
JEREMY RUDD
KARL WHELAN

Modeling Inflation Dynamics: A Critical Review of Recent Research

In recent years, a broad academic consensus has arisen that favors using rational expectations sticky-price models to capture inflation dynamics. 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 models fail to provide a useful empirical description of the inflation process.

JEL code: E31

Keywords: Phillips curve, sticky prices, rational expectations.

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 research on inflation, with most of it related to the so-called “new-Keynesian” Phillips curve (NKPC), 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 this edition 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

A more complete version of this paper (Rudd and Whelan 2005c) was presented at the FRB/JMCB conference “Quantitative Evidence on Price Determination.” We are grateful to our conference discussants, Laurence Ball and Bennett McCallum, and to two anonymous referees for a number of helpful comments. The views expressed are our own and do not necessarily reflect the views of the Board of Governors, the staff of the Federal Reserve System, or the Central Bank of Ireland.

JEREMY RUDD is a Senior Economist in the Division of Research and Statistics, Board of Governors of the Federal Reserve System (E-mail: jeremy.b.rudd@frb.gov). KARL WHELAN is a Deputy Head in the Economic Analysis and Research Department, Central Bank and Financial Services Authority of Ireland (E-mail: karl.whelan@centralbank.ie).

Received October 27, 2005; and accepted in revised form September 1, 2006.

Journal of Money, Credit and Banking, Supplement to Vol. 39, No. 1 (February 2007)

© 2007 The Ohio State University

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 modeling 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 our survey will reflect the answers to these questions that we have proposed in earlier work. 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.

We start by briefly reviewing the traditional econometric Phillips curve that has played a central role in applied macroeconomics in recent decades, and contrast this with the modern NKPC based on rational expectations. The paper next provides an empirical assessment of the NKPC. 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.

We also review the evidence for the so-called “hybrid” class of new-Keynesian models, which allow partial dependence of inflation on its own lags. These models are often viewed as striking a compromise between the need for rigorous micro-foundations 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 (1999) 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 these findings, 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.

1. OLD AND NEW PHILLIPS CURVES

Modern thinking about the Phillips curve can be traced to the seminal contributions of Phelps (1967) and Friedman (1968), who argued that there was no simple trade-off

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.”

between inflation and unemployment. Instead, the correct formulation of the inflation–unemployment relation was an “expectations-augmented” Phillips curve of the form:

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

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). Empirical implementations of the Phelps–Friedman model usually assumed that inflation expectations were formed adaptively, as a weighted average of recent inflation rates:

$$\pi_t = \alpha + \gamma U_t + \sum_{i=1}^N \beta_i \pi_{t-i} + \epsilon_t, \quad (2)$$

where the weights β_i were constrained to sum to unity.

Taken literally, this model has important implications for the conduct of macroeconomic policy. To the extent that lagged inflation captures true inertia in the price-setting process, the model implies that rapid reductions in inflation require a substantial increase in unemployment. Hence, the model supports using a gradualist approach to effect large reductions in inflation. In addition, there will be a long lag between when macroeconomic shocks (including policy actions) occur, and when they have their full effect on inflation. These implications are well known and part of the conventional wisdom in policy circles.

The modern NKPC is a relationship of the form:

$$\pi_t = \beta E_t \pi_{t+1} + \lambda mc_t, \quad (3)$$

which relates inflation, π_t , to next period’s expected inflation rate and to real marginal cost, mc_t . As shown by Roberts (1995), this equation can be derived from a number of different models of price rigidity, although the recent literature has tended to focus on a derivation that assumes Calvo (1983) pricing. The model implies that, absent any pricing frictions, firms would set prices as a fixed markup over marginal cost; thus, inflationary pressures are generated by a high ratio of marginal cost to price. In addition, under relatively general conditions, aggregate real marginal cost is proportional to the gap between actual and potential output. With this assumption, the NKPC becomes

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

where y_t is the output gap.

