Should Monetary Policy Target Labor's Share of Income?

Jeremy Rudd and Karl Whelan
Division of Research and Statistics
Federal Reserve Board*

Second Draft
February 8, 2002

Abstract

In recent work, Woodford (2001) presents evidence that using real unit labor costs (labor's share of income) as a driving variable in the new-Keynesian Phillips curve yields a superior fit for inflation relative to a model that uses deterministically detrended real GDP. This evidence leads him to conclude that the output gap—the deviation between actual and potential output—is better captured by the labor income share, in turn implying that the monetary authority should raise interest rates in response to increases in this variable. We document that the empirical case for the superiority of the labor's share version of the new-Keynesian model is actually quite weak, and conclude that there is little reason to view the labor income share as an appropriate target for monetary policy.

*Mail Stop 80, 20th and C Streets NW, Washington DC 20551. E-mail: jeremy.b.rudd@frb.gov.
We are grateful to Argia Sbordone and Michael Woodford for helpful comments on an earlier draft of the paper. The views expressed in this paper are our own and do not necessarily reflect the views of the Board of Governors or the staff of the Federal Reserve System.

1 Introduction

Recent years have seen an explosion in research aimed at assessing monetary policy rules using macroeconomic models built from explicit microfoundations. A crucial element of such models involves their specification of price setting behavior. The usual aggregate supply relationship implied by the types of optimizing sticky-price models employed in the optimal monetary policy literature is a new-Keynesian Phillips curve of the form

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

where mc_t denotes economy-wide real marginal cost—*i.e.*, nominal marginal cost divided by the price level. Real marginal cost therefore serves as the principal determinant of inflation in these models; moreover, the welfare loss associated with aggregate output's differing from potential—the level of output that would obtain under perfectly flexible prices—is related to the discrepancy between firms' nominal supply costs and their price.

As has recently been discussed by Woodford (2001), these sticky-price models may call into question the theoretical rationale for the Taylor rule, under which a central bank adjusts its short-term interest-rate target with reference to the levels of inflation and deterministically detrended real GDP (the "GDP gap"). In the context of these models, the effectiveness of this version of the Taylor rule depends upon the GDP gap's ability to serve as a satisfactory proxy for real marginal cost. Because the new-Keynesian Phillips curve generates a number of counterfactual empirical predictions when detrended GDP is used as a proxy for real marginal cost, several researchers have suggested using an alternative proxy—specifically, real unit labor costs (*i.e.*, labor's share of income). For example, Woodford cites the research of Sbordone (1998) as implying that the new-Keynesian Phillips curve "gives a very poor account of U.S. inflation when detrended real GDP is used as the gap measure but explains much of the medium-frequency variation when real unit labor costs are used instead."

The idea that labor's share of income represents a more appropriate measure of marginal cost—and hence that empirical models based upon this measure provide a good description of observed inflation dynamics—carries substantial implications for the optimal conduct of monetary policy. For example, one prescription that follows from this view is that central banks should ignore standard measures of the output gap and instead raise interest rates in response to increases in labor's share of income (which provides a more useful read on inflation pressures).[1] The unconventional nature of such a policy rule is readily apparent from Figure 1: Because the labor share of income has risen in every modern U.S. recession, a policy rule based on this variable instead of the GDP gap would have called for higher interest rates during each of these episodes.

In this paper, we provide some new evidence on the relative merits of the detrended output and labor share variants of the new-Keynesian Phillips curve. Specifically, we re-examine two pieces of evidence cited by Woodford (2001) as demonstrating the superiority of the labor's share version of this relationship, and find that the case for this view is far weaker than has been suggested.

First, we take a closer look at the empirical results presented by Woodford. These results indicated that an inflation series based on the discounted sum of expected future labor share values—the solution obtained from solving equation (1) forward—fits much better than a series based on detrended output, where the expectations of the driving terms are calculated using a reduced-form VAR. We show that this result is not robust; in particular, we demonstrate that the fit of the labor's share version of the new-Keynesian inflation equation is highly sensitive to small changes in the VAR that is used to forecast future labor's share values. For a broad range of reasonable specifications of this VAR, the model's fit is actually quite poor.

Second, we re-examine the evidence in Sbordone (1998). While this paper did not, in fact, compare the relative fits of inflation series based on the GDP gap and labor share versions of the new-Keynesian Phillips curve, it did show that it was possible to use the labor share version to construct an inflation series that fits the data well. We show, however, that the principal reason for the good fit obtained by Sbordone is not her use of the labor income share as a driving variable, but rather an additional assumption that produces a different closed-form solution for inflation than the one that is usually tested. Specifically, Sbordone's methodology is based upon the assumption that agents believe that the proxy for *nominal* marginal cost (*i.e.*, nominal unit labor costs) evolves independently of the price level, an assumption that we view as unrealistic. We find that when the same assumption

[1] The logic underlying this prescription is that a high value of labor's share provides a better signal that real marginal cost is high, and thus that output is running above its potential level.

is made for detrended output (specifically, that agents believe nominal output to evolve independently of prices), then this version of the new-Keynesian Phillips curve also fits the data well.

We conclude that the new-Keynesian Phillips curve provides a poor description of the inflation process even if labor's share is used to proxy for real marginal cost. We therefore find little empirical justification for including labor's share in a monetary policy rule.

