

Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?

Jeremy Rudd
Federal Reserve Board*

Karl Whelan
Central Bank of Ireland**

Second Draft
January 6, 2003

Abstract

The new-Keynesian Phillips curve is often criticized on the grounds that it fails to account for the dependence of inflation on its own lags. In response, a large number of recent studies have employed a so-called “hybrid” sticky-price specification in which inflation depends on a weighted average of lagged and expected future values of itself, in addition to a driving variable such as the output gap. In this paper, we consider some simple tests of the hybrid model, which exploit the model’s implication that the *change in inflation* depends on a geometrically discounted sum of current and expected future values of the driving term. Our results suggest that the hybrid model provides a poor description of empirical inflation dynamics, and that there is little evidence of the type of forward-looking behavior implied by the model. This conclusion is robust to the use of both detrended GDP and the labor share of income as the driving variable.

*Corresponding author. Mailing address: Mail Stop 80, 20th and C Streets NW, Washington, DC 20551. E-mail: jeremy.b.rudd@frb.gov.

**E-mail: karl.whelan@centralbank.ie. 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.

1 Introduction

In recent years, macroeconomics has moved toward analyzing business cycles and stabilization policy in the context of models that incorporate both nominal rigidities and rational optimizing agents.¹ One important way in which this “new-Keynesian” approach differs from earlier work in the Keynesian tradition involves the way in which expectations are assumed to affect price-setting behavior. In particular, rather than assuming adaptive expectations on the part of wage- and price-setters, recent work draws on the sticky-price models of Rotemberg (1982) and Calvo (1983) in order to motivate a forward-looking inflation equation (a “new-Keynesian Phillips curve”) of the form

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

where β is a parameter between zero and one, and y_t is a measure of the output gap. Under the assumption of rational expectations, this model yields the following closed-form expression for inflation:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k}, \quad (2)$$

which has the interpretation that current inflation is completely determined by price-setters’ expectations of future output gaps.

An important implication of this model is that inflation should be independent of its own lagged values. As a result, this specification is often criticized on the grounds that it cannot account for the important role played by lagged inflation in empirical inflation equations. In response to this critique, therefore, several researchers have suggested an alternative to the pure forward-looking model that is intended to better capture observed inflation inertia. This “hybrid” specification modifies (1) such that inflation depends on a weighted sum of its lag and its (rationally) expected future value,

$$\pi_t = (1 - \theta)\pi_{t-1} + \theta E_t \pi_{t+1} + \gamma y_t, \quad (3)$$

with the weights constrained to sum to unity in order to preclude the existence of a long-run level tradeoff between inflation and real activity.²

¹See Clarida, Galí, and Gertler (1999) for a survey of much of this work, and Woodford (2002) for a detailed treatment.

²Examples of studies that use this pricing equation include Rudebusch (2002), Ehrmann and Smets (2001), and Casares and McCallum (2000).

This specification has been motivated in a variety of ways. For example, Fuhrer and Moore (1995) employed an assumption that workers bargain over relative real wages in order to obtain a model with $\theta = 0.5$, while Roberts (1997) has argued that an equation such as (3) could result in an economy whose population was comprised of set fractions of rational and rule-of-thumb price setters. More recently, Christiano, Eichenbaum, and Evans (2001) have explicitly derived a specification similar to (3) using a variant of the Calvo model in which those firms that are unable to reoptimize their price instead index it to last period’s inflation rate.

In this paper, we assess whether the hybrid model provides a good empirical characterization of U.S. inflation behavior. The tests of the model that we consider are based on the observation that the hybrid specification (3) implies an expression for the *change* in inflation of the form

$$\Delta\pi_t = \frac{\gamma}{1-\theta} \sum_{k=0}^{\infty} \left(\frac{\theta}{1-\theta}\right)^k E_t y_{t+k}. \quad (4)$$

We focus on this prediction of the model, rather than on the model’s implications for the *level* of inflation, in order to derive tests that are capable of distinguishing the hybrid model from reasonable alternatives. In practice, inflation is well explained by its own lag; hence, incorporating lagged inflation into the inflation equation should allow the hybrid model to fit the level of inflation relatively well. However, such fit could be obtained by any model that features an important role for lagged inflation—including models that rely on non-rational, backward-looking expectations. In contrast, the hybrid model’s predictions for the evolution of $\Delta\pi_t$ are quite clear-cut, and allow us to precisely distinguish this model from a traditional backward-looking specification.

We consider two different methods for assessing whether equation (4) provides a good empirical description of the inflation process. The first employs the well-known methodology of Campbell and Shiller (1987), which entails estimating a VAR for the driving process y_t and using it to forecast the future values of this variable. The second method involves estimating the equation using GMM. Both methods turn out to yield useful insights—the first into the predicted time-series properties of $\Delta\pi_t$ that are implied by the model, and the second into the statistical significance of the model’s forward-looking component.