Despite its superficial similarity to an expectations-augmented 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. Perhaps the most important implication 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.² Instead, inflation is determined in a completely forward-looking manner, as can be seen by solving forward equation (4) under the assumption of rational expectations:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k}. \quad (5)$$

The model also provides a very different interpretation of reduced-form Phillips curve regressions: Lagged inflation terms enter econometric Phillips curves merely because they proxy for expectations of future values of the output gap. As the strength of this statistical correlation is likely to vary across monetary policy regimes, the NKPC also implies that reduced-form Phillips curves are likely to be subject to the Lucas critique.

2. EVIDENCE ON THE NEW-KEYNESIAN PHILLIPS CURVE

The NKPC's implications for policy and forecasting represent a dramatic departure from the conventional wisdom suggested by econometric Phillips curves. 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, the empirical case against the NKPC turns out to be quite strong.

2.1 Output Gap Models

Consider first the version of the model that is suggested by equation (4), 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, with inflation defined as the log-difference of the price index for nonfarm business sector output, and with the output gap constructed by detrending the log of real GDP with a quadratic polynomial in calendar time.³ As is well established, the results from this type of exercise are problematic for this variant of the model: The coefficient on the output gap is highly statistically significant but *negative*.⁴ One way to understand the failure of the output gap version of the NKPC is

2. If current and expected values of the output gap move sluggishly over time, inflation will still be autocorrelated (see below). However, it is important to distinguish between this and intrinsic inertia.

3. The instrument set includes inflation, the labor income share, the output gap, the 10-year and 3-month Treasury yield spread, 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 instead of four because the larger set fails Stock and Yogo's (2003) test for weak instruments. Also, Galí and Gertler used the Commodity Research Bureau's spot price index as their commodity price series; we do not because it has a number of discontinuities and narrower coverage.

4. See Fuhrer and Moore (1995) for an early discussion of the failure of NKPC-like models based on detrended output. In addition, Mankiw (2001) critiques these models on the grounds that they fail to generate reasonable responses to monetary policy shocks.

TABLE 1
GMM ESTIMATES OF "PURE" NEW-KEYNESIAN PHILLIPS CURVE

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

NOTES: Standard errors in parentheses. ** or * denotes significant at 1 or 5 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 10-year and 3-month Treasury yields), hourly compensation growth, and commodity price inflation (measured by the PPI for crude materials).

in terms of the model's closed-form solution (equation (5)), 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: Regressions of detrended GDP on its own lags and lags of inflation provide statistically significant evidence that higher inflation leads future *reductions* in the output gap.

2.2 Labor Share Proxies: Theoretical Issues

These poor results suggest two possible interpretations: Either the rational expectations NKPC provides a bad description of inflation, or this particular measure of the output gap is flawed. The latter explanation has proved to be 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 many kinds of shocks, and so could fluctuate significantly (and stochastically) from 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.

As is evident from equation (3), the theoretical models that underpin the NKPC predict that real marginal cost drives inflation. Hence, recent implementations of the NKPC 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 wL/Y (where w is hourly compensation, L is total hours, and Y is real output). The resulting proxy for *real* marginal cost, $(wL)/(pY)$, is therefore labor's share of income. In this section, we 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.

It is worth first briefly reviewing the theoretical case for the relationship underlying equation (4), which allows us to write the NKPC in terms of the gap between actual and potential output. As Woodford (2003) has argued, increases in output that are not