2 Two Measures of Real Marginal Cost

As noted in the introduction, the driving variable in the new-Keynesian inflation equation (1) is real marginal cost—a variable that is well defined in theory, but difficult to measure in practice. Two proxies for real marginal cost have typically been used to evaluate the new-Keynesian inflation equation: deterministically detrended real GDP, and real unit labor costs. We briefly consider the logic behind each measure.

The motivation for the traditional specification of the new-Keynesian Phillips curve, in which a measure of detrended real output is used as the driving term, stems from the assumption that marginal cost functions are upward-sloping—hence, higher levels of production relative to some level of potential output y_t^* will yield a higher level of real marginal costs.[2] While straightforward in principle, this idea is complicated in practice by the fact that we do not actually observe y_t^*. As a result, empirical models are instead typically based on the assumption that y^* evolves according to a deterministic trend.

The approach taken by Sbordone (1998) attempts to circumvent the fact that potential output is unobservable by instead directly focusing on the measurement of marginal cost. However, because there are also substantial difficulties surrounding the measurement of this variable, in practice this method relies on using average cost as a proxy for marginal cost. Specifically, Sbordone's approach involves using average unit labor costs (total nominal compensation divided by real output) as a proxy for nominal marginal cost. An implication of this procedure is that the proxy

[2] See Christiano, Eichenbaum, and Evans (2001) and Dotsey and King (2001) for a discussion of the relationship between real marginal cost and output in the context of standard pricing models.

for *real* marginal cost in this case is the labor income share (nominal compensation divided by nominal output). If real marginal cost is related to the deviation of output from its potential level, these considerations also suggest that the labor share can serve as a proxy for the gap between actual output and a stochastic series for potential output.

Once we have settled on a candidate measure of real marginal cost, we must decide on an empirical implementation of the inflation equation (1). In the following two sections, we outline two approaches—the methodology employed by Woodford (2001), and the approach taken by Sbordone (1998).

3 The Present-Value Method

The construction of an empirical series for inflation consistent with equation (1) requires some characterization of how agents formulate expectations of future price inflation, which in turn requires a specification of the the behavior of real marginal cost mc_t. In explaining the two methods that we consider in this paper, it is useful to express (log) real marginal cost in terms of nominal marginal cost n_t and the log of the price level p_t:

$$mc_t \equiv n_t - p_t. \tag{2}$$

Both approaches take equation (1) as their point of departure. Where they differ, however, is in the assumption each makes regarding the effect that prices have on nominal marginal cost n_t—which in turn has a profound impact on the implied process for mc_t and the resulting closed-form expression for inflation.

If there is a direct relationship between n_t and p_t such that real marginal cost mc_t follows a process that is exogenous to the price level or inflation rate, we can obtain a closed-form solution for inflation by using repeated substitutions to solve equation (1) forward, as follows:

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t mc_{t+k}. \tag{3}$$

Here, inflation is a function of a discounted sum of current and expected future real marginal costs.

The notion that the new-Keynesian Phillips curve implies this closed-form representation for inflation has been noted previously by a number of researchers, including Galí and Gertler (1999) and Goodfriend and King (2001). Intuitively, this expression reflects the fact that standard monopolistic-competition models imply that in the absence of frictions, firms would like to maintain a constant markup of prices over nominal marginal cost. Once frictions (sticky prices) are introduced, then the expectation that nominal marginal cost will rise relative to the aggregate price level in the future leads to higher price inflation today because firms attempt to offset, on average, the anticipated future erosion of their markup.

Empirical implementation of this equation requires an explicit calculation of the discounted sum of expected future real marginal costs. The procedure adopted by Woodford (2001) involves estimating an empirical process for mc_t, which allows him to express its expected future values in terms of variables observed today. Specifically, he defines mc_t as one of the variables in a multivariate VAR of the form

$$Z_t = AZ_{t-1} + \epsilon_t, \tag{4}$$

which implies that the vector of discounted sums of the variables in the VAR can be written as $e_i' (I - \beta A)^{-1} Z_t$ (where e_i' is a unit vector that extracts the discounted sum of our marginal cost proxy).[3] Given this discounted sum, we can then choose the value of γ that yields the best-fitting inflation series.

Results Using Baseline VARs: We start by considering the GDP gap version of the model, in which movements in real marginal cost mc_t are assumed to be proportional to the (log-) deviation of real output from a deterministic trend. Our data are defined for the U.S. nonfarm business sector, and cover the period 1960:Q1 to 2001:Q1. The labor income share (the series plotted in the upper panel of Figure 1) is defined as the ratio of nominal compensation to total nominal output, while the GDP gap (shown in the lower panel of Figure 1) is defined as the quadratically detrended log of real nonfarm GDP. Our VAR specification is informed by the work of Sbordone (2001), who uses a three-variable VAR consisting of current and lagged quarterly values of the labor income share, detrended GDP, and unit labor cost inflation.[4] This system is then employed to generate a discounted sum of current and

[3]This 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.

[4]See footnote 11 on page 6 of Sbordone (2001). Because equation (3) is actually derived as a loglinear approximation about a steady state, we include constant terms in the VAR and estimation

predicted detrended GDP values (in constructing the discounted sum, we assume a value for β of 0.99, but our conclusions are robust to the use of other values).

