reports values for Taylor contracting models with contract lengths of three to six periods; these models therefore have hazard functions with $(a_1, a_2, \ldots, a_{n-1}, a_n) = (0, 0, \ldots, 0, 1)$. Next, the middle two panels report values for two different versions of the truncated Calvo model—described by hazards of the form $(a_1, a_2, \ldots, a_{n-1}, a_n) = (\alpha, \alpha, \ldots, \alpha, 1)$—that correspond to two different values of the price reset probability, α. Finally, the bottom panel reports values for upward-sloping hazard functions, as generated by state-dependent frameworks such as the model of Dotsey, King, and Wolman (1999). For each of these models, all of the coefficients on lagged inflation are negative.

As is evident from these calculations, therefore, we will not be able to invoke rational expectations models based on these alternative formulations of pricing rigidities in order to explain why lagged inflation is present in empirical inflation equations. Indeed, the fact that these models make a prediction about the inflation process that is so clearly at variance with reality is strong *prima facie* evidence against their use.

8 Parameter Stability in Reduced-Form Phillips Curves

The preceding sections have demonstrated that there is limited evidence for the type of rational forward-looking behavior implied by a currently popular class of theoretical models. However, it is important to distinguish between the predictions of specific models, on the one hand, and the more general implications of rational expectations, on the other.

For instance, our failure to find an important empirical role for expected future values of the labor share may simply reflect the failure of this variable as a proxy for inflationary pressures rather than the failure of rational expectations *per se*. Even if today's vintage of rational expectations models are flawed, the Lucas-Sargent critique of econometric Phillips curves remains a powerful and intuitive theoretical argument. The mechanism by which the private sector formulates expectations is clearly crucial to the behavior of the inflation process, and it seems perfectly reasonable that this mechanism could change over time, particularly as monetary policy regimes change. In particular, as noted already, it is widely conjectured that the strong anti-inflation credibility established by the Federal Reserve during the Volcker-Greenspan era could have changed the way in which expectations of inflation are formed, so that expectations now depend less on lagged inflation. As a result, this component of a reduced-form Phillips curve could be unstable across policy regimes.

However, our own investigation of this issue suggests that standard econometric models remain useful tools. That dynamic simulations of reduced-form Phillips curves can match history over a long period of time (as shown in Figure 4) is one informal diagnostic that suggests that the parameters of the inflation process are unlikely to have changed too much. And, more formal econometric analysis indicates that the manner in which inflation depends on its own lags might be more stable than is commonly thought. We consider this next.

We conducted a rolling-regression analysis (based on 48-quarter windows) of a univariate

specification for NFB price inflation,

$$\pi_t = \alpha + \rho\pi_{t-1} + \sum_{k=1}^{N} \mu_k \Delta\pi_{t-k} + \epsilon_t, \tag{29}$$

and of the full triangle-style specification,

$$\pi_t = \alpha + \rho\pi_{t-1} + \sum_{k=1}^{N} \mu_k \Delta\pi_{t-k} + \lambda_1 y_{t-1} + \lambda_2 y_{t-2} + \varsigma\left(\pi_{t-1}^m - \pi_{t-1}\right) + \epsilon_t, \tag{30}$$

where y_t is the H-P filtered output gap and π_t^m is imported goods price inflation.[37] One important feature of our calculations is that we attempt to control for finite-sample biases. It is well known that OLS estimates of ρ for persistent series tend to be downward biased, and that these biases are particularly large in small samples (such as those being used in our rolling-window exercise).[38] Thus, we calculate median-unbiased estimates of ρ, together with 90 percent confidence intervals, using Bruce Hansen's (1999) grid-bootstrap methodology.

The upper panel of Figure 5 reports the results for the univariate model. For most of the twelve-year windows, the median-unbiased estimates of ρ are close to one. (The relative stability of estimates of this sort has been previously noted by Pivetta and Reis, 2004.) For the more recent twelve-year windows, the median-unbiased estimates of ρ decline noticeably, with the final estimate coming in as low as 0.55. However, it should also be noted that some of the twelve-year periods prior to 1994 also exhibited low values of ρ, and that this pattern did not persist. In addition, the 90 percent confidence intervals for ρ rarely exclude unity, even in these later periods.

Perhaps more importantly, the evidence for systematic change in the sum of the lag coefficients becomes even weaker when we consider a multivariate specification including lagged output gaps and supply shocks. In this case, the median-unbiased estimates display a flatter pattern over the sample, with the most-recent estimates remaining near one. On balance, then, we conclude that despite there being a theoretical possibility that shifts in monetary policy could have changed the reduced-form inflation process substantially, the evidence still points to a surprising degree of stability in the dependence of inflation on its own lags.[39]

---

[37] A value of $N = 6$ was chosen, consistent with applying the Akaike Information Criterion (AIC) to the full sample period. Our use of the AIC is informed by the work of Ng and Perron (2003), who show that it is important not to be too parsimonious in lag selection when assessing the value of ρ.

[38] Andrews (1993) calculates the size of these biases.

[39] O'Reilly and Whelan (2005) reach a similar conclusion about the Euro-area's inflation process.

9   Conclusions

The history of science provides many examples of theories that everyone knew were true, until they turned out to be false. At various points in history, intelligent people knew that the world was flat; knew that the sun revolved around the earth; and knew that there was an exploitable long-run tradeoff between inflation and unemployment. Today, one can meet many researchers who know that the new-Keynesian Phillips curve provides a good model of the inflation process once one uses a suitable proxy for real marginal cost; who know that forward-looking rational behavior dominates price setting; and who know that the U.S. inflation process has fundamentally changed in the last twenty years because of a shift in how monetary policy has been conducted. We hope the evidence presented in this paper will give at least some interested researchers cause to check these conclusions more fully against the available evidence.