Our results suggest that the hybrid model provides a very poor description of

empirical inflation dynamics. Specifically, we find that the empirical process for the change in inflation appears to bear very little resemblance to the expected discounted sum of current and future y_t values. Importantly, this conclusion holds both when we use detrended output as y_t , as well as when we use labor’s share of income (real unit labor costs), as has been suggested by Galí and Gertler (1999). In particular, even when the labor income share is used as the driving variable, we find that γ is not significantly different from zero, implying that the change in inflation is unrelated to the expectation of future labor shares and indicating that the type of rational forward-looking behavior hypothesized by the hybrid model is absent from the data.

The contents of the paper are as follows. Section 2 derives the present-value formulations of the new-Keynesian Phillips curve and the hybrid model and discusses how these models can be tested using VAR-based proxies for expectations. Section 3 uses this VAR-based test framework to confirm the poor performance of the pure new-Keynesian Phillips curve, while Section 4 assesses the hybrid model. Section 5 presents our GMM estimates of the hybrid model, and also considers whether the performance of this model can be improved by incorporating a more complex “rule-of-thumb” for backward-looking agents. Section 6 concludes.

2 Present-Value Formulations

In this section, we derive simple present-value based representations for both the pure new-Keynesian Phillips curve and the hybrid model alternative; we then discuss how the VAR-based techniques of Campbell and Shiller (1987) can be used to assess these models.

Begin by considering the new-Keynesian Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} + \gamma y_t.$$

Any empirical procedure that aims to assess the fit of this model must specify how the expectational term $E_t \pi_{t+1}$ is determined. The well-known approach of Campbell and Shiller (1987) assesses first-order stochastic difference equations of this type by first performing repeated substitutions to arrive at

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

and then using an econometric model to forecast all future values of y_t . Specifically, if we define y_t as the first variable in a multivariate VAR of the form

$$Z_t = AZ_{t-1} + \epsilon_t, \quad (5)$$

then we can express the discounted sum of current and future values of y_t as

$$\sum_{k=0}^{\infty} \beta^k E_t y_{t+k} = e'_1 (I - \beta A)^{-1} Z_t, \quad (6)$$

where e'_1 denotes a vector with one in the first row and zeroes elsewhere.³

Hence, one strategy for assessing the empirical performance of the new-Keynesian Phillips curve involves comparing π_t with the discounted sum of current and expected future y_t values that we obtain from a VAR such as (5), with a point estimate of γ in turn obtained by regressing π_t on the matrix expression in (6).⁴ Implementation of this method also requires us to have an estimate of β . The theory that underlies the new-Keynesian Phillips curve implies that this parameter is the rate at which firms discount future profits. In the calculations that we present in the next sub-section, we follow Woodford (2001) and set $\beta = 0.99$.

The hybrid sticky-price model of inflation, equation (3), can also be written as a first-order stochastic difference equation. To see this, note that this model implies that the first difference of inflation can be expressed as

$$\Delta\pi_t = \theta (E_t \pi_{t+1} - \pi_{t-1}) + \gamma y_t. \quad (7)$$

A couple of simple substitutions then yield

$$\Delta\pi_t = \frac{\theta}{1-\theta} E_t \Delta\pi_{t+1} + \frac{\gamma}{1-\theta} y_t. \quad (8)$$

So, by the same reasoning as before, the hybrid model implies that the *change* in inflation should equal a discounted sum of current and expected future values of y_t , with the discount rate in this case being $\frac{\theta}{1-\theta}$:

$$\Delta\pi_t = \frac{\gamma}{1-\theta} \sum_{k=0}^{\infty} \left(\frac{\theta}{1-\theta} \right)^k E_t y_{t+k}. \quad (9)$$

³This formula relies on the fact that $E_t Z_{t+k} = A^k Z_t$, and makes use of a matrix version of the standard geometric sum formula. See Sargent (1987, pp. 311-312) for more details.

⁴The reported standard error for γ will not be valid because the discounted sum is a generated regressor. We return to this issue in Section 5.

Again using a VAR like (5) to generate expectations of y_t yields

$$\Delta\pi_t = \frac{\gamma}{1-\theta} e'_1 \left(I - \frac{\theta}{1-\theta} A \right)^{-1} Z_t. \quad (10)$$

One new complication that arises when assessing this model empirically is that, unlike the case with the pure new-Keynesian Phillips curve, theory does not provide *a priori* guidance as to the appropriate value of θ . So, for the estimates reported in Section 4, we choose θ from a grid search aimed at obtaining the best-fitting hybrid model.⁵

3 The New-Keynesian Phillips Curve

Before examining the hybrid model, it will be useful to present some evidence illustrating exactly how the pure forward-looking new-Keynesian Phillips curve fails to match the empirical properties of inflation.