The results from the GDP gap version of this exercise are plotted in the upper panel of Figure 2; they are essentially identical to the results in Woodford (2001). The discounted sum of current and expected future values of detrended output does not do a good job of explaining inflation; indeed, this present-value series is *negatively* correlated with inflation (hence, in Figure 2 we follow Woodford in multiplying it by an arbitrary positive constant). It is quite apparent from the figure that the model completely fails to predict the high inflation rates of the 1970s, or the low inflation rates of the 1990s. Importantly, we have found that this conclusion—that the expected discounted sum of GDP gaps does poorly in explaining inflation—is robust across a wide range of VAR specifications.

Since the labor income share is among the variables in our VAR, it is a simple matter to use it in order to estimate the labor's share version of equation (3). The resulting inflation series is plotted in the lower panel of Figure 2. The performance of this variant of the model is slightly better than the GDP-gap version inasmuch as the expected discounted sum of labor income shares has the positive correlation with inflation that theory predicts. However, the model explains only a tiny fraction of the variation in inflation—the R^2 for the model is 0.01. Evidently, the new-Keynesian model fits badly no matter which choice of marginal cost proxy we use.

Our finding that the labor's share version of the new-Keynesian Phillips curve fits poorly might be considered surprising in light of the graphical evidence presented in Woodford (2001), which indicates that this version of the model tracks inflation relatively well. It turns out that the reason for this discrepancy stems from the use of a different VAR system to fit this version of the inflation equation. When calculating the discounted sum of current and future labor income shares, Woodford employed a different VAR system from the one used to calculate the present value of the GDP gap. Specifically, the model used by Woodford to construct forecasts of the labor income share was a *bivariate* VAR containing labor's share and nominal unit labor cost growth.[5] If we instead follow this procedure, we obtain a fitted

equations, and express all variables as logs or log-differences.

[5]Note that the specific procedure employed by Woodford is not discussed in his 2001 paper; we thank Professors Woodford and Sbordone for clarifying the details of these calculations in a set of personal communications.

inflation series (which we plot in Figure 3) that does track actual inflation more closely—the R^2 for this version of the model equals 0.44.

An immediate conclusion that can be drawn from these exercises is that the fit of the labor's share version of the new-Keynesian Phillips curve appears to be highly sensitive to how one specifies the forecasting VAR. However, in experimenting with various VAR specifications we have found that most deliver an expected present discounted sum of labor income shares that has a very low correlation with observed inflation. Table 1 reports results based on a number of different VAR systems, including the specification used to generate our Figure 2 (the fourth column of the table) and the bivariate VAR employed by Woodford (second column). The additional variables that we consider in the VAR models—detrended hours and the consumption-output ratio—are informed by the discussion in Sbordone (2001).

Several results from Table 1 are worth noting.

- First, excluding the GDP gap from the labor's share equation in the three-variable VAR—which is necessary in order to obtain the fitted inflation series plotted in Figure 3—is strongly rejected on statistical grounds (see column 4).

- Second, the improvement in inflation fit that occurs when we use the VAR described in column 2 of the table stems from the small, positive, and statistically insignificant coefficient that unit labor cost growth receives when the GDP gap is excluded from the VAR. This allows contemporaneous unit labor cost growth to receive a positive weight in the expression for the discounted sum of future labor income shares, and it is this term that accounts for more than two-thirds of the inflation equation's improved fit.

- In general, the results from the table indicate that the labor's share variant of the new-Keynesian Phillips curve explains a relatively small fraction of the observed variation in inflation.

- Finally, note that the poor fit of the model holds even for the VAR specifications summarized in columns 4 through 7—all of which include unit labor cost growth—and also holds if we employ VAR specifications with forecasting equations for labor's share that fit as well as or better than the equation from Woodford's bivariate VAR.

Results from VARs Including Inflation: While the VAR models used in the preceding analysis all exclude inflation from the system, there is no good reason to do so. In fact, according to equation (3), inflation contains important information about agents' expectations for future values of real marginal cost. If, as seems likely, the information set used by agents when formulating their expectations is larger than the set of variables we have used in our VAR, then including lagged inflation in the VAR should improve our ability to forecast the real marginal cost proxy.[6]

Table 2 summarizes the results that obtain from VAR models that include inflation. Importantly, we find no evidence that inflation Granger causes labor's share: For the bivariate model, the p-value for an F-test of the hypothesis that lagged inflation can be excluded from the labor's share equation equals 0.199. This is problematic for the joint hypothesis that real marginal costs are well captured by the labor income share and inflation is characterized by the new-Keynesian Phillips curve, since the most basic prediction of this model is that real marginal costs should be Granger caused by inflation.

In terms of the fit of the present-value series, the results in Table 2 are generally similar to those in Table 1: In most cases, the expected present value series explains only a tiny fraction of the observed variation in inflation. One new result is worth noting—namely, the fitted inflation series derived from a bivariate VAR in inflation and labor's share receives an R^2 of 0.415. However, almost all of this fit comes from the fact that the expression for the expected present value of current and future labor income shares places small positive weights on both lagged and contemporaneous inflation, the very variable we are attempting to explain. Since the Granger causality tests indicate that there is no statistical reason to include inflation in the VAR for forecasting the labor share, there is also no statistical reason to prefer the fitted inflation series with an R^2 of 0.415 to the other series with much poorer fits.