At the start of this paper, we posed two fundamental questions about inflation: first, what is a suitable measure of inflationary pressures; and second, how are inflation expectations formed. Regarding the former issue, we believe the evidence presented here militates strongly against the use of labor's share of income as a useful proxy for inflationary pressures. In general, we find very little evidence for a strong link between inflation and current or expected values of this variable. Just as important to note in this context is the often-forgotten fact that traditional output gaps remain highly significant explanatory variables in reduced-form inflation regressions.

Regarding expectations formation, we view the failures of the new-Keynesian Phillips curve as likely reflecting the weakness of rational expectations as a description of reality. The new-Keynesian inflation equation makes three assumptions about price-setting behavior: first, that prices are sticky; second, that agents optimize their behavior given that their prices are fixed; and third, that agents' expectations are formulated in a *rational—i.e.*, model-consistent—manner. Empirical studies suggest that a significant degree of price stickiness is present in the U.S. economy, and this in turn almost certainly implies that firms attempt to make some prediction about future inflation when determining their current price. What appears to be less reasonable, however, is the assumption that these predictions are formulated in the manner implied by the new-Keynesian model under rational expectations, at least when either a traditional output gap or labor's share of income is employed as a proxy for real marginal cost.

Of course, we would not wish to overstate the implications of our results for the merits

of the rational expectations hypothesis; in particular, they do not rule out the possibility that the rational expectations approach might better fit the data with an alternative proxy for real marginal cost. Nevertheless, we believe there is little evidence at present that structural modelling of inflation in a rational expectations framework provides a clearly superior approach relative to traditional models of inflation dynamics. Reduced-form econometric models of inflation are surprisingly stable, and these models remain useful tools for forecasting. This may be because simple models of adaptive expectations, as described by Friedman and Phelps, still provide decent approximations to reality. Or it could instead be that existing rational expectations mechanisms do not capture the richness of how expectations are formulated in the real world. Further research on how expectations evolve and interact with the policy environment—such as recent work by Sargent (1999) and Orphanides and Williams (2005)—may well prove useful in advancing our understanding of these vital questions.

References

[1] Andrews, Donald (1993). "Exactly Median-Unbiased Estimation of First-Order Autoregressive/Unit Root Models," *Econometrica*, 61, 139-165.

[2] Ball, Laurence (1994). "Credible Disinflation with Staggered Price Setting," *American Economic Review*, 84, 282-289.

[3] Ball, Laurence and N. Gregory Mankiw (2002). "The NAIRU in Theory and Practice," *Journal of Economic Perspectives*.

[4] Ball, Laurence and Robert Moffitt (2002). "Productivity Growth and the Phillips Curve" in *The Roaring Nineties: Can Full Employment Be Sustained?* edited by Alan Krueger and Robert Solow. Russell Sage Foundation.

[5] Bakhshi, Hasan, Pablo Burriel-Llombart, Hashmat Khan, and Barbara Rudolf (2003). "Price-Setting Behavior and Inflation Dynamics," mimeo, Bank of England.

[6] Bils, Mark and Peter Klenow (2004). "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947-985.

[7] Blinder, Alan (1998). *Central Banking in Theory and Practice*. MIT Press.

[8] Braun, Steven (1984). "Productivity and the NIIRU (And Other Phillips Curve Issues)," Federal Reserve Board of Governors Working Paper Series No. 34.

[9] Calvo, Gulliermo A. (1983). "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, 12, 383-398.

[10] Campbell, John and Robert Shiller (1987). "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062-1088.

[11] Campbell, John and N. Gregory Mankiw (1989). "Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence," *NBER Macroeconomics Annual*, 3, 185-216.

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

[13] Clarida, Richard, Jordi Galí, and Mark Gertler (1999). "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661-1707.

[14] Cogley, Timothy and Thomas Sargent (2002). "Evolving Post-World War II Inflation Dynamics" in *NBER Macroeconomics Annual*, Vol. 16.

[15] 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, 655-690.

[16] Dotsey, Michael (2002). "Pitfalls in Interpreting Tests of Backward-Looking Pricing in New-Keynesian Models," *Federal Reserve Bank of Richmond Quarterly*, Winter, 37-50.

[17] Driscoll, John and Steinar Holden (2003). "Inflation Persistence and Relative Contracting," *American Economic Review*, 93, 1369-1372.

[18] Eichenbaum, Martin and Jonas Fisher (2004). "Evaluating the Calvo Model of Sticky Prices," mimeo, Northwestern University.

[19] Fischer, Stanley (1977). "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy*, 85, 191-205.

[20] Friedman, Milton (1968). "The Role of Monetary Policy," *American Economic Review*, 58, 1-17.

[21] Fuhrer, Jeffrey and George Moore (1995). "Inflation Persistence," *Quarterly Journal of Economics*, 110, 127-159.

[22] Fuhrer, Jeffrey (1997). "The (Un)Importance of Forward-Looking Behavior in Price Specifications," *Journal of Money, Credit, and Banking*, 29, 338-350.

[23] Galí, Jordi (1999). "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations," *American Economic Review*, 89, 249-271.

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

[25] Galí, Jordi, Mark Gertler, and David Lopez-Salido (2005). "Robustness of Estimates of the Hybrid New Keynesian Phillips Curve," *Journal of Monetary Economics*, 52, September.

[26] Goodfriend, Marvin and Robert King (1997). "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual*, 12, 231-283.

[27] Gordon, Robert J. (1976). "Recent Developments in the Theory of Inflation and Unemployment," *Journal of Monetary Economics*, 2, 185-220.

[28] Gordon, Robert J. (1988). "The Role of Wages in the Inflation Process," *American Economic Review*, Papers and Proceedings, 78, 276-283.

[29] Gordon, Robert J. (1998). "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU," *Brookings Papers on Economic Activity*, Volume 2, 297-346.

[30] Hansen, Bruce (1999). "The Grid Bootstrap and the Autoregressive Model," *Review of Economics and Statistics*, 81, 594-607.