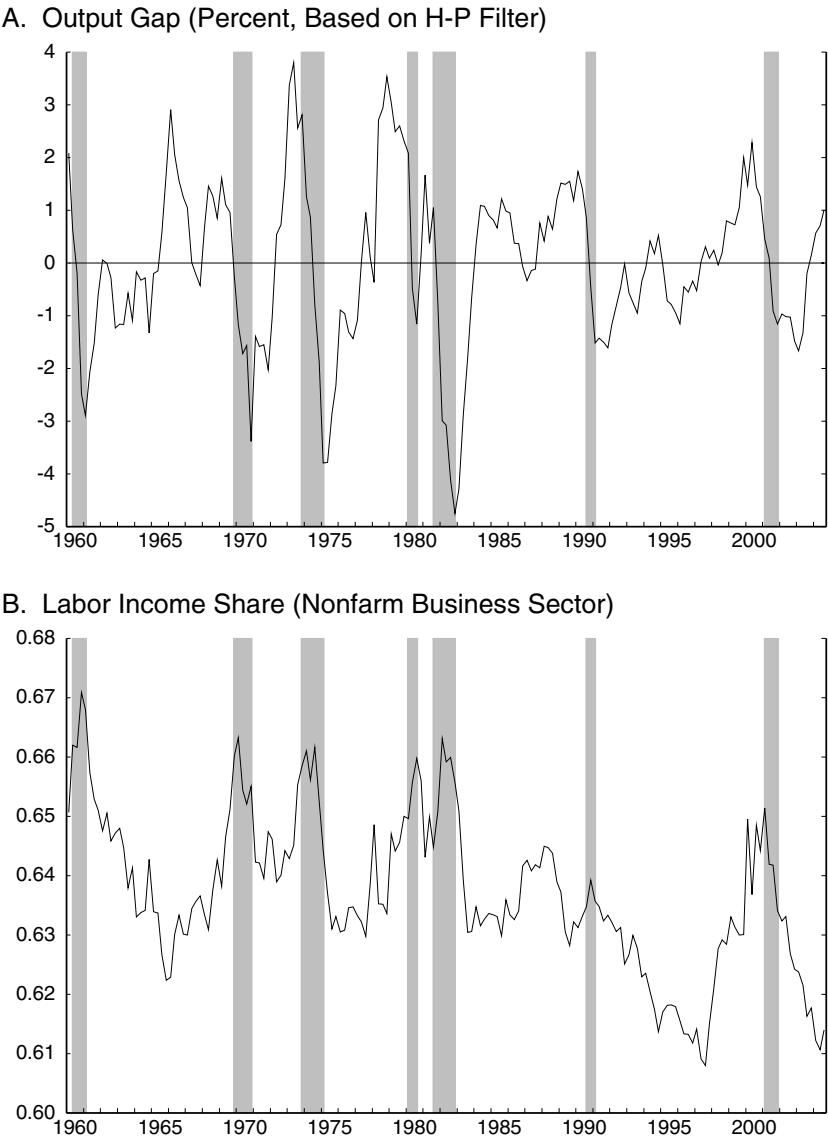


FIG. 1. Two Potential Measures of Inflationary Pressures (NBER Recession Dates Shaded).

driven by increases in technological efficiency will tend to raise nominal marginal costs by more than prices in a broad class of models, as workers require higher real wages in order to supply more hours. In light of the theoretical case for procyclicality, an examination of Figure 1 reveals a potential problem with using labor's share as a proxy for real marginal cost. The figure plots detrended GDP and the labor income 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.

The negative correlation between labor's share and traditional output gap measures has been noted by those who advocate using labor's share to measure mc_t , and is often cited as evidence that traditional GDP gaps might be negatively correlated with the "true" gap between actual and potential output. 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 the figure—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. These are generally viewed as times when output was below potential, with this assessment 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 latter measure is consistent with the NBER concept, in that this gap manifests substantial peak-to-trough declines in each recession. Furthermore, it should be remembered that this measure of the output gap works well as an explanatory variable in traditional inflation regressions: 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. This suggests that the NKPC itself, rather than the traditional output gap, may be the problem.

In contrast to the HP-filtered output gap, the labor share tends to jump up to a local peak near the onset of a recession. 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 actual and potential output, we would have to conclude from the behavior of the labor share that output was actually *above* potential during each postwar recession. While this interpretation is at least theoretically possible, and might find support among those who view adverse technology shocks as the principal cause of recessions, we see the evidence to the contrary—as outlined, for instance, by Galí (1999)—as more compelling: Technology shocks do not appear to be the dominant source of business cycles.

The discrepancy between the observed behavior of the labor share and the theoretical prediction that real marginal cost should be procyclical most likely arises because average unit labor costs are a poor proxy for nominal marginal costs. Indeed, as Rotemberg and Woodford (1999) discuss in detail, a number of factors—such as the existence of overtime premia, adjustment costs, or overhead labor—could lead marginal and average cost to manifest different cyclical patterns. It seems likely, therefore, that real marginal cost is procyclical and is related to the gap between actual and potential output, but that the labor share proxy is simply not capturing these cyclical movements correctly.