Comparisons with Reduced-Form Inflation Equations: Even if we view the results based on our best-fitting VARs (the bivariate VARs in s_t and Δulc_t or s_t and

[6]This is a common feature of a number of rational expectations models whose solution takes the form of equation (3), such as the permanent-income model of consumption and the expectations theory of the term structure; see Campbell and Deaton (1989) for a representative example.

π_t) as valid (that is, we accept the fitted inflation series in Figure 3 as providing a fair representation of the fit of the labor's share variant of the new-Keynesian model), there are additional, more substantive reasons to doubt that this version of the model provides an adequate characterization of inflation dynamics.

As an empirical matter, U.S. inflation dynamics are well represented by a reduced-form (traditional) Phillips curve of the form

$$\pi_t = \alpha y_t + A(L)\pi_{t-1}, \tag{5}$$

where y_t is usually defined to be the GDP gap or a related measure. Estimates of this reduced-form equation invariably find that the sum of the coefficients on lagged inflation is large (typically around 0.9, and often statistically indistinguishable from one). Hence, if the new-Keynesian model (1) is the correct *structural* description of inflation dynamics, then the interpretation that we must give to the role played by lagged inflation in the empirical model (5) is that lagged inflation is serving as a proxy for $E_t\pi_{t+1}$—or, more precisely, as a proxy for the discounted sum of expected current and future real marginal costs that is the true determinant of current inflation. As a result, if the new-Keynesian model is correct, there should be little role for lagged inflation in an equation like

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t mc_{t+k} + B(L)\pi_{t-1}. \tag{6}$$

In practice, however, this turns out not to be the case. Even the discounted sum terms that, on their own, generate the best-fitting inflation series—*i.e.*, those based on the $(s_t, \Delta ulc_t)$ or (s_t, π_t) VARs—do little to reduce the sum of the coefficients on lagged inflation in equation (6). Indeed, that the sums are affected at all probably reflects the fact that these expected present value measures place some weight on current and lagged inflation or current and lagged unit labor cost growth, both of which are highly (or perfectly!) correlated with inflation.[7] It is also worth noting

[7] In the discussion of this section, we have ignored the fact that the present value estimation methodology suffers from a generated regressor problem in that the expectation of the discounted sum of current and future real marginal costs is proxied for with a VAR-based forecast. In related work, Rudd and Whelan (2001) use instrumental-variables methods—which address the generated regressor problem—to assess how much of the importance of lagged inflation in empirical inflation regressions comes from its proxying for the discounted sum, and conclude that the inclusion of the discounted sum does little to change the role played by lagged inflation.

that even the best-fitting inflation series in Tables 1 and 2 fit far less well than the inflation series that are generated by simple regressions of inflation on its own lagged values (which typically receive an R^2 of 0.7 and above).

On balance, then, a closer examination suggests that the empirical results in Woodford (2001) do not provide persuasive support for his claim that "real unit labor cost is a much better measure of the true output gap, at least for purposes of explaining inflation variation" in a forward-looking model. We therefore turn to a consideration of a second piece of evidence on this point that has been cited by Woodford and others, namely, the results reported by Sbordone (1998).

4 Sbordone's Methodology

The empirical methodology of Sbordone (1998) proceeds from the assumption that nominal marginal cost, n_t, evolves independently of current, past, and expected future values of the price level. To illustrate the effect that this has on the solution for price inflation, we start with the observation that equation (1) can be re-written as

$$p_t - p_{t-1} = \beta E_t p_{t+1} - \beta p_t + \gamma n_t - \gamma p_t. \tag{7}$$

This equation is, of course, algebraically equivalent to equation (1). However, the assumption that n_t—not mc_t—is the exogenous "forcing variable" profoundly changes the nature of the solution to the model. The assumption of rational expectations implies that agents understand that, in this case, real marginal cost $n_t - p_t$ is jointly determined by the exogenous process for n_t and the endogenous behavior of price setters. Rational price-setters take into account the effect that a higher current price level has on real marginal cost, and the simultaneous feedback effect this has on the price level itself. Hence, under this alternative assumption about the determination of nominal marginal cost, the new-Keynesian pricing equation changes from a first-order stochastic difference equation in *inflation* with *real* marginal cost as the forcing variable, to a second-order difference equation in the *price level* with *nominal* marginal cost as the forcing variable.

Under the assumption that n_t evolves exogenously, standard techniques yield the following closed-form solution for the price level:

$$p_t = \lambda_1 p_{t-1} + (1 - \lambda_1) \left[(1 - \lambda_2) \sum_{i=0}^{\infty} \lambda_2^i E_t n_{t+i} \right], \tag{8}$$

where λ_1 and λ_2 are obtained from the roots of the characteristic equation of (7).[8]

The intuition behind this relationship is simple, though it differs from the intuition given for equation (3). Once again, the starting point is the idea that firms want to maintain a markup of prices over nominal marginal cost that is, on average, as close as possible to its constant optimal frictionless value. Given that nominal marginal cost follows an exogenous process, the decision rule that satisfies this criterion involves the price level's moving each period toward a markup over a moving average of current and expected future nominal marginal costs. It is important to note that if we use the labor income share as our proxy for real marginal cost under this setup, then the implicit assumption that we are making by applying this method is that nominal unit labor costs evolve independently of the price level. Similarly, if we instead assume that movements in detrended real output capture changes in real marginal cost, then the application of this method relies on the assumption that nominal output evolves exogenously to the price level (this could occur, for example, if the monetary authority targeted nominal output growth).