[31] Koenig, Evan (2003). "Is the Markup a Useful Real-Time Predictor of Inflation?"" *Economics Letters*, 80, 261-267.

[32] Lucas, Robert (1972a). "Econometric Testing of the Natural Rate Hypothesis" in Otto Eckstein (ed.), *The Econometrics of Price Determination Conference*, reprinted in *Studies in Business Cycle Theory*, 1981, Cambridge: MIT Press.

[33] Lucas, Robert (1972b). "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103-124.

[34] Lucas, Robert (1976). "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Series on Public Policy*, 1, 19-46.

[35] McCallum, Bennett (1976). "Rational Expectations and the Natural Rate Hypothesis: Some Consistent Estimates," *Econometrica*, 44, 43-52.

[36] McCallum, Bennett (1994). "Identification of Inflation-Unemployment Tradeoffs in the 1970s: A Comment," *Carnegie-Rochester Conference Series on Public Policy*, 41, 231-241.

[37] Ng, Serena and Pierre Perron (2001). "Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power," *Econometrica*, 69, 1519-1554.

[38] Okun, Arthur M. (1975). "Inflation: Its Mechanics and Welfare Costs," *Brookings Papers on Economic Activity*, 2, 351-390.

[39] O'Reilly, Gerard and Karl Whelan (2005). "Has Euro-Area Inflation Persistence Changed Over Time?" *Review of Economics and Statistics*, 87(4).

[40] Orphanides, Athanasios and John C. Williams (2005). "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in *The Inflation Targeting Debate*, edited by Ben Bernanke and Michael Woodford, University of Chicago Press.

[41] Phelps, Edmund (1967). "Phillips Curves, Expectations of Inflation, and Optimal Inflation Over Time," *Economica*, 135, 254-281.

[42] Phillips, A.W. (1958). "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, 25, 283-99.

[43] Pivetta, Frederic, and Ricardo Reis (2004). "The Persistence of Inflation in the United States," mimeo, Harvard University (February).

[44] Roberts, John M. (1995). "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit, and Banking*, 27, 975-984.

[45] Romer, Christina and David Romer (2002). "The Evolution of Economic Understanding and Postwar Stabilization Policy," in *Rethinking Stabilization Policy*, Federal Reserve Bank of Kansas City.

[46] Rotemberg, Julio and Michael Woodford (1999). "The Cyclical Behavior of Prices and Costs," in *The Handbook of Macroeconomics*, edited by John Taylor and Michael Woodford. North-Holland.

[47] Rudd, Jeremy and Karl Whelan (2003). "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" Federal Reserve Board Finance and Economics Discussion Series No. 2003-46 (forthcoming, *American Economic Review*).

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

[49] Rudd, Jeremy and Karl Whelan (2005b). "New Tests of the New-Keynesian Phillips Curve," *Journal of Monetary Economics*, 52, 1167-1181.

[50] Samuelson, Paul and Robert Solow (1960) "Analytical Aspects of Anti-Inflation Policy," *American Economic Review*, 50, 177-94.

[51] Sargent, Thomas (1971). "A Note on the 'Accelerationist' Controversy," *Journal of Money, Credit, and Banking*, 3, 721-725.

[52] Sargent, Thomas (1999). *The Conquest of American Inflation*, Princeton University Press.

[53] Sbordone, Argia (2002). "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics*, 49(2), 265-292.

[54] Sbordone, Argia (2005). "Do Expected Future Marginal Costs Drive Inflation Dynamics?" *Journal of Monetary Economics*, September.

[55] Smets, Frank and Raffeal Wouters (2003). "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economics Association,* 1, 1123-1175.

[56] Solow, Robert (1968). "Recent Controversy on the Theory of Inflation: An Eclectic View," in *Inflation: Its Causes, Consequences, and Control*, edited by Stephen W. Rousseas. New York: New York University.

[57] Solow, Robert (1969). *Price Expectations and the Behavior of the Price Level*, University of Manchester Press.

[58] Solow, Robert (1976). "Down the Phillips Curve with Gun and Camera," in *Inflation, Trade and Taxes: Essays in Honor of Alice Bourneuf*, David Belsey (ed.), Columbus: Ohio State University Press.

[59] Staiger, Douglas, James Stock, and Mark Watson (1997a). "How Precise Are Estimates of the Natural Rate of Unemployment?" in Christina Romer and David Romer (eds.), *Reducing Inflation*, Chicago: University of Chicago Press.

[60] Staiger, Douglas, James Stock, and Mark Watson (1997b). "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives*, 11, Winter, 33-49.

[61] Stock, James and Motohiro Yogo (2003). "Testing for Weak Instruments in Linear IV Regression," mimeo, Harvard University (February draft).

[62] Taylor, John (1979). "Staggered Wage Setting in a Macro Model," *American Economic Review*, Papers and Proceedings, Vol. 69, 108-113.

[63] Taylor, John (1998). "Monetary Policy Guidelines for Unemployment and Inflation Stability," in *Inflation, Unemployment, and Monetary Policy*, edited by Robert Solow and John Taylor, MIT Press.

[64] Tobin, James (1972). "The Wage-Price Mechanism: Overview of the Conference," in Otto Eckstein (ed.), *The Econometrics of Price Determination Conference*.

[65] Whelan, Karl (2004). "Staggered Price Contracts and Inflation Persistence: Some General Results," ECB Working Paper No. 417.

[66] Woodford, Michael (2003). *Interest and Prices*, Princeton University Press.

