Research Repository UCD

Title	Does labor's share drive inflation?
Authors(s)	Rudd, Jeremy, Whelan, Karl
Publication date	2005-04
Publication information	Rudd, Jeremy, and Karl Whelan. "Does Labor's Share Drive Inflation?" 37, no. 2 (April, 2005).
Publisher	Blackwell - published on behalf of The Ohio State University
Item record/more information	http://hdl.handle.net/10197/243

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

JEREMY RUDD
KARL WHELAN

Does Labor's Share Drive Inflation?

A number of researchers have recently argued that the new-Keynesian Phillips curve matches the empirical behavior of inflation well when the labor income share is used as a driving variable, but fits poorly when deterministically detrended output is used. The theoretical motivation for these results rests on the idea that the output gap—the deviation between actual and potential output—is better captured by the labor income share, in turn implying that central banks should raise interest rates in response to increases in this variable. We show that the empirical evidence generally suggests that the labor share version of the new-Keynesian Phillips curve is a very poor model of price inflation. We conclude that there is little reason to view the labor income share as a good measure of the output gap, or as an appropriate variable for incorporation in a monetary policy rule.

JEL code: E31

Keywords: Phillips curve, sticky prices, marginal cost, output gap.

RECENT YEARS have seen an explosion in research aimed at assessing monetary policy rules using macroeconomic models built from explicit microfoundations. In many versions of these models, pricing behavior is described by a “new-Keynesian Phillips curve,” which relates inflation to expected future inflation and the output gap x_t :

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

Under rational (i.e., model-consistent) expectations, this equation can be re-written as

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

The authors are grateful to Dale Henderson, John Leahy, Argia Sbordone, Michael Woodford, and two anonymous referees for helpful comments on earlier drafts of the paper. The views expressed are our own and do not necessarily reflect the views of the Board of Governors, the staff of the Federal Reserve System, or the Central Bank of Ireland.

JEREMY RUDD is an economist in the Division of Research and Statistics, Federal Reserve Board (E-mail: jeremy.b.rudd@frb.gov). KARL WHELAN is a senior economist at the Central Bank of Ireland (E-mail: karl.wheLAN@centralbank.ie.).

Received December 9, 2002; and accepted in revised form April 21, 2003.

Journal of Money, Credit, and Banking, Vol. 37, No. 2 (April 2005)

Copyright 2005 by The Ohio State University

where this latter expression is obtained from applying repeated substitution to Equation (1). Hence, the new-Keynesian Phillips curve implies that inflation is a purely forward-looking “jump” variable.

A key shortcoming of this model is that empirical implementations that use deterministically detrended output to measure the output gap have been shown to provide a poor description of the actual inflation process. For example, as Fuhrer and Moore (1995) have noted, the model’s prediction that inflation is a jump variable appears inconsistent with the empirical evidence from reduced-form inflation regressions, which indicates that inflation depends importantly on its own lags. Moreover, the empirical role played by these lags cannot be explained by their acting as proxies for expectations of future values of detrended output: The prediction that higher inflation Granger causes increases in detrended output is firmly rejected by the data.

One potential explanation for the empirical failure of the new-Keynesian Phillips curve relies on the idea that the theory is correct but that deterministically detrended GDP is an inappropriate measure of the output gap for the purposes of explaining inflation. In theoretical terms, the sticky-price models that underpin the new-Keynesian Phillips curve imply that the correct driving variable in this equation is actually real marginal cost (nominal marginal cost divided by the price level). The use of detrended output by previous authors stems from the assumption that firms’ 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. However, to the extent that y_t^* does not evolve according to a deterministic trend, such an approach could be problematic.

As an alternative to the standard approach, Galí and Gertler (1999), Sbordone (2002), and others have suggested using average unit labor costs (nominal compensation divided by real output) as a proxy for nominal marginal cost.¹ The resulting proxy for *real* marginal cost—and thus the output gap—is the labor share of income (nominal compensation divided by nominal output).