We will consider two versions of the model. The first version equates y_t with a traditional output gap measure, defined here as the deviation of log real nonfarm GDP from a quadratic trend. The second version follows Galí and Gertler (1999) in using the labor share of income (again defined for the nonfarm business sector). The motivation for this latter y_t proxy stems from the observation that the sticky-price models underpinning the new-Keynesian Phillips curve imply that the correct driving variable for inflation is actually the log of real marginal cost (nominal marginal cost divided by the price level). Because the theoretical restrictions required in order for real marginal cost to move in line with the traditional output gap are very restrictive, Galí and Gertler (and others) have instead proposed using average unit labor costs—nominal compensation divided by real output—as a proxy for nominal marginal cost. The resulting measure of *real* marginal cost is labor’s share of income (nominal compensation divided by nominal output). Of course, it should also be

⁵Note that the solution given by (9) is only valid when $\frac{\theta}{1-\theta}$ is less than one. Were this term greater than one, then the solution to the stochastic difference equation implied by the hybrid model would instead require us to “solve forward” using the unit root. (This will yield a convergent discounted sum so long as the output gap y_t is a zero-mean variable.) While this solution cannot be ruled out on theoretical grounds, in practice we find strong empirical evidence that the forward root is less than one (in particular, the fit of the models that we consider deteriorates sharply as the value of the forward root approaches unity).

kept in mind that the theoretical conditions under which the labor share can be equated with real marginal cost are themselves quite restrictive.

Output Gap Model: To forecast future values of the output gap, we use a standard two-lag, three-variable VAR that includes the gap, the federal funds rate, and inflation, which we measure as the log-difference of the price deflator for the non-farm business sector.⁶ The sample period extends from 1960:Q1 to 2002:Q1. This simple VAR forecasts the output gap quite well and has been used in a number of papers, including Cochrane (1994), Fuhrer and Moore (1995), and Rotemberg and Woodford (1997).

Panel A of Figure 1 demonstrates that the output gap version of the pure new-Keynesian Phillips curve provides a very poor empirical model of inflation. The discounted sum of output gaps is actually *negatively* correlated with inflation, which directly contradicts the model's predictions. In particular, the model fails to capture the combination of high inflation and deep output gaps that prevailed throughout the mid 1970s and early 1980s; it also predicts that inflation should have risen sharply during the long expansion of the 1990s.⁷ This finding—that the output gap version of the model performs poorly—is robust across a wide range of specifications of the underlying forecasting VAR.

Labor Share Model: To test this version of the model, we augment our existing three-variable VAR with the log of the labor share. Panel B of Figure 1 shows, however, that the discounted sum of labor shares does not do a significantly better job predicting inflation; specifically, although this series is marginally positively correlated with the inflation rate, it only explains about 1 percent of its overall variation.

Unlike the output gap case, however, this finding of a very poor fit is somewhat

⁶All VARs and estimation equations include constant terms.

⁷The recent behavior of the estimated discounted sum of expected output gaps is rather interesting. This series generally declines in recessions, reflecting the positive autocorrelation of the output gap. However, despite a 3.6 percentage point decline in our measure of the output gap from its peak in 2000:Q2 to our final observation in 2002:Q1, the estimated discounted sum has actually increased; this reflects the substantial decline in the federal funds rate and inflation over this period.

sensitive to the choice of underlying VAR. In particular, excluding the output gap from the forecasting system yields a discounted sum that explains a more respectable fraction (around 55 percent) of the observed variation in inflation. This accounts for the difference between our Figure 1 and the results presented by Woodford (2001), who argued that the labor share model fits quite well. Woodford’s estimates were based on a bivariate VAR in the labor share and the log-difference of unit labor costs, with detrended output excluded from the model.⁸ However, there are several reasons to question whether the improvement in fit that can be obtained by excluding the output gap from the VAR should be considered a victory for the new-Keynesian Phillips curve.

First, the hypothesis that detrended output can be excluded from our VAR system is strongly rejected on statistical grounds (this is also true in the context of Woodford’s VAR). By contrast, there is no evidence that inflation Granger-causes the labor income share (which in itself contradicts an important prediction of the new-Keynesian model); moreover, while lagged inflation and the Federal funds rate play almost no role in helping to forecast the labor income share, their inclusion is responsible for most of the model’s ability to fit inflation (if a univariate process for labor’s share is used instead, the resulting discounted sum explains only about 17 percent of the variation in the inflation rate). Finally, the poor performance of the labor share version of the new-Keynesian Phillips curve turns out to be robust across a wide range of VAR specifications that include the output gap.⁹

The Role of Lagged Inflation: In explaining the poor empirical performance of the new-Keynesian Phillips curve, a useful starting point is the observation that both the output gap and labor share variants of the model fail to account for the important role played by lags of inflation in a reduced-form inflation equation like

$$\pi_t = A(L)y_t + B(L)\pi_{t-1}.$$

From equation (2) it is evident that the model predicts that lagged dependent

⁸The specific VAR system that Woodford used is not explicitly discussed in his 2001 paper; we are grateful to Professor Woodford for clarifying the details of these calculations in a personal communication. Note that, because the log-difference of unit labor costs can be written as a linear combination of our measure of price inflation and the changes in the log of the labor share, Woodford’s forecasting VAR is nested within our specification.