A A General Hazard Model of Price Rigidity

This appendix outlines the model underlying the calculations reported in Section 7. We consider an economy consisting of a continuum of imperfectly competitive firms who have demand functions derived from a Dixit-Stiglitz preference structure. The economy is therefore characterized by a continuum of firms such that firm i with price P_{it} has demand function

$$Y_{it} = Y_t \left(\frac{P_{it}}{P_t} \right)^{-\theta},$$

where Y_t is total output and P_t is the aggregate price level. All firms in the model are identical, so each firm that sets a price at time t will set the same price. We label the new price set at time t as X_t. Thus, the aggregate price level consistent with Dixit-Stiglitz preferences is defined as

$$P_t = \left(\sum_{k=0}^{n-1} s_{kt} X_{t-k}^{1-\theta} \right)^{\frac{1}{1-\theta}},$$

where s_{kt} is the fraction of prices that are k periods old at time t.

Under the assumption that the economy has converged to a point at which the shares s_{kt} are constant over time, this price-level equation log-linearizes to yield

$$p_t = \sum_{k=0}^{n-1} s_k x_{t-k}.$$

Given the market structure and the hazard function for price changes, firms that set a price at time t choose X_t to maximize the discounted sum of expected real profits over the life of the contract. Note that f_k (the probability that the price will last as far as period $t+k$) is directly proportional to s_k (the steady-state share of firms whose price was set k periods ago). Thus, the firm's problem consists of maximizing

$$E_t \left[\sum_{k=0}^{n-1} \beta^k s_k \left(Y_{t+k} P_{t+k}^{\theta-1} X_t^{1-\theta} - \frac{1}{P_{t+k}} C_{t+k} \left(Y_{t+k} P_{t+k}^{\theta} X_t^{-\theta} \right) \right) \right],$$

where β is the firm's discount factor and C_t is its nominal cost function at time t. Solving this problem yields the following formula for the optimal contract price:

$$X_t = \frac{\theta}{\theta-1} \frac{E_t \left(\sum_{k=0}^{n-1} s_k \beta^k Y_{t+k} P_{t+k}^{\theta-1} \widetilde{MC}_{t+k} \right)}{E_t \left(\sum_{k=0}^{n-1} s_k \beta^k Y_{t+k} P_{t+k}^{\theta-1} \right)},$$

where $\widetilde{MC}_t$ is the firm's nominal marginal cost at time t. Log-linearizing this expression around a constant output level and a zero inflation rate, we obtain

$$x_t = \frac{E_t\left[\sum_{k=0}^{n-1}\beta^k s_k \widetilde{mc}_{t+k}\right]}{\sum_{k=0}^{n-1}\beta^k s_k},$$

where lower-case letters correspond to logged variables. Defining real marginal cost as

$$MC_t = \frac{\widetilde{MC}_t}{P_t},$$

and assuming a simple relationship between the log of real marginal cost and the output gap, $mc_t = \gamma y_t$, the optimal contract price becomes

$$x_t = \frac{E_t\left[\sum_{k=0}^{n-1}\beta^k s_k \left(p_{t+k} + \gamma y_{t+k}\right)\right]}{\sum_{k=0}^{n-1}\beta^k s_k}.$$

The first step in solving for the process for the contract price is to substitute out the expected future price levels in terms of future and past contract prices to obtain the following expression for the contract price:

$$x_t = \frac{E_t\left[\sum_{k=0}^{n-1}\beta^k s_k \left(\sum_{r=0}^{n-1} s_r x_{t+k-r} + \gamma y_{t+k}\right)\right]}{\sum_{k=0}^{n-1}\beta^k s_k}.$$

This is a $(2n-2)$th-order stochastic difference equation in x_t and the properties of its solution underlie the properties of aggregate price inflation in this model. Defining the following polynomial:

$$\sigma\left(x\right) = \sum_{k=1}^{n-1}\left(\sum_{r=0}^{n-k-1}\beta^r s_r s_{r+k}\right)x^k,$$

it turns out that the contract-price process can be re-written in terms of lag and forward operators as

$$E_t\left[\sigma\left(\beta F\right) - \left(\sum_{k=0}^{n-1}\beta^k s_k\left(1-s_k\right)\right) + \sigma\left(L\right)\right]x_t = -\gamma Z_t,$$

where

$$Z_t = \sum_{k=0}^{n-1}\beta^k s_k y_{t+k}.$$

The characteristic polynomial of this difference equation has a unit root, and also has the property that if λ_i is a solution, then $\left(\beta^{-1}\lambda_i^{-1}\right)$ is also a solution. These properties allow the contract-price process to be re-written as

$$E_t\left[s_0 s_{n-1}\beta^{n-1}(F-1)(F-\beta^{-1})\left\{\prod_{i=1}^{n-2}\left(F-\lambda_i\right)\left(F-\frac{1}{\beta\lambda_i}\right)\right\}L^{n-1}x_t\right] = -\gamma Z_t.$$

A stable solution exists when there are $n-1$ roots on or inside the unit circle. Let $\lambda_1, \lambda_2, ..., \lambda_{n-2}$ be the roots inside the unit circle and let $\lambda_{n-1} = 1$. Using the general principle of solving stable roots backwards and unstable roots forward, we need to apply the $n-1$ lag operators L^{n-1} to the roots on or inside the unit circle to leave only one possible non-explosive solution. Thus, the solution can be written as:

$$s_0 s_{n-1} \beta^{n-1} \left\{ \prod_{i=1}^{n-2} (1 - \lambda_i L) \right\} \Delta x_t = -\gamma E_t \left[\left\{ \prod_{i=1}^{n-1} (F - \frac{1}{\beta \lambda_i}) \right\}^{-1} Z_t \right].$$

Note also that

$$\frac{1}{F - (\beta \lambda_i)^{-1}} = -\frac{\beta \lambda_i}{1 - \beta \lambda_i F} = -\beta \lambda_i \sum_{k=0}^{\infty} \beta^k \lambda_i^k F^k.$$