The important measurement issues that surround the use of average unit labor cost as a proxy for marginal cost have generally been ignored in the recent literature on

the NKPC. Indeed, many proponents of the labor share approach refer to labor's share as "real marginal cost," as if these two concepts were identical. Occasionally, 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 (which they argue is already good when labor's share is used) can only be improved with a more sophisticated measure of real marginal cost. This line of reasoning seems suspect, however, since a good proxy for mc_t is likely to have a cyclical pattern that is the opposite of that displayed by labor's share. Thus, if the labor share version of the model were to fit well, it would be unlikely that a model based on a marginal cost proxy manifesting completely different cyclical behavior would also do so.

2.3 Empirical Performance of the Labor Share Model

Despite our theoretical reservations, any assessment of the labor's share version of the NKPC must ultimately rest upon the model's ability to match the data. If this variant of the NKPC works well empirically, then perhaps the issues just discussed are not critical. One result that is often cited in favor of the model is Galí and Gertler's (1999) finding that GMM estimation of

$$\pi_t = \beta E_t \pi_{t+1} + \gamma s_t, \quad (6)$$

where s_t denotes the log of labor's share of income, yields 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 for our updated sample.⁵ Importantly, however, the results also indicate that a key finding 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. 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.⁶

A more concrete way to illustrate the weakness of this model is to construct explicit measures of $E_t s_{t+k}$ with a forecasting model. For instance, one can fit the equation

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t s_{t+k}, \quad (7)$$

using a vector autoregression (VAR) to forecast the values of s_{t+k} . However, the resulting discounted sum of labor shares delivers a very poor empirical model of inflation.

5. Because the labor share is probably a poor proxy for real marginal cost (and thus inflationary pressures), a discussant (Laurence Ball) suggested that this positive sign likely reflects two offsetting errors: the NKPC's failure to correctly predict the lead-lag relationship between inflation and the output gap, and the labor share's failure to co-move positively with the correct gap concept.

6. Revisions to the labor share data appear to be driving this result: Using Galí and Gertler's original data and the *current* vintage of labor share data also overturns the finding of a significant γ .

Consider, for instance, 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. Both equations in the VAR fit well; in particular, detrended GDP enters as a highly significant predictor of future labor shares. However, the discounted sum of expected labor shares (generated using a value of β equal to 0.99) is almost uncorrelated with inflation: This version of the NKPC receives an R^2 of only 0.13. In interpreting this result it should be kept in mind that the NKPC is a structural model and should be expected to fit worse than a reduced-form econometric model. However, since a simple $AR(4)$ model of inflation receives an \bar{R}^2 of 0.75, it is fair to say that the NKPC clearly fails to provide a useful interpretation of this important reduced-form fact.

2.4 Granger Causality

The AR model's ability to provide 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 practice, Rudd and Whelan (2005a) show that VARs incorporating inflation itself usually provide a somewhat improved fit (though still well short of that obtained by reduced-form models). However, for such a forecasting approach to be legitimate, there must be evidence that inflation helps to forecast future values of the labor share. It turns out, however, that inflation does not Granger cause the labor share (this result also holds in the updated data set used here and is robust across a wide range of specifications). In itself, this finding suggests that the essential story behind this model—that inflation is determined by expectations of future labor shares—is not at all evident in the data.

3. "HYBRID" NEW-KEYNESIAN MODELS

The debate over the adequacy of the NKPC is ongoing. However, even some of the NKPC's more enthusiastic supporters 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

$$\pi_t = \gamma_f E_t \pi_{t+1} + \gamma_b \pi_{t-1} + \kappa x_t. \quad (8)$$

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 directly addressing a well-known empirical deficiency of the pure forward-looking model.⁷ Despite these models' relatively weak microfoundations, there is now a

7. There are at least three ways to motivate the hybrid model theoretically. First, Fuhrer and Moore (1995) assume an alternative contracting specification in which workers bargain over relative real wages. Second, Christiano, Eichenbaum, and Evans (2005) allow all prices to change each period, though only a fraction can do so optimally (the rest are indexed to past inflation). Finally, Galí and Gertler (1999) assume the existence of a group of "rule-of-thumb" price setters. In each case, the resulting models are arguably more *ad hoc* than microfounded.

near-consensus among new-Keynesian economists that a hybrid model featuring a large γ_f and a small γ_b provides a sensible and empirically plausible theoretical framework for analyzing the behavior of inflation. Here, we offer some counterarguments to this view.

