

5/RT/03

August 2003

CENTRAL BANK AND FINANCIAL
SERVICES AUTHORITY OF IRELAND
Research Technical Paper

*Can Rational Expectations Sticky-Price
Models Explain Inflation Dynamics?*

by

Jeremy Rudd
Federal Reserve Board

Karl Whelan
CBFSAI*

Economic Analysis and Research Department
P.O. Box 559, Dame Street
Dublin 2
Ireland
<http://www.centralbank.ie/research.asp>

*The views expressed in this paper are the personal responsibility of the authors. They are not necessarily held either by the Central Bank and Financial Services Authority of Ireland, the ESCB or the Board of Governors of the Federal Reserve System. The authors are grateful to Dale Henderson, Frank Smets, and William Wascher for helpful comments on earlier drafts of this paper. They can be contacted at jeremy.rudd@frb.gov and karl.whelan@centralbank.ie.

Abstract

The canonical inflation specification in sticky-price rational expectations models (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, many recent studies have employed a “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 that are derived from the model’s closed-form solution. 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 rational forward-looking behavior implied by the model.

1 Introduction

In recent years, there has been a trend in macroeconomics toward analyzing business cycles and stabilization policy in the context of models that incorporate both nominal rigidities and optimizing agents with rational (*i.e.*, model-consistent) expectations.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 inflation 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 close to or equal to 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 has often been criticized on the grounds that it cannot account for the important role played by lagged dependent variables in inflation regressions. In response to this critique, 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

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

a long-run level tradeoff between inflation and real activity.²

Within the class of papers employing variants of the hybrid specification (3), the best-known studies have featured models in which $\theta \leq 1/2$. For example, the well-known model of Fuhrer and Moore (1995) employs an assumption that workers bargain over relative real wages in order to obtain an equation with $\theta = 1/2$. 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 their framework, θ equals $\frac{\beta}{1+\beta}$ (where β is the factor used to discount firms' profits); this directly implies that θ will be less than $1/2$.

In this paper, we assess whether hybrid models of this sort can provide 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) with $\theta \leq 1/2$ implies an expression for the *change* in inflation of the form

$$\Delta\pi_t = \lambda_1 \sum_{k=0}^{\infty} \lambda_2^k E_t y_{t+k}, \quad (4)$$

where $\lambda_2 \leq 1$. 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 can be predicted well from its own lagged value; hence, incorporating lagged inflation into the inflation equation should allow the hybrid model to fit the level of inflation relatively well. However, such a fit could also 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

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

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.

While variants of the hybrid specification in which $\theta \leq 1/2$ have received a large amount of attention in recent work, there is no *a priori* reason to rule out the possibility that price setting is characterized by a preponderance of forward-looking behavior. We therefore also consider versions of the hybrid model with $\theta > 1/2$, which imply the following closed-form solution:

$$\pi_t = \mu_1 \sum_{k=0}^{\infty} E_t y_{t+k} + \mu_2 \pi_{t-1}. \quad (5)$$

Here, the *level* of current inflation is related to lagged inflation (with $\mu_2 < 1$) and current and expected future values of the driving term, where these receive a unit weight in all periods. Again, the presence of lagged inflation ensures that this model will be able to fit π_t relatively well; hence, the relevant question here concerns what contribution the forward-looking terms make to explaining inflation dynamics.

Taken as a whole, our results suggest that the hybrid model provides a 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. Moreover, we find that the coefficient on the discounted sum (λ_1 or μ_1) is not significantly different from zero for *any* variant of the hybrid model that we consider, implying that inflation is unrelated to the expectation of future values of the driving term and indicating that the type of rational forward-looking behavior hypothesized by the hybrid model is absent from the data. Importantly, these conclusions hold 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).

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. Finally, Section 6 repeats our analysis for the version of the model that obtains when $\theta > 1/2$, and Section 7 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 with $\theta \leq 1/2$; we then discuss how the VAR-based techniques of Campbell and Shiller (1987) can be used to assess these models.