Thus, the contract-price process is

$$\left\{ \prod_{i=1}^{n-2} (1 - \lambda_i L) \right\} \Delta x_t = \frac{\gamma}{s_0 s_{n-1}} E_t \left[\left(\prod_{i=1}^{n-2} (-\lambda_i) \right) \left(\sum_{k=0}^{\infty} \beta^k \lambda_1^k F^k \right) ..... \left(\sum_{k=0}^{\infty} \beta^k \lambda_{n-1}^k F^k \right) Z_t \right].$$

Letting

$$\delta(L) = \left\{ \prod_{i=1}^{n-2} (1 - \lambda_i L) \right\},$$

we obtain the solution for the rate of change of the new contract price as

$$\delta(L) \Delta x_t = \frac{\gamma}{s_0 s_{n-1}} \sum_{k=0}^{\infty} \kappa_k E_t Z_{t+k}.$$

Now let

$$\alpha(L) = \left(\sum_{j=0}^{n-1} s_j L^j \right)$$

be the operator that translates contract prices into the aggregate price level; aggregate price inflation can then be written as

$$\pi_t = \alpha(L) \Delta x_t.$$

Because $\delta_0 = 1$, this can be re-written as

$$\pi_t = \psi(L) \pi_{t-1} + \frac{\gamma \alpha(L)}{s_0 s_{n-1}} \left[\sum_{k=0}^{\infty} \kappa_k E_t Z_{t+k} \right].$$

Aggregate price inflation is therefore related to two factors. The first is current and past expectations about the future paths of the driving variable y_t. The second factor is inflation's own lagged values: As long as contracts are longer than two periods in length (*i.e.*, as long as $n > 2$), then inflation will be directly affected by its own lags.

Table 1: GMM Estimates of "Pure" New-Keynesian Phillips Curve

Driving variable	γ	β
1. Detrended GDP	-0.056**	1.039**
	(0.019)	(0.029)
2. Labor income share	0.005	1.001**
	(0.022)	(0.028)

Note: Standard errors in parentheses. **/*/^{a} denotes significant at 1/5/10 percent level, respectively. Estimation period is 1960:Q1 to 2004:Q3. Inflation is defined as the log-difference of the price index for nonfarm business output. Instrument set consists of two lags each of inflation, labor's share, an output gap, an interest rate spread (defined as the difference between ten-year and three-month Treasury yields), hourly compensation growth, and commodity price inflation (measured by the PPI for crude materials).

Table 2: Inflation Granger Causality Tests from Bivariate VAR Systems

	Number of lags in VAR system				
	1	2	3	4	5
A. GDP gap equation, (y_t, π_t) VAR					
Coefficient sum	-0.097^{**}	$\mathbf{-0.078^{**}}$	-0.062^{*}	-0.061^{*}	-0.050^{a}
	(0.022)	$\mathbf{(0.024)}$	(0.025)	(0.026)	(0.027)
Exclusion restriction	0.000	$\mathbf{0.006}$	0.050	0.099	0.049
B. Labor share equation, (s_t, π_t) VAR					
Coefficient sum	0.030	$\mathbf{0.028}$	0.025	0.025	0.024
	(0.023)	$\mathbf{(0.025)}$	(0.025)	(0.025)	(0.026)
Exclusion restriction	0.203	$\mathbf{0.474}$	0.380	0.500	0.605

Note: Table gives coefficient sums and exclusion restriction p-values for lags of inflation from the specified VAR equation. Boldface entries are estimates from models with lag lengths selected by the Schwartz criterion. Standard errors in parentheses. $^{**}/^{*}/^{a}$ denotes significant at 1/5/10 percent level, respectively.

Table 3: GMM Estimates for Closed Form of Hybrid Model

Instrument set	μ	δ_1	δ_2^{-1}	γ_b	γ_f	κ
1. Galí-Gertler (1999), two lags						
$N = 12$	0.065	0.757**	0.737**	0.486**	0.473**	0.042
	(0.060)	(0.037)	(0.265)	(0.065)	(0.110)	(0.044)
$N = 8$	0.061	0.729**	0.759**	0.469**	0.489**	0.039
	(0.045)	(0.044)	(0.163)	(0.046)	(0.072)	(0.032)
$N = 4$	0.024	0.549**	0.981**	0.357**	0.637**	0.016
	(0.018)	(0.085)	(0.030)	(0.036)	(0.038)	(0.012)
2. Galí, et al. (2005)						
$N = 12$	0.051	0.708**	0.860**	0.440**	0.535**	0.032
	(0.035)	(0.035)	(0.137)	(0.033)	(0.057)	(0.024)
$N = 8$	0.018	0.519**	0.971**	0.345**	0.646**	0.012
	(0.015)	(0.073)	(0.025)	(0.033)	(0.033)	(0.010)
$N = 4$	0.005	0.481**	1.007**	0.324**	0.678**	0.003
	(0.018)	(0.107)	(0.028)	(0.049)	(0.051)	(0.012)
3. Rudd-Whelan (2003), two lags						
$N = 12$	0.057	0.720**	0.821**	0.452**	0.516**	0.036
	(0.064)	(0.038)	(0.234)	(0.052)	(0.095)	(0.044)
$N = 8$	0.010	0.445**	0.983**	0.310**	0.684**	0.007
	(0.016)	(0.102)	(0.020)	(0.049)	(0.047)	(0.011)
$N = 4$	0.012	0.535**	1.032**	0.345**	0.665**	0.007
	(0.020)	(0.135)	(0.033)	(0.055)	(0.056)	(0.013)

Note: Standard errors in parentheses. **/*/[a] denotes significant at 1/5/10 percent level, respectively. See text for additional details.