As with the methodology discussed in the previous section, the empirical implementation of this approach involves specifying a process for the exogenous driving variable (in this case, n_t); a simulated inflation series can then be estimated by differencing the predicted price series. In her implementation of this method, Sbordone (1998) re-arranges equation (8) to obtain

$$p_t = \lambda_1 p_{t-1} + (1 - \lambda_1) n_t + (1 - \lambda_1) \left[\sum_{i=1}^{\infty} \lambda_2^i E_t \Delta n_{t+j} \right], \tag{9}$$

and then constructs forecasts for Δn_t (the rate of change of nominal marginal cost) using a VAR that includes this variable.[9] If we assume that labor's share is the

[8] If we assume any dependence of n_t on lags or expected leads of the price level and some other exogenous driving term, then the stochastic difference equation for p_t would not be the same as (7). Rather, it would explicitly incorporate the effects of these additional price level terms, and would contain a different forcing variable. In this case, equation (8) would no longer characterize the closed-form solution to the model.

[9] Technically, because the term inside the square bracket in equation (9) starts at $i = 1$, we measure this discounted sum using $A(I - \beta A)^{-1} Z_t$ instead of $(I - \beta A)^{-1} Z_t$.

appropriate real marginal cost proxy, then Δn_t corresponds to the growth rate of (nominal) unit labor costs. This implies that we can use the same three-variable VAR that we used to construct Woodford's fitted inflation series, since this VAR included unit labor cost growth as one of the variables of the system. If we instead assume that detrended real GDP is the candidate proxy for real marginal cost, then we must replace unit labor cost growth with the first difference of the nominal GDP gap (where the nominal GDP gap is defined as detrended log real GDP plus the log of the price level), which is the corresponding Δn_t concept in this case. Once the two measures of the expected discounted sum of Δn_t are in hand, we can then choose the values of λ_1 and λ_2 in equation (9) that yield the best-fitting series for inflation.[10]

The resulting inflation series are plotted in Figure 4; they demonstrate that Sbordone's method produces an inflation series that fits well *no matter which* measure of marginal cost we use. For the labor's share version of the model (the upper panel of Figure 4), we find $\lambda_1 = 0.77$ and $\lambda_2 = 0.72$; the R^2 for the fitted inflation series is 0.80. Likewise, for the model that uses detrended GDP (the lower panel), we have $\lambda_1 = 0.92$ and $\lambda_2 = 0.94$, with an R^2 for the fitted inflation series of 0.73. While the labor's share version of the inflation series fits slightly better than the GDP gap version, the principal message of these figures is clearly that both series fit well. (This result—that the predicted inflation series fit well when either measure of marginal cost is used—is robust across various specifications of the VAR system.)

We conclude, then, that the fact that we can obtain a good fit for inflation under Sbordone's methodology when labor's share is equated with real marginal cost—the finding cited by Woodford—should not be considered compelling evidence that the labor income share is a substantially superior proxy for real marginal cost. Combined with the results of the previous section, it is apparent that the success or failure of the new-Keynesian Phillips curve in fitting actual inflation does not appear to significantly depend on the choice of real marginal cost proxy.

While our principal purpose is to assess the relative merits of detrended GDP

[10]This differs slightly from Sbordone (1998), who chooses these parameters to maximize the fit of the simulated price-unit labor cost ratio (*i.e.*, the inverse of the labor share). While we consider our choice of estimation to be somewhat more natural in the context we are discussing, our point—that the fit of both marginal cost proxies is good when this method is used—holds just as well if we use her approach to estimate λ_1 and λ_2.

and the labor income share as proxies for real marginal cost, it is also worth asking whether the impressive fit that obtains under Sbordone's methodology can be taken as empirical evidence in favor of the new-Keynesian Phillips curve. On balance, we think not. On theoretical grounds, we view the key assumption underlying her method (that price-setters expect nominal marginal cost to evolve exogenously vis-à-vis the price level) as providing an unappealing description of how the measures of nominal marginal cost that we have examined are determined. That price-setters would expect nominal unit labor costs to evolve independently of the price level appears to run counter to the idea that workers bargain in terms of real wages. And, while one might invoke nominal-income targeting by the central bank in order to motivate the notion that nominal output is exogenous, such a policy rule does not provide an accurate description of how U.S. monetary policy is actually practiced.