Begin by considering the pure 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

$$Z_t = AZ_{t-1} + \epsilon_t, \tag{6}$$

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

$$\sum_{k=0}^{\infty} \beta^k E_t y_{t+k} = e'_1 (I - \beta A)^{-1} Z_t, \tag{7}$$

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 (6), with a point estimate

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.

of γ in turn obtained by regressing π_t on the matrix expression in (7).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 discount factor applied by firms to future profits. In the calculations that we present in the next sub-section (which use quarterly data), 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

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

A couple of simple substitutions then yield

$$\Delta\pi_t = \frac{\theta}{1-\theta} E_t \Delta\pi_{t+1} + \frac{\gamma}{1-\theta} y_t. \quad (9)$$

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 ,

$$\Delta\pi_t = \lambda_1 \sum_{k=0}^{\infty} \lambda_2^k E_t y_{t+k}, \quad (10)$$

with the “discount factor” λ_2 equal to $\frac{\theta}{1-\theta}$ in this case.⁵ Again using a VAR like (6) to generate expectations of y_t yields

$$\Delta\pi_t = \lambda_1 e_1' (I - \lambda_2 A)^{-1} Z_t. \quad (11)$$

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 λ_2 . Thus, for the estimates reported in Section 4, we use a grid search to choose the value of λ_2 that yields the best-fitting hybrid equation. In addition, we should note that the solution given by (10) is only valid when $\theta \leq 1/2$, as this implies that λ_2 (which equals $\frac{\theta}{1-\theta}$) will be less than or equal to one, thus ensuring that the term on the right-hand-side is not explosive. We defer a consideration of the case where $\theta > 1/2$ to Section 6.

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.

⁵See the Appendix for details.

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 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 follows Galí and Gertler (1999) in using (the log of) labor’s 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 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 be 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 output gap, the federal funds rate, and inflation, which we measure as the log-difference of the price deflator for the nonfarm 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

⁶All VARs and estimation equations include constant terms.

the mid 1970s and early 1980s; it also predicts that inflation should have risen sharply during the long expansion of the 1990s. 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 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.⁷ 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 good news 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); moreover, 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.⁸ Second, there is no evidence that inflation Granger-causes the labor income share, which in itself provides direct evidence

⁷The 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 changes in the log of the labor share, Woodford’s forecasting VAR is nested within our specification.

⁸See Rudd and Whelan (2002) for these results.

against the model’s prediction that inflation summarizes agents’ expectations about future values of the driving term y_t . Finally, 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).

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 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:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k} + A(L)\pi_{t-1}. \quad (12)$$

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 (12) with the labor share as y_t and two lags of inflation, the sum of the coefficients on the lags equals 0.90, which is 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 (12) 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 even 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.⁹ However, despite their inability to account for the important role played by lagged inflation, our empirical implementations of the new-Keynesian 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 with $\theta \leq 1/2$ can also be assessed empirically using the Campbell-Shiller method, this time applied to the first-difference of infla-

⁹Fuhrer and Moore (1995), Taylor (1999), and Guerrieri (2002) are 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.

tion. An important difference in this case, however, is that the “discount factor” associated with the infinite sum— λ_2 in equation (4)—is equal to $\frac{\theta}{1-\theta}$, 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 λ_2 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 λ_2 , implying a model that reduces to $\Delta\pi_t = \lambda_1 y_t$. In this model, then, expectations of future output gaps do nothing to improve the equation’s fit.

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 in quarterly data—specifically, over the sample period considered here, this model explains only about 3-1/2 percent of the variance in the first-difference in inflation. This mediocre 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.¹⁰

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 λ_2 of 0.97, so the discounted sum does not

¹⁰The 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.

vanish. 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 income 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 λ_2 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 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 important aspects of the empirical behavior of inflation is not at all shared by the hybrid sticky-price model.¹¹

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 to what degree 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 the *change* in commodity price inflation to be negatively autocorrelated, and this pattern does indeed hold in the data.¹² 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 pattern of negative coefficients on the lagged changes in inflation.¹³

Results Using Annual Data: An additional factor that could contribute to the negative autocorrelation that we observe in $\Delta\pi_t$ is the presence of serially uncorrelated measurement error (or some other type of transitory high-frequency shock) in inflation. Noise of this sort would have an effect similar to that described above for commodity prices, and could act to obscure any relationship between the first-difference of inflation and the discounted sum of the driving variable.

To test this possibility, we use annual data to re-estimate the output gap and

¹¹See Staiger, Stock, and Watson (1997) and Gordon (1998) for two typical implementations of a reduced-form Phillips curve.