Table 4: Lagged Inflation Coefficients from $\psi(L)$ Polynomial

	ψ_1	ψ_2	ψ_3	ψ_4
Taylor $(\beta = 0.99)$				
$n = 3$	-0.269			
$n = 4$	-0.430	-0.121		
$n = 5$	-0.533	-0.235	-0.068	
$n = 6$	-0.606	-0.328	-0.147	-0.044
Truncated Calvo $(\beta = 0.99, \alpha = 0.25)$				
$n = 3$	-0.263			
$n = 4$	-0.415	-0.113		
$n = 5$	-0.510	-0.216	-0.061	
$n = 6$	-0.574	-0.296	-0.127	-0.036
Truncated Calvo $(\beta = 0.99, \alpha = 0.5)$				
$n = 3$	-0.235			
$n = 4$	-0.355	-0.086		
$n = 5$	-0.419	-0.151	-0.037	
$n = 6$	-0.455	-0.192	-0.069	-0.017
Upward-Sloping Hazards $(\beta = 0.99)$				
$(0.1, 0.4, 1)$	-0.232			
$(0.1, 0.2, 0.7, 1)$	-0.342	-0.058		
$(0.1, 0.2, 0.3, 0.7, 1)$	-0.453	-0.153	-0.026	
$(0.1, 0.2, 0.3, 0.4, 0.8, 1)$	-0.507	-0.213	-0.065	-0.009

Notes: Refers to coefficients in equation (25) for various different hazard functions. n is the maximum possible contract length in the Taylor or truncated Calvo models, β is the discount rate, α is the probability of price changes in truncated Calvo model before period n. All equations obtained by log-linearizing around a zero-inflation steady state.
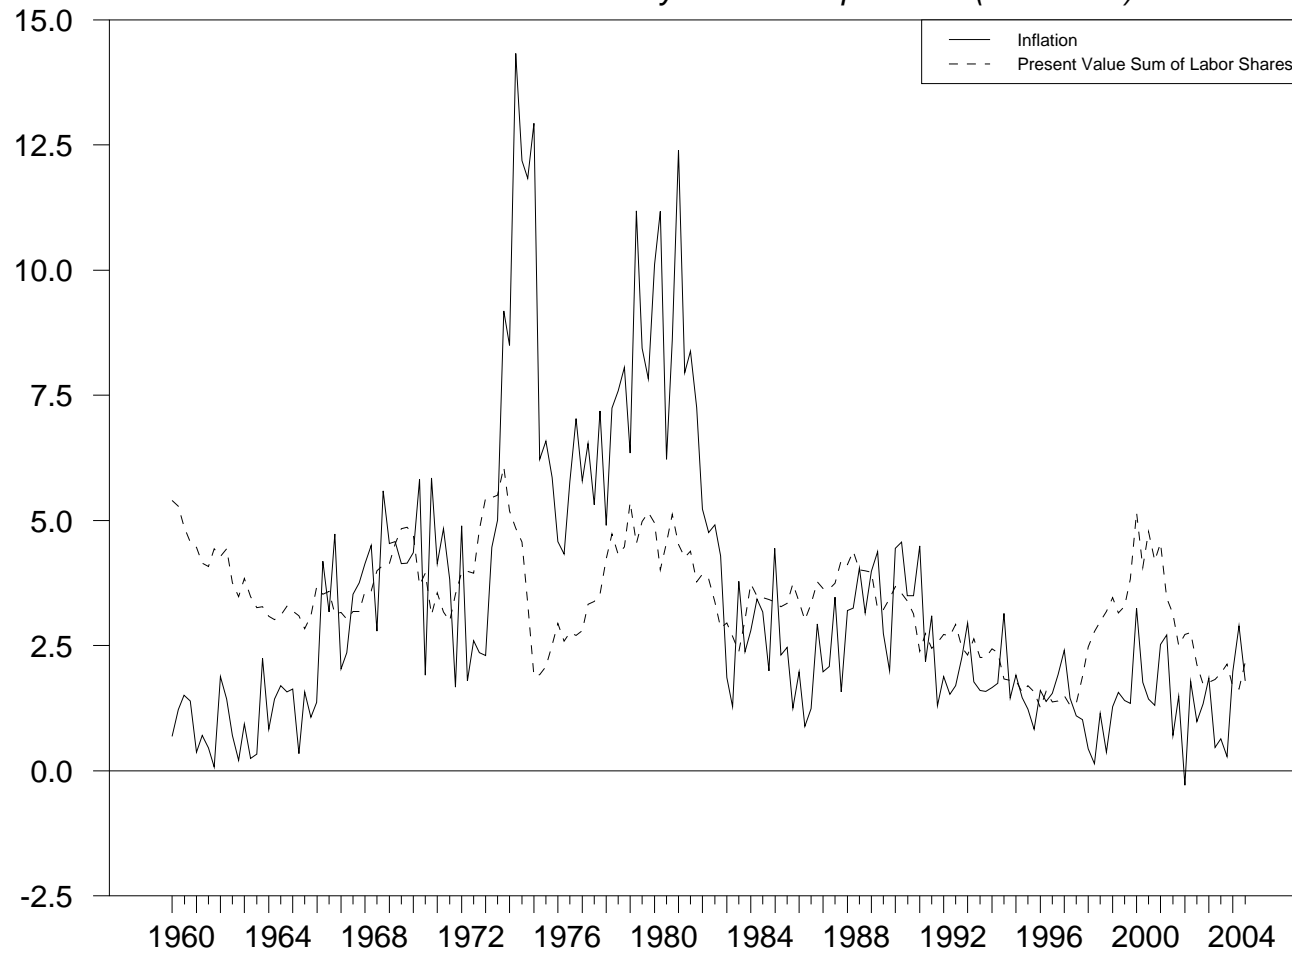
Figure 1