In support of this alternative approach, a number of researchers have recently presented evidence that the labor share variant of the new-Keynesian Phillips curve provides a good empirical fit for inflation. For example, Woodford (2001) argues that the labor share “is a much better measure of the true output gap, at least for purposes of explaining inflation variation.” Were this conclusion true, it would carry substantial implications for the conduct of monetary policy. For example, it would imply that central banks should not follow traditional variants of the Taylor rule, in which short-term interest rates are set with reference to the levels of inflation and deterministically detrended output. Rather, the reformulated version of the Taylor rule suggested by these results would call for the monetary authority to ignore standard measures of the output gap and instead raise interest rates in response to increases in labor’s share of income. The unconventional nature of such a policy

1. This can be justified if the production technology exhibits a constant elasticity of output with respect to labor (as will be the case if the production function is Cobb-Douglas).

rule is clear from Figure 1: Because the labor share has spiked upward in every modern U.S. recession, a policy rule based on this variable instead of detrended output would have called for higher interest rates during each of these episodes.

In this paper, we present new evidence on the merits of the labor share version of the new-Keynesian Phillips curve. We conclude that the empirical case for this model is, in fact, very weak.

In the first part of the paper, we use VAR-based methods to assess the empirical fit of Equation (2). Specifically, we use a reduced-form VAR to calculate $E_t x_{t+k}$ and then compare a weighted sum of these terms with actual inflation. This approach is identical to the one followed by Woodford (2001); however, we come to very different conclusions regarding the merits of the labor share model. In particular, Woodford reports that a predicted inflation series based on detrended output is negatively correlated with actual inflation, while a series based on the labor share fits well. However, we show that this latter result is not robust inasmuch as the fit of the labor share version of the model is highly sensitive to small changes in the VAR that is used to forecast future values of the labor income share. For a broad range of VAR specifications, the model's fit is actually very poor. In addition, we show that inflation does not contain useful information for forecasting future values of the labor share, which implies that one can reject the hypothesis that the labor share version of the new-Keynesian Phillips curve can account for the role of lagged inflation in reduced-form inflation regressions.

The second part of the paper reconciles our conclusion that the labor share version of the model fits poorly with the evidence presented by Sbordone (2002), who concludes that the model fits well. The difference turns out to lie in Sbordone's use of a procedure that does not involve directly testing Equation (2). We show, however, that her procedure yields a good fit for inflation even in cases where inflation is determined by a backward-looking model (and so the new-Keynesian Phillips curve is clearly false). Relatedly, we demonstrate that her method yields an inflation series that fits well regardless of which output gap proxy is used. Thus, Sbordone's procedure has no power against relevant alternatives, and her results cannot be interpreted as providing useful evidence in favor of the new-Keynesian model.

We conclude that the new-Keynesian Phillips curve provides a poor description of the inflation process even when the labor income share is used as a proxy for real marginal cost. As a result, the relative fits of the various versions of the new-Keynesian Phillips curve should not be considered an appropriate metric for determining whether the true output gap is better measured by the labor share or by detrended output, and certainly cannot support the conclusion that the labor income share should supplant conventional gap measures in monetary policy rules.

1. INFLATION AND EXPECTATIONS OF REAL MARGINAL COST

In this section, we describe empirical implementations of Equation (2)—which expresses inflation as a function of a discounted sum of current and expected future

A. Labor Income Share

B. Quadratically Detrended Log GDP

FIG. 1. Output Gap Concepts, U.S. Nonfarm Business Sector (NBER recession dates shaded)

output gaps—using both the labor share and detrended output as proxies for the output gap. Our data are for the U.S. nonfarm business sector, and cover the period 1960:Q1 to 2001:Q1. The labor income share is plotted in the upper panel of Figure 1. Our detrended output series (shown in the lower panel of Figure 1) is defined as the deviation of the log of real nonfarm GDP from a quadratic trend.

The construction of an empirical inflation series consistent with Equation (2) requires some characterization of how agents formulate expectations of future values of x_t . The procedure adopted by Woodford (2001) involves specifying x_t as one of the variables in a multivariate VAR of the form

$$Z_t = AZ_{t-1} + \varepsilon_t. \quad (3)$$

This allows expected future values to be expressed in terms of variables observed today. Specifically, the discounted sum of the output gap proxy can be written as $e'_i(I - \beta A)^{-1}Z_t$, where e'_i is a suitably defined unit vector.² Given this discounted sum, we can then obtain the value of γ that yields the best-fitting inflation series by regressing actual inflation on the discounted sum.³ (Hence, a different value of γ is chosen for each of the estimated discounted sums that we consider.)