¹²For 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 .

¹³This 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.

labor share variants of the hybrid model. When we do so, we find that none of our principal conclusions are altered; in particular, we still find that the expected discounted sum of the labor income share explains very little of the variance in $\Delta\pi_t$ while the best-fitting value of λ_2 in the version of the hybrid model that uses detrended GDP remains zero (thus implying that forward-looking behavior is completely absent from the model).¹⁴

The reason for the hybrid model's inability to fit annual data is closely related to the source of the model's failure in quarterly data. Recall that, in quarterly data, $\Delta\pi_t$ was negatively autocorrelated while the estimated discounted sum of the driving term was highly positively autocorrelated. Using annual data smooths away much of the high-frequency variation in $\Delta\pi_t$, and leaves the first difference of inflation essentially uncorrelated with its own lags. However, the estimated discounted sums of both the output gap and labor's share remain strongly positively autocorrelated in annual data. Hence, our demonstration of the hybrid model's inability to provide a good characterization of the $\Delta\pi_t$ process does not depend on the use of quarterly data.

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

- The popular class of hybrid models for which $\theta \leq 1/2$ can generate predicted series for the *level* of inflation that are both highly correlated with actual inflation (for either driving variable, this correlation equals 0.85 in quarterly data) and highly autocorrelated.
- However, there appears to be very little evidence that the models' success in matching the level of inflation requires any of the rational forward-looking behavior posited by the hybrid models. In particular, the prediction of these models that distinguishes them from backward-looking alternatives—that the change in inflation should move with a discounted sum of output gaps or labor income shares—is strongly rejected.
- Moreover, these specifications completely fail to capture important features of the data that can be summarized by simple reduced-form Phillips curves

¹⁴Note, however, that a model relating $\Delta\pi_t$ to the level of detrended GDP fits somewhat better in annual data.

that feature the output gap and several lags of inflation.

These results still leave some important questions unanswered. The first involves the certainty with which we can rule out the presence of forward-looking behavior in the hybrid inflation specifications: Beyond the weak correlation that we found between the change in inflation and the VAR-based discounted sums, we have not been able to formally assess the statistical significance of the forward-looking terms. The second issue relates to whether a patched-up version of the class of hybrid models with $\theta \leq 1/2$ —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. Finally, there is the question of how models based on the assumption of $\theta > 1/2$ perform. These questions are addressed next.

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—in particular, we cannot determine whether the driving term’s discounted sum makes a statistically significant contribution to observed inflation dynamics. 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 $\Delta\pi_t$ (and, hence, to assess the model’s fit), 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 both λ_1 and λ_2 , together with their associated standard errors.

5.1 The Basic Hybrid Model

We use GMM to estimate our basic relationship (equation 10), 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 (10) and the discounted sum's realized value; hence, the following orthogonality condition

$$E \left[\left(\Delta\pi_t - \lambda_1 \sum_{k=0}^{\infty} \lambda_2^k y_{t+k} \right) Z_t \right] = 0, \quad (13)$$

should hold in the data. One practical issue that must be dealt with involves the presence of an infinite sum in (13); 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

$$E \left[\left(\Delta\pi_t - \lambda_1 \sum_{k=0}^K \lambda_2^k y_{t+k} - \lambda_2^{K+1} \Delta\pi_{t+K+1} \right) Z_t \right] = 0. \quad (14)$$

The estimates of λ_1 and λ_2 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 two 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 λ_1 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.¹⁵ 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:

$$\pi_t = \theta_1\pi_{t-1} + \theta_2\pi_{t-2} + (1 - \theta_1 - \theta_2)E_t\pi_{t+1} + \gamma y_t. \quad (15)$$

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 (15) has the following closed-form solution:

$$\Delta\pi_t = \lambda_1 \sum_{k=0}^{\infty} \lambda_2^k E_t y_{t+k} + \lambda_3 \Delta\pi_{t-1}, \quad (16)$$

where the parameters λ_1 , λ_2 , and λ_3 represent nonlinear functions of the underlying parameters θ_1 , θ_2 , and γ . In Table 3, we report GMM estimates of λ_1 , λ_2 , and λ_3 that are obtained using the same procedure and the same instrument sets that were used in estimating equation (14). Again, the key question is whether we obtain statistically significant and economically sensible values for λ_1 and λ_2 (*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, λ_1 , is statistically significant, but the estimated forward root λ_2 is *negative*, which is not reasonable in this context. For the labor