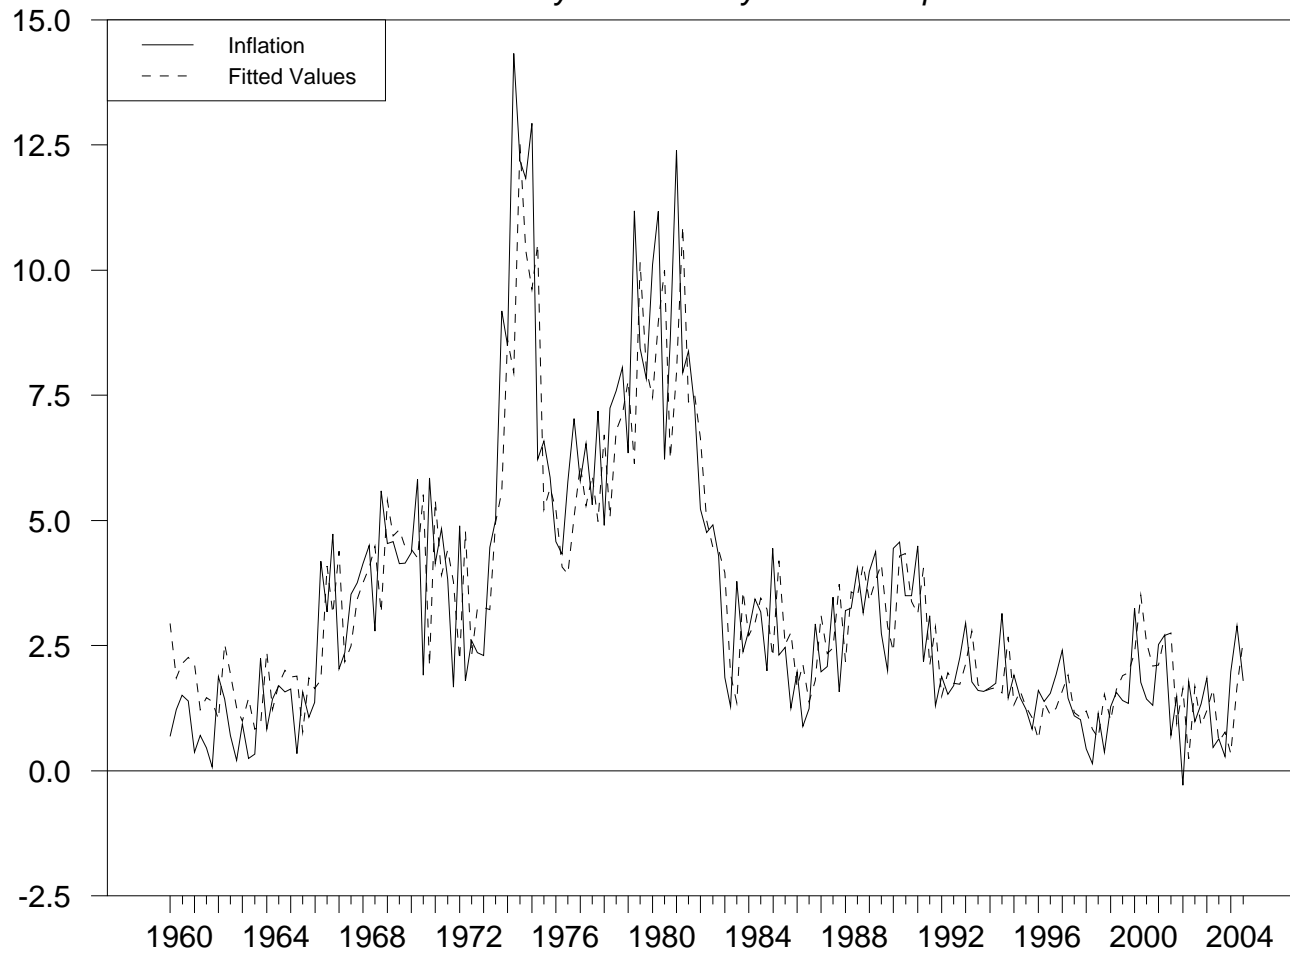*Two Potential Measures of Inflationary Pressures (NBER Recessions Shaded)*

Output Gap (Based on HP Filter)

Labor Share (Nonfarm Business Sector)
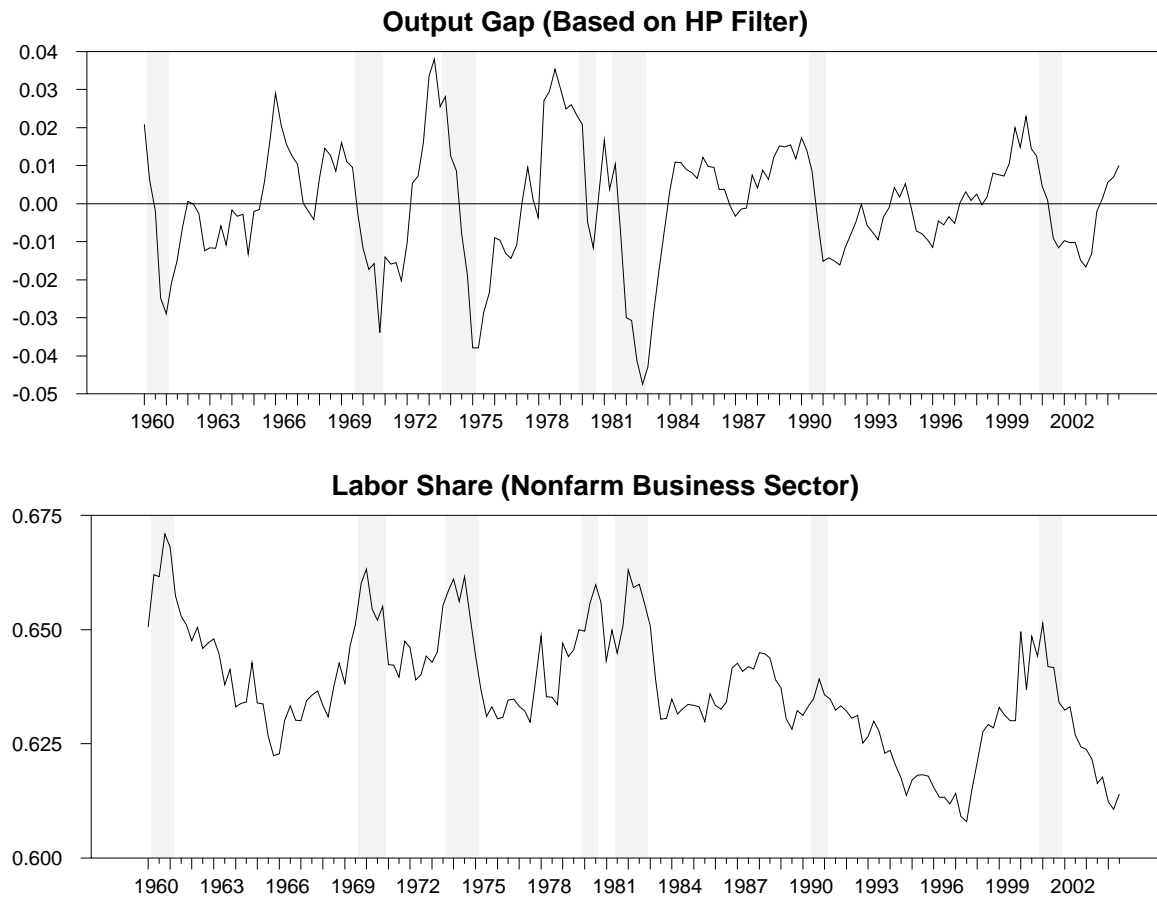
Figure 2