Finally, it is also of interest to ask why Sbordone's methodology yields an inflation series that fits so well. Mechanically, the price-level equation that underpins her approach (equation 9) implies an inflation equation of the form

$$\pi_t = \lambda_1 \pi_{t-1} + (1-\lambda_1)\,\Delta n_t + (1-\lambda_1)\left[\sum_{i=1}^{\infty} \lambda_2^i E_t \Delta n_{t+j} - \sum_{i=1}^{\infty} \lambda_2^i E_{t-1} \Delta n_{t+j-1}\right], \quad (10)$$

in which inflation is related to its own lag, unit labor cost growth, and a term that is intended to capture updates to agents' expectations of future unit labor cost growth (the expression in square brackets). Seen in this light, it should not be surprising that this estimation method yields a well-fitting inflation series. In practice, even a single lag of inflation explains a large fraction of the variation in the series, and unit labor cost growth also contains some incremental explanatory power for inflation. Indeed, by themselves these terms can explain more than 72 percent of the variation in inflation over this period, implying that the term in square brackets—which truly distinguishes this equation from traditional reduced-form empirical inflation equations—contains only a small amount of incremental explanatory power for inflation. (Moreover, one should be careful in interpreting even this small amount of additional explanatory power as evidence of the type of rational forward-looking behavior underlying the new-Keynesian Phillips curve, because this equation can only be considered an accurate representation of the model under the unrealistic identifying assumption that is used in its derivation.)

5 Conclusions

In this paper, we have critically assessed the claim that the new-Keynesian Phillips curve performs poorly when detrended real GDP is used as the driving variable, but fits well when real unit labor costs (labor's share of income) is used. We find that the robust conclusion that emerges is that *neither* variable allows the new-Keynesian model to fit well unless a highly unrealistic assumption is used to derive the estimation equation; in this latter case, either variable works well.

Our relatively negative assessment of the new-Keynesian Phillips curve is closely related to Fuhrer and Moore's (1995) critique of standard sticky-price models, which highlights the inconsistency between the forward-looking new-Keynesian inflation equation and the empirical finding that lags of inflation play an important role in inflation regressions. One way to reconcile these findings would be to show that lagged inflation proxies for future values of the output gap; alternatively, one could follow Galí and Gertler (1999) and Goodfriend and King (2001) and argue that these lags of inflation proxy for expectations of future labor shares. However, the evidence presented in this paper suggests that neither possibility is correct. We find no evidence that inflation Granger causes the labor share of income, and the discounted sum of current and expected future labor shares generally explains very little of the empirical variation in inflation.

Thus, we believe that the evidence provides a firm answer to the question posed in the title of our paper: There does not appear to be a strong case for including the labor income share in a monetary policy rule. Indeed, given the historical behavior of labor's share, there are compelling reasons not to associate this series with the "output gap"—for a start, doing so necessarily implies that every postwar U.S. recession has actually been a *boom* relative to the prevailing level of potential output. It seems unlikely to us that even those who believe in an important role for technology shocks in driving business cycles would defend this view.

Finally, we note that our conclusions should not be interpreted as implying that forward-looking inflation models based on real marginal cost cannot work, inasmuch as both driving variables considered here may be very poor proxies for marginal cost. For example, Rotemberg and Woodford (1999) detail a number of reasons—such as the existence of overhead labor, overtime premia, and adjustment costs for labor—why real marginal cost could be procyclical even though real unit labor costs are not.

Thus, the increases in average cost that are observed during recessions are likely to be poor indicators of marginal cost pressures. On balance, then, we conclude that it remains possible that some forward-looking model based on a measure of real marginal cost provides a good description of the inflation process, but this conjecture can by no means be considered proven.

References

[1] Campbell, John and Angus Deaton (1989). "Why is Consumption So Smooth?" *Review of Economic Studies*, 56, 357-373.

[2] Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans (2001). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," NBER Working Paper No. 8403.

[3] Dotsey, Michael and Robert King (2001). "Pricing, Production, and Persistence," NBER Working Paper No. 8407.

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

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

[6] Goodfriend, Marvin and Robert King (2001). "The Case for Price Stability," NBER Working Paper No. 8423.

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

[8] 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.

[9] Sargent, Thomas J. (1987). *Macroeconomic Theory* (2nd edition), Academic Press.

[10] Sbordone, Argia (1998). "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics*, forthcoming.

[11] Sbordone, Argia (2001). "An Optimizing Model of U.S. Wage and Price Dynamics," March draft, Rutgers University.

[12] Woodford, Michael (2001). "The Taylor Rule and Optimal Monetary Policy," *American Economic Review*, 91(2), 232-237.

Table 1: Results from Alternative VAR Forecasting Models for Labor's Share

	VAR specifications						
	$[s_t]$	$\begin{bmatrix} s_t \\ \Delta ulc_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \\ h_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \\ c_t/y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \\ h_t \\ c_t/y_t \end{bmatrix}$
	A. R^2 from inflation equation						
	0.162	0.437	0.129	0.014	0.040	0.040	0.001
	B. Exclusion restriction p-values (s_t equation)						
s_t	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Δulc_t		0.288		0.901	0.483	0.559	0.480
y_t			0.000	0.000	0.019	0.018	0.614
h_t					0.252		0.133
c_t/y_t						0.118	0.063
	C. $\bar{R}^2$ from labor's share VAR equation						
	0.848	0.849	0.864	0.862	0.863	0.864	0.866

Key: $s_t \equiv$ labor's share, $\Delta ulc_t \equiv$ unit labor cost growth, $y_t \equiv$ detrended real GDP, $h_t \equiv$ detrended hours, $c_t/y_t \equiv$ detrended consumption-output ratio. See text for additional details.