3.1 Direct GMM Estimation

As before, one can estimate the hybrid specification (8) directly with GMM using variables dated time t or earlier as instruments. In this case, because equation (8) 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.

However, such estimates do not really allow us to distinguish between forward- and backward-looking models of inflation: 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 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 likely that some of these omitted variables 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.

In this connection, Rudd and Whelan (2005b) show that Galí and Gertler's estimates are also consistent with the true model's being a completely backward-looking formulation that contains some of the variables employed in (or correlated with) the instrument set. Galí, Gertler, and López-Salido (2005) argue that this point is based on an “extreme scenario” in which there is no forward-looking behavior in price setting. In actuality, though, the extreme nature of this scenario helps to make a more general point: To the extent that this procedure can yield large estimates of γ_f even when forward-looking behavior is completely absent, such estimates can hardly be viewed as compelling evidence that such behavior is dominant.

3.2 Closed-Form Solutions

Given that direct GMM estimation of equation (8) sheds little light on the relative importance of the forward- and backward-looking terms, it is useful to consider other ways to examine this question. Fortunately, one can go beyond the simple hybrid equation in order to test for the presence of rational forward-looking behavior. This is because the assumption of rational expectations 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,

$$\pi_t = \delta_1\pi_{t-1} + \mu \sum_{k=0}^{\infty} \delta_2^{-k} E_t s_{t+k}, \quad (9)$$

where δ_1 and δ_2 are the roots of the polynomial equation

$$\gamma_f z^2 - z + \gamma_b = 0. \quad (10)$$

Equation (9) provides a better way to distinguish between the hybrid model and a traditional Phillips curve. Both models allow lagged inflation to matter. However, while standard Phillips curves use a conventional output gap measure to capture inflationary pressures, the hybrid model considered here sees the expected discounted sum of s_{t+k} as the relevant driving term. Hence, the key question is whether this term is an important determinant of empirical inflation dynamics. The evidence suggests not.

One way to reveal both the strengths and weaknesses of the labor share version of the hybrid model is to construct proxies for $E_t s_{t+k}$ using a VAR. To do this, we again used a two-lag VAR in labor's share and H-P filtered GDP. Consider the predicted values of π_t that obtain from the best-fitting parameterization of the equation (where this is determined by searching over a grid of values for the parameter δ_2^{-1} ranging from zero to one). The resulting model does a relatively good job fitting the inflation process: Its R^2 is 0.713, which is considerably higher than what the pure forward-looking NKPC receives. Indeed, Galí and Gertler (1999) also reported obtaining a good fit from this type of exercise (this is the message of the "fundamental inflation" series shown in Figure 2 of their paper), and presented it as an important endorsement of the hybrid model. The picture is gloomier, however, when one focuses on the question we are examining: Almost none of the model's ability to fit the data comes from including expectations of future labor shares. 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. So, the fit of this model can hardly be attributed to the contribution made by expected future values of s_t .

To be more precise about whether these terms play any role in explaining inflation, we also applied a GMM estimation methodology to the closed-form solution of the hybrid model, equation (9). To deal with the infinite sum in equation (9), we can