Fit of Labor-Share-Based New-Keynesian Phillips Curve (beta=0.99)

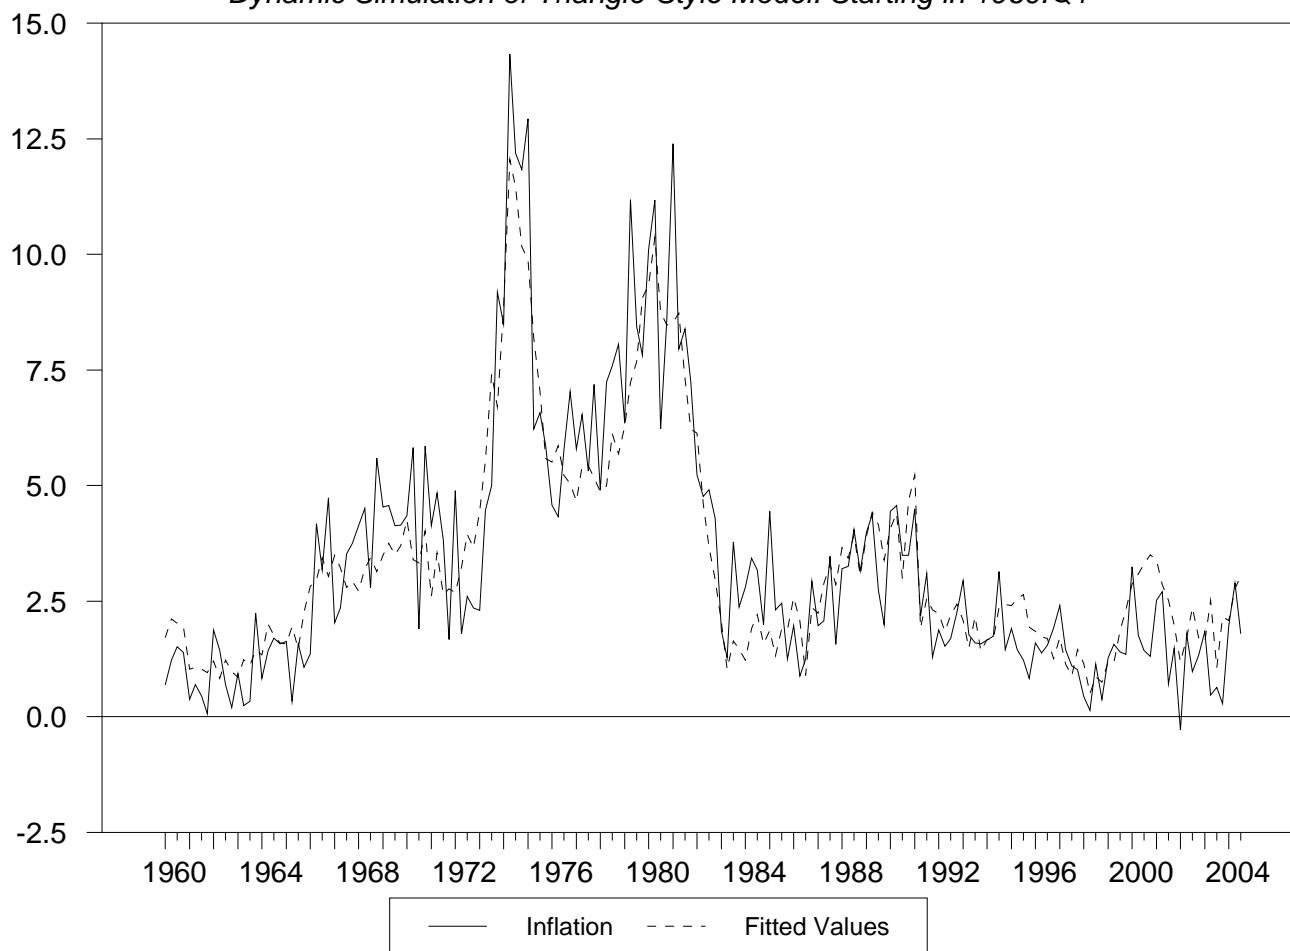# Figure 3

Fit of Labor Share Hybrid New-Keynesian Phillips Curve

Figure 4

Dynamic Simulation of Triangle-Style Model: Starting in 1960:Q1

Figure 5

Median Unbiased Estimates of Sum of AR Coefficients (With 90% Confidence Intervals)

NFB Price Inflation, Univariate Model

NFB Price Inflation, Triangle-Style Model