⁹See Rudd and Whelan (2002) for these results.

variables will play a role in inflation regressions only to the extent that they are proxying for future values of y_t . Thus, if the model were correct, there should be little role for lagged inflation when estimating the following specification:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k} + A(L)\pi_{t-1}. \quad (11)$$

This is particularly true in this case because our VAR systems include lagged inflation—hence, we have already accounted for any role it plays in forecasting future values of y_t .

In practice, however, this prediction of the model does not hold. For example, if one uses our baseline forecasting VAR to estimate equation (11) with the labor share as y_t and two lags of inflation, the sum of the coefficients on the lags equals 0.90, almost exactly what is obtained in a reduced-form inflation regression. Moreover, this conclusion—that the new-Keynesian model fails to explain the important role played by lagged inflation—is robust even if we use a forecasting VAR that yields a discounted sum of labor shares that is more highly correlated with inflation than the sum obtained from our baseline VAR. For example, if we drop the output gap from our VAR, the resulting estimate of the discounted sum can alone explain more than half the variation in inflation, but the sum of the coefficients on lagged inflation in an equation like (11) is still 0.70. More generally, simple regressions of inflation on its own lags yield R^2 statistics of around 0.75, far in excess of what we can obtain with the best-fitting discounted sum of labor income shares.

The Persistence Problem: It is important to stress that it is this result—the failure of the pure forward-looking model to account for the empirical importance of lagged inflation—that defines the so-called *persistence problem* faced by the new-Keynesian Phillips curve. We make this observation because discussions of inflation persistence have commonly focused on the high *autocorrelation* of inflation, with the implication being that it is this property of the data that sticky-price models should seek to match.¹⁰ However, despite their inability to account for the important role played by lagged inflation, our empirical implementations of the new-Keynesian

¹⁰Fuhrer and Moore (1995), Taylor (1999), and Guerrieri (2002) provide three examples of papers that discuss the new-Keynesian Phillips curve’s “persistence” problem in terms of its ability to match high autocorrelations for inflation.

Phillips curve still predict that inflation should be highly autocorrelated. For example, the first-order autocorrelation coefficients for the discounted sums of the output gap and labor share are 0.95 and 0.92, respectively. These are both *higher* than inflation’s empirical autocorrelation coefficient (of 0.84). Of course, it is unsurprising that these discounted sums are highly autocorrelated given that detrended output and the labor income share are themselves highly autocorrelated variables.

These findings suggest that it is the failure to capture the *inertia* in inflation, given fundamentals, that characterizes the pure forward-looking model’s persistence problem. Put differently, the persistence problem stems from the fact that lagged inflation enters reduced-form inflation equations with coefficients that sum close to one *even after* we have conditioned on a driving variable (such as the output gap) that is itself highly autocorrelated. This suggests that hybrid variants of the basic sticky-price model, which directly allow for a lagged inflation term, may perform better empirically. We now examine these models.

4 The Hybrid Model

As discussed above, the hybrid model can also be assessed empirically using the Campbell-Shiller method, this time applied to the first-difference of inflation. An important difference in this case, however, is that the “discount factor” associated with the infinite sum is now $\frac{\theta}{1-\theta}$, where θ is the weight on expected future inflation in equation (3). The approach that we take here involves using a grid search (over the interval zero to one) to obtain the value of $\frac{\theta}{1-\theta}$ that yields the highest correlation between the resulting discounted sum and the first-difference of inflation.

Output Gap Model: Our results indicate that augmenting the pure forward-looking model with a backward-looking component can reverse our earlier finding of a negative coefficient on the output gap. However, this extension to the model does little to endorse the existence of forward-looking behavior: The grid search reveals that zero is the best-fitting non-negative value of θ , implying a model that reduces to $\Delta\pi_t = \gamma y_t$. This model contains no forward-looking term; hence, expectations of future output gaps do nothing to improve the fit of the model.

In addition, although this simple relationship between the change in inflation

and the output gap is often used as a textbook example of the traditional “accelerationist” Phillips curve, its fit is actually rather poor—specifically, for the data and sample period considered here, this model explains only about 3-1/2 percent of the variance in the first-difference in inflation. This poor fit is illustrated graphically in Figure 2. The top panel of the figure plots the time series for the first-difference of inflation along with the time series for the model’s fitted values; because the change in inflation is such a volatile series, it is somewhat difficult to accurately assess the model’s fit from this chart. Hence, the lower panel of the figure presents a simple scatter diagram; as can be seen from the almost random distribution of the data points, the ability of this model to predict even the *sign* of the change in inflation is quite poor.¹¹

Labor Share Model: The results for the labor share version of the hybrid model are not much more encouraging. In this case, the grid search reveals that the best-fitting hybrid model implies a value for θ of 0.492, which is consistent with approximately equal weights on lagged and expected future inflation. However, as is illustrated in Figure 3, this model does an even worse job than the output gap model in fitting the first-difference of inflation (its R^2 is only 0.01). In addition, a simple regression of $\Delta\pi_t$ on the discounted sum of labor shares yields a t -statistic of only 1.40. Because the explanatory variable in this case is a generated regressor and because we are arbitrarily treating θ as known, this statistic cannot be interpreted as being drawn from a standard distribution (an issue that we will address in the next section). But, together with the model’s low R^2 , these results serve to question whether there is statistical evidence for *any* link between the first-difference of inflation and current and future values of the labor income share.