TABLE 2
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 (0.060))	(0.757** (0.037))	(0.737** (0.265))	(0.486** (0.065))	(0.473** (0.110))	(0.042 (0.044))
$N = 8$	(0.061 (0.045))	(0.729** (0.044))	(0.759** (0.163))	(0.469** (0.046))	(0.489** (0.072))	(0.039 (0.032))
$N = 4$	(0.024 (0.018))	(0.549** (0.085))	(0.981** (0.030))	(0.357** (0.036))	(0.637** (0.038))	(0.016 (0.012))
2. Galí, Gertler, and López-Salido (2005)						
$N = 12$	(0.051 (0.035))	(0.708** (0.035))	(0.860** (0.137))	(0.440** (0.033))	(0.535** (0.057))	(0.032 (0.024))
$N = 8$	(0.018 (0.015))	(0.519** (0.073))	(0.971** (0.025))	(0.345** (0.033))	(0.646** (0.033))	(0.012 (0.010))
$N = 4$	(0.005 (0.018))	(0.481** (0.107))	(1.007** (0.028))	(0.324** (0.049))	(0.678** (0.051))	(0.003 (0.012))
3. Rudd and Whelan (2006), two lags						
$N = 12$	(0.057 (0.064))	(0.720** (0.038))	(0.821** (0.234))	(0.452** (0.052))	(0.516** (0.095))	(0.036 (0.044))
$N = 8$	(0.010 (0.016))	(0.445** (0.102))	(0.983** (0.020))	(0.310** (0.049))	(0.684** (0.047))	(0.007 (0.011))
$N = 4$	(0.012 (0.020))	(0.535** (0.135))	(1.032** (0.033))	(0.345** (0.055))	(0.665** (0.056))	(0.007 (0.013))

NOTES: Standard errors in parentheses. ** or * denotes significant at 1 or 5 percent level, respectively. See text for additional details.

partially solve forward equation (8) to obtain an *exact* expression in a finite number of terms,

$$\pi_t = \delta_1 \pi_{t-1} + \mu \sum_{k=0}^N \delta_2^{-k} E_t s_{t+k} + E_t [\delta_2^{-(N+1)} (\pi_{t+N+1} - \delta_1 \pi_{t+N})], \quad (11)$$

which can be directly estimated with GMM. Table 2 reports the estimated parameters obtained for three different values of N (four, eight, and twelve quarters) and three different instrument sets.⁸ In every case, the estimated μ coefficient, which describes the influence of the 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 equation. Similarly, while the labor's share version of the hybrid model can fit actual

8. The first instrument set is described in note 3. The second instrument set is the one used by Galí, Gertler, and López-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 (2006), and is the same as the Galí, Gertler, and López-Salido set but with two lags of inflation instead of four.

inflation reasonably well, this 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.

3.3 Comparison with Galí, Gertler, and López-Salido (2005)

Of course, there is always more than one way to interpret any set of empirical results. Galí, Gertler, and López-Salido (2005) estimate equations similar to these, and arrive at conclusions that apparently contradict our own. For instance, these authors conclude that backward-looking behavior, “while statistically significant, is 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 discrepancy 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, Gertler, and López-Salido's (2005) employing 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 assert 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 δ_1 and δ_2 parameters and the implied values of γ_f and γ_b is very complex (it occurs through the second-order polynomial equation (10)). Moreover, it turns out that the parameters that Galí, Gertler, and López-Salido 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 that this is the case, consider Table 2 again. 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.⁹ To understand how this pattern occurs, consider the case in which $\gamma_f + \gamma_b = 1$, a restriction that is directly imposed by several popular hybrid models, and one that conforms closely to the estimates reported in Table 2. In this case, if $\gamma_f \leq 0.5$, then the closed-form solution is given by

$$\pi_t = \pi_{t-1} + \mu \sum_{k=0}^{\infty} \left(\frac{\gamma_f}{1 - \gamma_f} \right)^k E_t s_{t+k}, \quad (12)$$

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

$$\pi_t = \left(\frac{1 - \gamma_f}{\gamma_f} \right) \pi_{t-1} + \mu \sum_{k=0}^{\infty} E_t s_{t+k}. \quad (13)$$

9. For completeness, the table also reports the values of κ (the coefficient on the driving variable in the hybrid inflation equation (8)) 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 also statistically indistinguishable from zero.

Now consider the properties of the estimates of γ_f that will be obtained from the closed-form solution. First, 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), Galí, Gertler, and López-Salido's (2005) procedure will estimate γ_f to be greater than one-half. 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, Galí, Gertler, and López-Salido's metric will still imply that rational forward-looking behavior is “dominant.” And, even when the true model is completely backward looking, this procedure will imply the opposite.