Table 2: Results from VAR Forecasting Models That Include Inflation

	VAR specifications						
	$\begin{bmatrix} s_t \\ \pi_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ \Delta ulc_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ \Delta ulc_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ \Delta ulc_t \\ y_t \\ h_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ \Delta ulc_t \\ y_t \\ c_t/y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \pi_t \\ \Delta ulc_t \\ y_t \\ h_t \\ c_t/y_t \end{bmatrix}$
	A. R^2 from inflation equation						
	0.415	0.364	0.001	0.026	0.113	0.121	0.040
	B. Exclusion restriction p-values (s_t equation)						
s_t	0.000	0.000	0.000	0.000	0.000	0.000	0.000
π_t	0.199	0.401	0.921	0.631	0.501	0.925	0.965
Δulc_t		0.518		0.635	0.564	0.662	0.555
y_t			0.000	0.000	0.018	0.019	0.372
h_t					0.463		0.273
c_t/y_t						0.297	0.186
	C. $\bar{R}^2$ from labor's share VAR equation						
	0.849	0.849	0.862	0.863	0.863	0.864	0.864

Key: $s_t \equiv$ labor's share, $\pi_t \equiv$ inflation, $\Delta ulc_t \equiv$ unit labor cost growth, $y_t \equiv$ detrended real GDP, $h_t \equiv$ detrended hours, $c_t/y_t \equiv$ detrended consumption-output ratio. Lag lengths chosen using Schwarz criterion. See text for additional details.

Figure 1
Output Gap Concepts, U.S. Nonfarm Business Sector
(NBER Recession Dates Shaded)

A. Labor Income Share

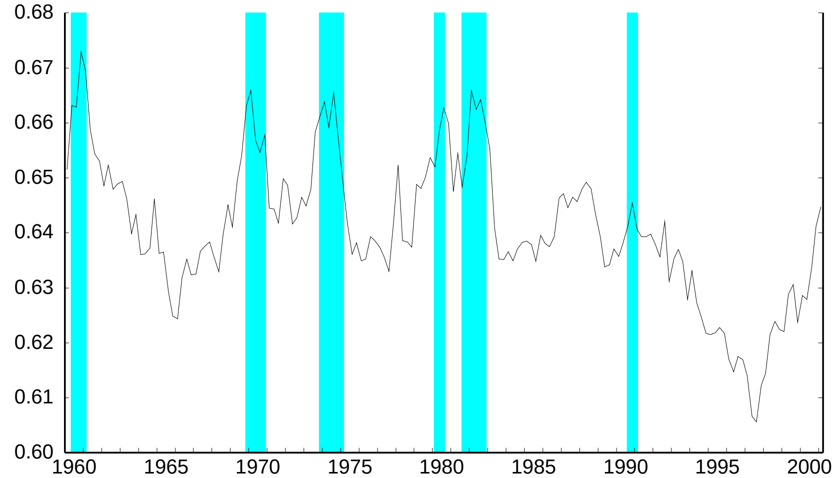

B. Quadratically Detrended Log GDP

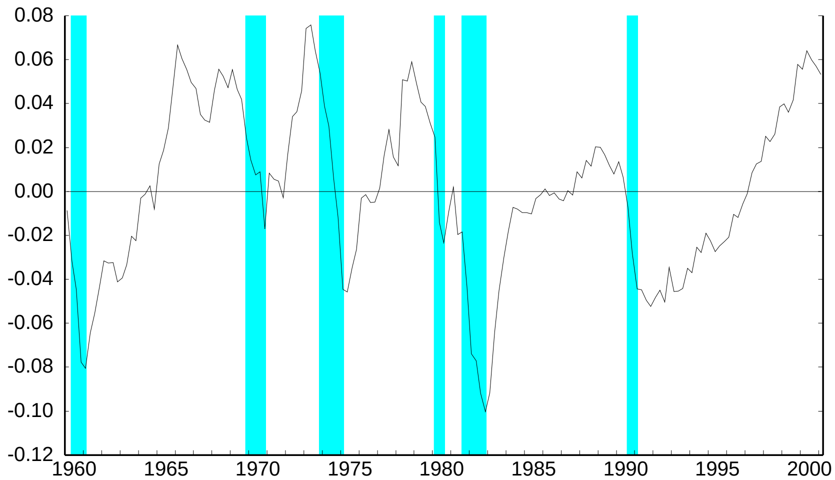

Figure 2
Actual and Predicted Inflation--Present-Value Method
(VAR models include GDP gap, labor's share, and unit labor cost growth)