Comparison with Reduced-Form Regressions: Of course, because the first-difference of inflation is such a volatile variable, we would not necessarily expect such relatively parsimonious models as these to fit very well. That said, a useful benchmark that illustrates just how poorly the hybrid models fit the data can be obtained from a simple regression of $\Delta\pi_t$ on a constant and its own lag. This

¹¹The fact that the model cannot predict the magnitude of these inflation changes can also be seen from the scatterplot: While the x -axis, which plots actual changes in inflation, has a range of 15 percentage points, the fitted values on the y -axis have a range of less than 2 percentage points.

regression has an adjusted R^2 of 0.14; its fit is illustrated graphically in Figure 4. While it is difficult to predict the exact magnitudes of quarterly changes in inflation, this model does much better than either of the hybrid models in matching the direction and size of these changes.

The simple regression achieves this improvement in fit by capturing an important feature of inflation dynamics that is absent from the hybrid model. The coefficient on the lagged change in inflation in this regression is -0.38 , which reflects the fact that the change in inflation is negatively autocorrelated. In contrast, the discounted sums of the output gap (which here is merely the output gap itself) and the labor income share are both highly positively autocorrelated, with first-order autocorrelation coefficients that exceed 0.9. Hence, the discounted sums fundamentally fail to describe a key feature of the $\Delta\pi_t$ process.

Table 1 reports some additional reduced-form regressions for $\Delta\pi_t$. Adding a second lag (column 2) raises the regression's \bar{R}^2 a touch, to 0.15. More interestingly, the inclusion of the output gap also improves the fit of this regression: For the two-lag case, the \bar{R}^2 is 0.22 and the output gap's t -statistic equals 4.06. In contrast, the addition of the labor income share (column 4) yields essentially no improvement in the fit of this regression. These patterns demonstrate that the ability of a standard reduced-form Phillips curve regression—which relates the level of inflation to its own lags (restricting the sum to one) and a measure of slack such as the output gap—to replicate the empirical behavior of inflation is not at all shared by the hybrid sticky-price model.¹²

Finally, column 5 of Table 1 reports the effects of adding two lags of commodity price inflation to the basic reduced-form specification, where commodity prices are defined as the Producer Price Index for crude materials. The purpose of adding this variable is to assess whether the observed negative autocorrelation in $\Delta\pi_t$ reflects volatility in commodity prices. It seems unlikely that the types of frictions envisaged by sticky-price models hold for these types of prices, which are often determined in auction markets. And, as might be expected for a competitively determined price, changes in commodity prices are quite random (there is little correlation between commodity price inflation and its own lags). As a result, one would expect

¹²See Staiger, Stock, and Watson (1997) and Gordon (1998) for two typical implementations of a reduced-form Phillips curve.

the *change* in commodity price inflation to be negatively autocorrelated, and this pattern does indeed hold in the data.¹³ Table 1 shows, however, that while including commodity prices improves the fit of the reduced-form regression, with the \bar{R}^2 rising to 0.32 (see also Figure 5), it does little to alter the coefficients on the lagged changes in inflation.¹⁴

Summary: The results of this section can be summarized as follows.

- The hybrid models can generate a predicted series for the *level* of inflation that is both highly correlated with actual inflation (for either driving variable, this correlation equals 0.85) and highly autocorrelated.
- However, there appears to be very little evidence that the model’s success in matching the level of inflation requires any of the rational forward-looking behavior posited by the hybrid models. In particular, the precise prediction of these models—that the change in inflation should move with a discounted sum of output gaps or labor income shares—is strongly rejected.
- Moreover, the hybrid specifications completely fail to capture important features of the data that can be summarized by simple reduced-form Phillips curves that feature the output gap and several lags of inflation.

These results still leave some important questions unanswered. The first is with what certainty can we reject the forward-looking behavior described in the basic hybrid model? Beyond the weak correlations between the change in inflation and the VAR-based discounted sums, we have not been able to place appropriate standard errors on any of the estimated parameters. The second issue relates to whether a patched-up version of the hybrid model—based, for example, on an alternative rule-of-thumb for backward-looking agents—can do better in matching the data, perhaps thereby revealing an important role for forward-looking behavior. These questions are addressed next.

¹³For example, if commodity price inflation is a white-noise process, then its first-difference will follow an $MA(1)$ process with a first-order autocorrelation coefficient of -0.5 .