For these reasons, the estimates of γ_f and γ_b that are generated by this method are not useful for assessing the importance of the forward-looking component of the hybrid model. In the absence of a 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. Importantly, this critique is not restricted to GMM-based estimates, but also applies to so-called maximum-likelihood estimation of rational-expectations pricing models, in which the “structural parameter values” implied by model-consistent solutions such as these closed-form equations are computed.¹⁰

Beyond estimation issues, the substantial focus on point estimates of γ_f and γ_b obtained from hybrid inflation equations has, in our opinion, led many researchers to misapprehend how badly the pure NKPC model fits the data, and its failure to explain the 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 seen as striking a compromise between the standard econometric 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 traditional model.

In addition, even if one wishes to accept Galí, Gertler, and López-Salido's (2005) parameter estimates (which imply 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 (9), 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.

4. 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

10. See Fuhrer (1997) for an early example.

knew that the world was flat, knew that the sun revolved around the earth, and knew that there was an exploitable long-run trade-off between inflation and unemployment. Today, one can meet many researchers who know that the NKPC provides a good model of the inflation process once one uses a suitable proxy for real marginal cost and who know that forward-looking rational behavior dominates price setting. We hope the preceding discussion will give at least some interested researchers cause to check these conclusions more fully against the available evidence.

We began with two fundamental questions about inflation: first, what is a suitable measure of inflationary pressures; and second, how are inflation expectations formed. Regarding the former, there appears to be little evidence that labor's share is a suitable driving variable for inflation. Regarding expectations formation, our results contradict standard rational expectations models based on either traditional output gaps or labor's share. Of course, this does 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 modeling of inflation in a rational expectations framework provides an adequate description of the empirical inflation process.

LITERATURE CITED

- Calvo, Guillermo A. (1983) "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, 12, 383–98.
- 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.
- Fuhrer, Jeffrey. (1997) "The (Un)Importance of Forward-Looking Behavior in Price Specifications." *Journal of Money, Credit and Banking*, 29, 338–50.
- Fuhrer, Jeffrey, and George Moore. (1995) "Inflation Persistence." *Quarterly Journal of Economics*, 110, 127–59.
- Friedman, Milton. (1968) "The Role of Monetary Policy." *American Economic Review*, 58, 1–17.
- Galí, Jordi. (1999) "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–71.
- Galí, Jordi, and Mark Gertler. (1999) "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, 44, 195–222.
- Galí, Jordi, Mark Gertler, and David López-Salido. (2005) "Robustness of Estimates of the Hybrid New Keynesian Phillips Curve." *Journal of Monetary Economics*, 52, 1107–18.
- Mankiw, N. Gregory. (2001) "The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment." *Economic Journal*, 111, 45–61.
- Phelps, Edmund. (1967) "Phillips Curves, Expectations of Inflation, and Optimal Inflation Over Time." *Economica*, 135, 254–81.
- Roberts, John M. (1995) "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking*, 27, 975–84.

- 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, pp. 1051–1135. Amsterdam: North-Holland.
- Rudd, Jeremy, and Karl Whelan. (2005a) "Does Labor's Share Drive Inflation?" *Journal of Money, Credit and Banking*, 37, 297–312.
- Rudd, Jeremy, and Karl Whelan. (2005b) "New Tests of the New-Keynesian Phillips Curve." *Journal of Monetary Economics*, 52, 1167–81.
- Rudd, Jeremy, and Karl Whelan. (2005c) "Modelling Inflation Dynamics: A Critical Review of Recent Research." Federal Reserve Board Finance and Economics Discussion Series No. 2005–66.
- Rudd, Jeremy, and Karl Whelan. (2006) "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" *American Economic Review*, 96, 303–20.
- Sbordone, Argia M. (2002) "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics*, 49, 265–92.
- 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, pp. 1–17. New York: New York University.
- Solow, Robert. (1976) "Down the Phillips Curve with Gun and Camera." In *Inflation, Trade and Taxes: Essays in Honor of Alice Bourneuf*, edited by David Belsey, pp. 3–22. Columbus: Ohio State University Press.
- Stock, James, and Motohiro Yogo. (2003) "Testing for Weak Instruments in Linear IV Regression." Mimeo, Harvard University (February draft).
- Woodford, Michael. (2003) *Interest and Prices*. Princeton, NJ: Princeton University Press.

Copyright of *Journal of Money, Credit & Banking* (Blackwell) is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.