A. Present Value of GDP Gaps from VAR System

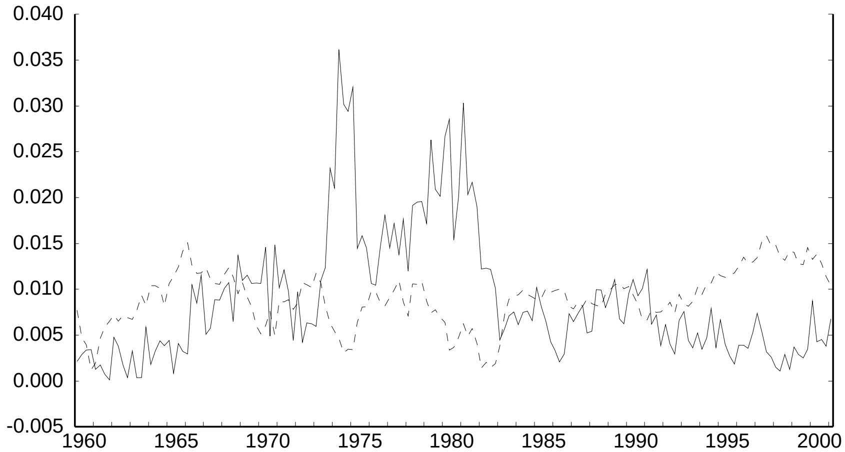

B. Present Value of Labor Income Shares from VAR System

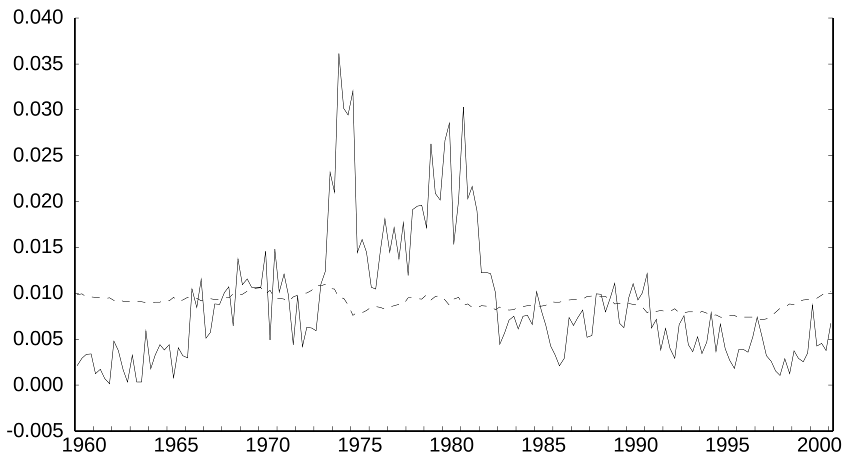

Figure 3
Actual and Predicted Inflation--Present-Value Method (alt. VAR)
(VAR model includes labor's share and unit labor cost growth only)

Present Value of Labor Income Shares from VAR System

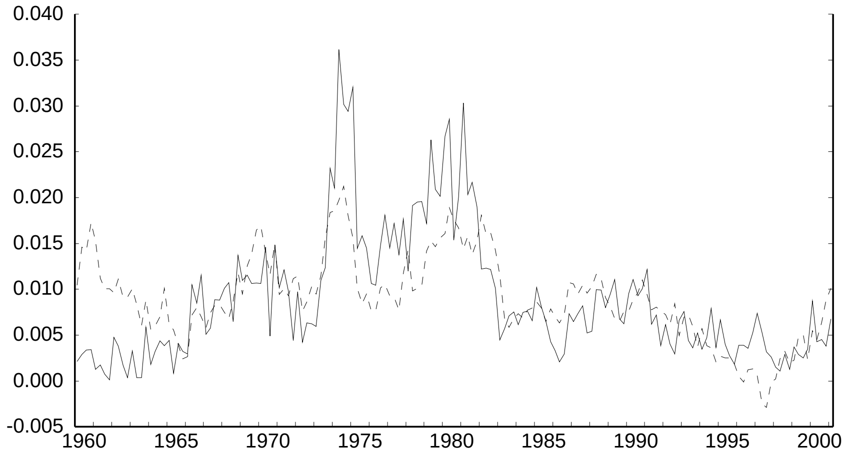

Figure 4
Actual and Predicted Inflation--Sbordone Method
(VAR models include GDP gap, labor's share, and ULC or nominal GDP gap growth)

A. Price-Level Equation Using Expected ULC Growth

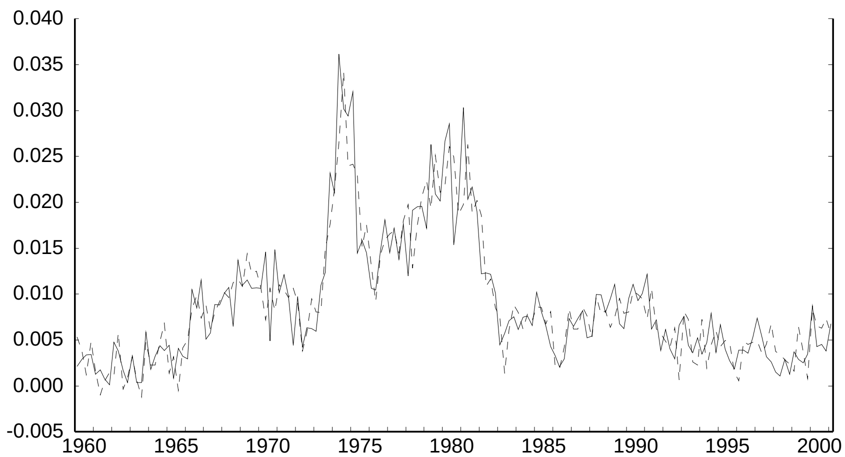

B. Price-Level Equation Using Expected Nominal GDP Gap Growth