¹⁴This result is quite general; it obtains if we employ different specifications for the commodity price term (*e.g.*, if we measure it as a relative price change), or include alternative “supply-shock” measures (such as energy or import prices) in the regression.

5 GMM Estimation

The usefulness of the Campbell-Shiller approach comes from its ability to provide an explicit prediction for the values of $\Delta\pi_t$ that are implied by the hybrid model. However, one drawback of this method is that it cannot be used to derive statistical inferences about the model's parameters γ and θ . An alternative methodology that does not suffer from this problem involves using the generalized method of moments technique (GMM) to estimate the hybrid model. While GMM will no longer permit us to construct a predicted series for $\Delta\pi_t$, it enjoys a distinct advantage over the Campbell-Shiller procedure in that it does not require us to specify an explicit process for the driving term y_t . And, of course, GMM allows us to consistently estimate all of the model's parameters, together with their associated standard errors.

5.1 The Basic Hybrid Model

We use GMM to estimate our basic relationship (equation 9), which relates the change in inflation to a discounted sum of current and expected future values of y_t . This procedure requires us to specify a set of instruments Z_t that are known by agents at time t . Under rational expectations, these instruments will be uncorrelated with the difference between the time- t expectation of the discounted sum in (9) and the discounted sum's realized value; hence, the following orthogonality condition

$$E \left[\left(\Delta\pi_t - \frac{\gamma}{1-\theta} \sum_{k=0}^{\infty} \left(\frac{\theta}{1-\theta} \right)^k y_{t+k} \right) Z_t \right] = 0, \quad (12)$$

should hold in the data. One practical issue that must be dealt with involves the presence of an infinite sum in (12); we address this problem by following the approach of Rudd and Whelan (2001), who noted that orthogonality conditions of this form can also be written as

$$E \left[\left(\Delta\pi_t - \frac{\gamma}{1-\theta} \sum_{k=0}^K \left(\frac{\theta}{1-\theta} \right)^k y_{t+k} - \left(\frac{\theta}{1-\theta} \right)^{K+1} \Delta\pi_{t+K+1} \right) Z_t \right] = 0. \quad (13)$$

The estimates of γ and θ that we obtain using this procedure are reported in Table 2. For the models that use labor's share as a proxy for y_t , the instrument set Z_t consists of four lags each of the change in inflation, the output gap, the labor share, and wage inflation (measured as the log-difference in nonfarm compensation

per hour). When detrended output is used as the driving term, we replace log-differenced hourly compensation—which makes no contribution to first-stage fit—with the federal funds rate, which is a highly significant predictor in the first-stage regressions. We set K equal to 12. (The results are not sensitive to the choice of instruments, to the number of lags of each variable that are included in Z_t , or to the specific value of K that we assume.)

The results in Table 2 confirm the problem with the hybrid model that was suggested by our Campbell-Shiller exercises: For both the output gap and labor share versions of the models, the estimated values of γ are not statistically different from zero. Hence, not only does the discounted sum of future labor shares or output gaps explain very little of the variation in $\Delta\pi_t$, it actually appears to have no statistically discernable influence on this variable whatsoever.

5.2 More General Hybrid Models

Our earlier results suggest one potential route for improving the performance of the hybrid model. Table 1 showed that an implicit assumption underlying the simple hybrid specification—namely, that incorporating a single lag of inflation would allow the model to match the empirical nature of inflation inertia—was incorrect. In particular, the negative autocorrelation of $\Delta\pi_t$ implies that the underlying model for the level of inflation should include more than one lagged dependent variable. One way to address this in the context of the hybrid model is to assume that the underlying structural equation contains an additional inflation lag, thereby taking the form:

$$\pi_t = \theta_1\pi_{t-1} + \theta_2\pi_{t-2} + (1 - \theta_1 - \theta_2)E_t\pi_{t+1} + \gamma y_t. \quad (14)$$

Such a specification could be motivated, for example, by assuming a fraction of non-rational price-setters who use the last two observations of inflation to formulate their expectations, or—within the Christiano, Eichenbaum, and Evans (2001) framework—a more complex indexation rule for those firms who do not set an optimal price this period.

Equation (14) has the following closed-form solution:

$$\Delta\pi_t = \lambda\Delta\pi_{t-1} + \alpha \sum_{k=0}^{\infty} \delta^k E_t y_{t+k}, \quad (15)$$

where the parameters λ , α , and δ will be highly nonlinear functions of the underlying parameters, θ_1 , θ_2 , and γ . In Table 3, we report GMM estimates of λ , α , and δ that are obtained using the same procedure and the same instrument sets that were used in estimating equation (14). We report direct estimates of these coefficients rather than complicating the analysis by attempting to back out estimates of θ_1 , θ_2 , and γ ; this is because, from our perspective, the key issue is whether this approach produces statistically significant and economically sensible values for α and δ (*i.e.*, whether allowing for extra lags of inflation improves the case for the existence of a forward-looking rational expectations term).