¹⁵Put differently, it is possible to demonstrate that the first difference of an inflation measure will be negatively autocorrelated if its *level* can be empirically characterized by an inflation equation with more than one lagged dependent variable (where the lag coefficients sum close to unity). This is why $\Delta\pi_t$ tends to be negatively autocorrelated in quarterly data for virtually all measures of final-goods prices—including the core CPI, the core PCE price index, and the GDP deflator—as well as for the measure that we use in our estimation exercises (the price deflator for nonfarm business output), *even if* we control for such transitory factors as supply shocks.

share version, the estimated forward root is positive, but the coefficient on the discounted sum receives a t -statistic of only 0.65. 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.¹⁶

6 The Hybrid Specification with $\theta > 1/2$

The versions of the hybrid model that we have considered up to this point involve values of θ that are less than or equal to one-half. We now examine whether a statistically significant role for forward-looking behavior can be found in versions of the hybrid model for which $\theta > 1/2$.

As we demonstrate in the Appendix, it is relatively straightforward to show that, in this case, the solution to the expectational difference equation implied by the hybrid model (equation 3) will involve a root equal to unity and a root μ_2 equal to $\frac{1-\theta}{\theta}$ (which is less than one). Since the output gap y_t is a zero-mean variable, a convergent solution will result even if we solve forward using the unit root. Hence, in this case we can derive the following closed-form solution for the *level* of inflation,

$$\pi_t = \mu_1 \sum_{k=0}^{\infty} E_t y_{t+k} + \mu_2 \pi_{t-1}. \quad (17)$$

Table 4 presents the results that obtain from estimating equation (17) using the same methodology and instrument set as before (note, though, that here we use π_t as an instrument in lieu of $\Delta\pi_t$). Once again, we find that μ_1 is not statistically significant no matter which measure of the driving variable we use, which in turn implies that forward-looking behavior (as summarized by the sum of current and expected future values of y_t) plays no discernable empirical role in determining inflation. Moreover, this result obtains even if we allow the equation to include additional lags of inflation, as in

$$\pi_t = \mu_1 \sum_{i=0}^{\infty} E_t y_{t+i} + \mu_2 \pi_{t-1} + \mu_3 \pi_{t-2}, \quad (18)$$

in that μ_1 remains statistically insignificant in this specification (see Table 5).¹⁷

¹⁶Note that we obtain essentially identical results if we use our GMM procedure to fit equations (10) and (16) to annual data.

¹⁷As can be seen from the table, the statistical significance of μ_1 in the labor share version

7 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 generates precise predictions about the inflation process that are 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 conclusion that can be drawn from these results is that the hybrid model’s approach to patching up the new-Keynesian Phillips curve—which involves a direct attempt to deal with its persistence problem—may merely be addressing a symptom of what is in fact a much more deeply rooted problem with this type of model. Specifically, our findings suggest that pricing models of this sort suffer from

of the model rises somewhat when an additional lag of inflation is included (though μ_1 remains insignificant at the 10 percent level). However, this turns out to be extremely sensitive to our specific choice of K (the number of leads of y_t and terminal value of π_t that we employ in the estimation equation): Varying K by even a small amount causes the coefficient’s p -value to drop sharply.

a more serious (and less easily addressed) weakness; namely, 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 the assumption that these predictions are formulated in the manner implied 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 the current generation of rational expectations sticky-price models would predict. This conclusion does not rule out a role for *some* sort of rational optimizing behavior in explaining inflation dynamics; indeed, there may be an optimization-based rationale for why the reduced-form Phillips curve models discussed in this paper fit so well. For example, in the absence of any agreement amongst economists on what the correct models for inflation (or the rest of the economy) actually are, and given most individuals' limited ability to understand or model these uncertainties, a model in which agents base their expectations for future inflation on extrapolations of the recent past may itself constitute a form of optimizing behavior.