Not surprisingly, Table 3 indicates that the coefficient on $\Delta\pi_{t-1}$ is negative and highly statistically significant. But this exercise still fails to produce any convincing evidence for forward-looking behavior. For the output gap version of the model, the coefficient on the discounted sum, α , is statistically significant, but the estimated forward root δ is *negative*, which is not reasonable in this context. For the labor share version, the estimated forward root is positive, but the coefficient on the discounted sum receives a t -statistic of only 1.04. On the whole, then, these results do little to endorse the presence of forward-looking rational expectations, and thus the case for a more complex hybrid model featuring extra lags of inflation.

6 Conclusions

The observation that lagged inflation plays an important role in empirical inflation regressions has posed a major challenge to the rational expectations sticky-price models that underpin the new-Keynesian Phillips curve. Indeed, it has now become relatively well accepted that purely forward-looking models of inflation cannot account for the degree of inflation inertia that we actually observe in the data, and that this failure significantly reduces their usefulness in assessing practical policy questions. In response, researchers have increasingly adopted hybrid pricing specifications, in which lagged inflation is allowed to have an explicit role in price setting. This class of model is widely seen as striking a reasonable compromise between the desire to fit a key empirical characteristic of the inflation process (its inertia), and the desire to preserve an important role for forward-looking, rational expectations in price setting.

The goal of this paper has been to determine whether this reformulation of the

basic sticky-price model yields a pricing specification that is capable of capturing empirical inflation behavior. We have shown that the hybrid specification implies a strong prediction about the inflation process that is easily tested—and firmly rejected. In fact, we find no evidence in postwar U.S. data that inflation dynamics reflect the type of rational forward-looking behavior that the model hypothesizes. Hence, while the addition of a lagged inflation term permits the hybrid model to better capture certain features of the inflation process, ultimately this fix is cosmetic in that the feature of the model that truly distinguishes it from alternative models of inflation—such as a traditional Phillips curve based on backward-looking expectations—appears to be empirically invalid.

One possible conclusion is that the problem with these models stems from their reliance on a strict form of rational expectations. 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 thus that firms almost surely attempt to make some prediction about future inflation when determining their current price. What appears to be less reasonable, however, is that these predictions are formulated in the manner suggested by the new-Keynesian model under rational expectations. Put differently, it may well be that $E_t\pi_{t+1}$ is a key influence on current inflation. But if this is so, the evidence indicates that this expectation is not determined in the manner that a rational expectations sticky-price model would predict.

We conclude, then, that further research in this area is probably best aimed toward developing models that deviate from the standard rational-expectations framework in favor of alternative descriptions of how agents process information and develop forecasts. Work in this vein by Sims (1998) and Mankiw and Reis (2002) may prove to be a promising start in this direction.

References

- [1] Calvo, Guillermo A. (1983). “Staggered Prices in a Utility Maximizing Framework,” *Journal of Monetary Economics*, 12, 383-398.
- [2] Campbell, John and Robert Shiller (1987). “Cointegration and Tests of Present Value Models,” *Journal of Political Economy*, 95, 1062-1088.
- [3] Casares, Miguel and Bennett T. McCallum (2000). “An Optimizing IS-LM Framework with Endogenous Investment,” NBER Working Paper No. 7908.
- [4] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (2001). “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” NBER Working Paper No. 8403.
- [5] Clarida, Richard, Jordi Galí, and Mark Gertler (1999). “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661-1707.
- [6] Cochrane, John (1994). “Shocks,” *Carnegie-Rochester Conference Series on Public Policy*, 41, 295-364.
- [7] Ehrmann, Michael, and Frank Smets (2001). “Uncertain Potential Output: Implications for Monetary Policy,” ECB Working Paper No. 59.
- [8] Fuhrer, Jeffrey C. and George R. Moore (1995). “Inflation Persistence,” *Quarterly Journal of Economics*, 110, 127-159.
- [9] Galí, Jordi and Mark Gertler (1999). “Inflation Dynamics: A Structural Econometric Analysis,” *Journal of Monetary Economics*, 44, 195-222.
- [10] Gordon, Robert J. (1998). “Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU,” *Brookings Papers on Economic Activity*, 297-346.
- [11] Guerrieri, Luca (2002). “The Inflation Persistence of Staggered Contracts,” Federal Reserve Board, International Finance Discussion Paper No. 734.
- [12] Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117, 1295-1328.