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, 2003) 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] Rotemberg, Julio (1982). "Sticky Prices in the United States," *Journal of Political Economy*, 60, 1187-1211.
- [14] Rotemberg, Julio and Michael Woodford (1997). "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297-346.
- [15] 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*).
- [16] 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 (forthcoming, *Journal of Money, Credit, and Banking*).
- [17] Rudebusch, Glenn (2002). "Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty," *Economic Journal*, 112, 402-432.
- [18] Sargent, Thomas J. (1987). *Macroeconomic Theory* (2nd ed.), Academic Press.
- [19] Sims, Christopher A. (1998). "Stickiness," *Carnegie-Rochester Conference Series on Public Policy*, 49, 317-356.
- [20] Sims, Christopher A. (2003). "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665-690.
- [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.

A Solutions to the Hybrid Model

Re-writing the hybrid model in terms of the lag operator L , we obtain that

$$\left[\frac{1-\theta}{\theta} L^2 - \frac{1}{\theta} L + 1 \right] \pi_t = -\frac{\gamma}{\theta} L y_t.$$

The lag polynomial can be factored as

$$(1 - z_1 L)(1 - z_2 L) \pi_t = -\frac{\gamma}{\theta} L y_t.$$

It is straightforward to apply the quadratic formula to show that one root of the characteristic polynomial will always equal one, while the other will equal $\frac{\theta}{1-\theta}$. Hence, the solution to the difference equation implied by the hybrid model can be written as

$$(1 - L) \left(1 - \frac{1-\theta}{\theta} L \right) \pi_t = -\frac{\gamma}{\theta} L y_t.$$

When $\theta \leq 1/2$, the stable solution is found by multiplying through by the “forward inverse” of $(1 - \frac{1-\theta}{\theta} L)$, which is

$$\frac{-\frac{\theta}{1-\theta} L^{-1}}{1 - \frac{\theta}{1-\theta} L^{-1}}.$$

This yields an equation of the form

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

which is the first solution we considered.

Alternatively, when $\theta > 1/2$, the stable solution is found by multiplying through by the forward inverse of $(1 - L)$, which is $\frac{-L^{-1}}{1-L^{-1}}$. We then obtain

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

which is the solution we examined in Section 6.

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</i> (y_t)	λ_1	λ_2
Detrended output	0.039 (0.035)	0.614 ^a (0.372)
Labor income share	0.017 (0.030)	0.769 (0.498)

Note: Table gives estimated values for the parameters from the basic hybrid model $\Delta\pi_t = \lambda_1 \sum_{i=0}^{\infty} \lambda_2^i E_t y_{t+i}$. 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</i> (y_t)	λ_1	λ_2	λ_3
Detrended output	0.146** (0.048)	-0.990** (0.050)	-0.364* (0.141)
Labor income share	0.024 (0.036)	0.764 (0.465)	-0.392** (0.053)

Note: Table gives estimated values for the parameters from the augmented hybrid model $\Delta\pi_t = \lambda_1 \sum_{i=0}^{\infty} \lambda_2^i E_t y_{t+i} + \lambda_3 \Delta\pi_{t-1}$. Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively.

Table 4: GMM Estimates of Hybrid Inflation Equation with $\theta > 1/2$

<i>Driving variable (y_t)</i>	μ_1	μ_2
Detrended output	-0.007 (0.005)	0.622** (0.069)
Labor income share	0.017 (0.014)	0.485** (0.072)

Note: Table gives estimated values for the parameters from the alternative basic hybrid model $\pi_t = \mu_1 \sum_{i=0}^{\infty} E_t y_{t+i} + \mu_2 \pi_{t-1}$. Standard errors in parentheses; **/*/^a denotes significant at 1/5/10 percent level, respectively.


Table 5: GMM Estimates of Augmented Hybrid Inflation Equation with $\theta > 1/2$

<i>Driving variable (y_t)</i>	μ_1	μ_2	μ_3
Detrended output	0.001 (0.005)	0.429** (0.051)	0.325** (0.051)
Labor income share	0.025 (0.015)	0.337** (0.071)	0.247** (0.057)

Note: Table gives estimated values for the parameters from the alternative augmented hybrid model $\pi_t = \mu_1 \sum_{i=0}^{\infty} E_t y_{t+i} + \mu_2 \pi_{t-1} + \mu_3 \pi_{t-2}$. 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)