- [13] Roberts, John M. (1997). “Is Inflation Sticky?” *Journal of Monetary Economics*, 39, 173-196.
- [14] Rotemberg, Julio (1982). “Sticky Prices in the United States,” *Journal of Political Economy*, 60, 1187-1211.
- [15] Rotemberg, Julio and Michael Woodford (1997). “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” *NBER Macroeconomics Annual*, 12, 297-346.
- [16] Rudd, Jeremy and Karl Whelan (2001). “New Tests of the New-Keynesian Phillips Curve,” Federal Reserve Board, Finance and Economics Discussion Series Paper No. 2001-30 (forthcoming, *Journal of Monetary Economics*).
- [17] Rudd, Jeremy and Karl Whelan (2002). “Does the Labor Share of Income Drive Inflation?” Federal Reserve Board, Finance and Economics Discussion Series Paper No. 2002-30.
- [18] Rudebusch, Glenn (2002). “Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty,” *Economic Journal*, 112, 402-432.
- [19] Sargent, Thomas J. (1987). *Macroeconomic Theory* (2nd ed.), Academic Press.
- [20] Sims, Christopher A. (1998). “Stickiness,” *Carnegie-Rochester Conference Series on Public Policy*, 49, 317-356.
- [21] Staiger, Douglas, James Stock, and Mark Watson (1997). “How Precise Are Estimates of the Natural Rate of Unemployment?” in Christina Romer and David Romer (eds.), *Reducing Inflation*, Chicago: University of Chicago Press.
- [22] Taylor, John (1999). “Staggered Price and Wage Setting in Macroeconomics,” in John Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, North-Holland, 1009-1050.
- [23] Woodford, Michael (2001). “The Taylor Rule and Optimal Monetary Policy,” *American Economic Review*, 91(2), 232-237.
- [24] Woodford, Michael (2002). *Interest and Prices*, manuscript in preparation.

Table 1: Estimated Reduced-Form Models for $\Delta\pi_t$

<i>Included variables</i>	Specification				
	1	2	3	4	5
$\Delta\pi_{t-1}$	-0.378** (0.072)	-0.422** (0.077)	-0.488** (0.075)	-0.425** (0.077)	-0.490** (0.077)
$\Delta\pi_{t-2}$		-0.119 (0.078)	-0.179* (0.076)	-0.122 (0.078)	-0.167* (0.071)
y_t			0.122** (0.030)		0.080** (0.029)
s_t				2.302 (5.666)	
π_{t-1}^{com}					0.031** (0.007)
π_{t-2}^{com}					0.006 (0.008)
\bar{R}^2	0.138	0.145	0.218	0.141	0.315

Note: $y_t \equiv$ detrended output, $s_t \equiv$ labor's share of income, $\pi_t^{com} \equiv$ commodity price inflation. Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively.

Table 2: GMM Estimates of Hybrid Inflation Equation

<i>Driving variable (y_t)</i>	γ	θ
Detrended output	0.041 (0.052)	0.269 (0.295)
Labor income share	0.008 (0.011)	0.439** (0.077)

Note: Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively.

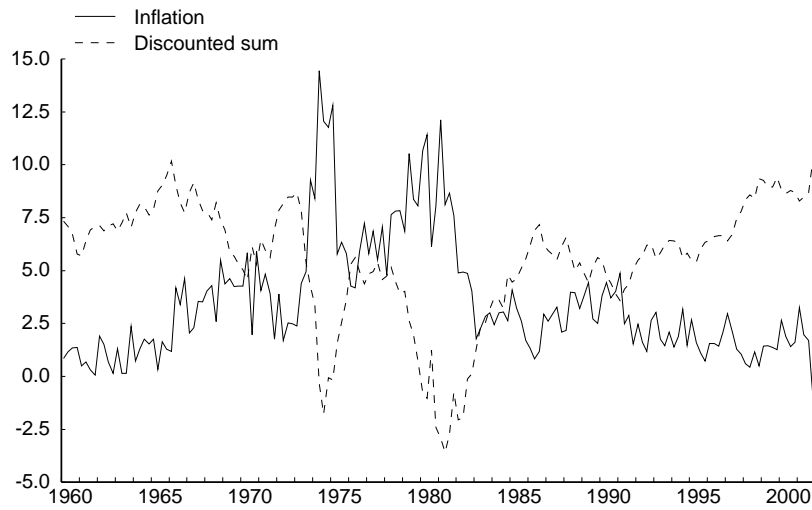
Table 3: GMM Estimates of Augmented Hybrid Inflation Equation

<i>Driving variable (y_t)</i>	λ	α	δ
Detrended output	-0.440** (0.079)	0.163** (0.037)	-0.955** (0.034)
Labor income share	-0.388** (0.037)	0.022 (0.021)	0.778** (0.222)

Note: Table gives estimated reduced-form parameters from the augmented hybrid model $\Delta\pi_t = \lambda\Delta\pi_{t-1} + \alpha\sum_{i=0}^{\infty}\delta^i E_t y_{t+i}$. Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively.

Figure 1
Fit of New-Keynesian Phillips Curve

A. Output Gap Version (beta=0.99)



B. Labor Share Version (beta=0.99)




Figure 2
Fit from Regressing Change in Inflation on Detrended Output




Figure 3
Fit for Change in Inflation, Labor Share Hybrid Model




Figure 4
Fit for Change in Inflation, AR(1) Model




Figure 5
Fit for Change in Inflation (Reduced-Form Model with Commodity Prices)