1.1 Results Using Baseline VARs

We start with the results for the detrended output version of the model. The VAR specification that we initially consider is identical to the one used by Woodford (2001) for this exercise; it consists of current and lagged values of detrended output, the labor income share, and unit labor cost inflation.⁴ This system is then employed to generate a discounted sum of current and predicted values of detrended output (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 this exercise are plotted in the upper panel of Figure 2; they are essentially identical to results reported by Woodford for this version of the model. The expected discounted sum of detrended output values does a very poor job of explaining inflation; indeed, this series is actually *negatively* correlated with inflation (hence, in Figure 2 we multiply it by an arbitrary positive constant). It is apparent from the figure that this model completely fails to predict the high inflation rates of the 1970s, or the low inflation rates of the 1990s. Moreover, this conclusion—that the expected discounted sum of detrended output values does poorly in explaining inflation—is robust across a wide range of VAR specifications.

Because the labor income share is among the variables included in the VAR, it is a simple matter to use this same system to construct the expected discounted sum

2. 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.

3. Because Equation (2) is actually derived as a loglinear approximation about a steady state, we include constant terms in the VAR and estimation equations, and express all variables as logs or log-differences.

4. The specific VAR system employed was not explicitly discussed in Woodford's (2001) paper (the specification used follows a similar exercise in Sbordone, 2001). We thank Professors Woodford and Sbordone for clarifying the details of these calculations in a series of personal communications.

A Present Value of GDP Gaps from VAR System

B Present Value of Labor Income Shares from VAR System

FIG. 2. Actual and Predicted Inflation, Present-value Method. (VAR models include GDP gap, labor's share, and unit labor cost growth.) Actual inflation given by solid line; predicted inflation given by dashed line

Present Value of Labor Income Shares from VAR System

FIG. 3. Actual and Predicted Inflation, Present-value Method with Alternative VAR (VAR model includes labor's share and unit labor cost growth only.) Actual inflation given by solid line; predicted inflation given by dashed line

of labor shares. The resulting inflation series is plotted in the lower panel of Figure 2. The performance of this variant of the model is perhaps slightly better than the GDP-gap version inasmuch as the discounted sum of labor shares has the positive correlation with inflation predicted by the theory. However, the model explains only a tiny fraction of the variation in inflation—the R^2 for the model is 0.01.

How can this finding be reconciled with the evidence presented in Woodford's paper, which indicated that a discounted sum of labor shares tracks inflation relatively well? It turns out that the reason for this discrepancy stems from Woodford's use of a different VAR system to fit the labor share 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 discounted sum of detrended output values; specifically, the system used in the former case was a *bivariate* VAR containing the labor share and nominal unit labor cost growth. If we instead follow this procedure, we also obtain a fitted inflation series (plotted in Figure 3) that tracks actual inflation more closely—the R^2 is 0.44 for this version of the model.⁵

An immediate conclusion that can be drawn from these exercises is that the fit of the labor share version of the new-Keynesian Phillips curve appears to be highly

5. This fit is a bit lower than that reported in Woodford's paper for the same exercise. We believe the difference reflects our use of a more recent vintage of data; in particular, our dataset incorporates the substantial revisions made by the Bureau of Labor Statistics to the published labor income shares for the late 1990s.

TABLE 1
RESULTS FROM ALTERNATIVE VAR FORECASTING MODELS FOR LABOR'S SHARE

	VAR specifications						
	$[s_t]$	$\begin{bmatrix} s_t \\ \Delta ulc_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ y_t \\ h_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \\ c_t/y_t \end{bmatrix}$	$\begin{bmatrix} s_t \\ \Delta ulc_t \\ y_t \\ h_t \\ c_t/y_t \end{bmatrix}$
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 Share VAR Equation	0.848	0.849	0.864	0.862	0.863	0.864	0.866
D. Sum of Lagged π Coefficients in Inflation Equation (If Included)	0.871	0.766	0.855	0.890	0.883	0.882	0.898

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

sensitive to how one specifies the forecasting VAR. However, in experimenting with various VAR specifications we have found that most generate an expected discounted sum of labor income shares that has a very low correlation with inflation. Table 1 reports results based on several different VAR systems, including the bivariate system employed by Woodford (the second column of the table) and the specification used to generate our Figure 2 (the fourth column).⁶ The two other variables that we include in the additional VAR specifications—namely, detrended hours and the consumption-output ratio—are used in the VARs that Sbordone (2001) considers.

Several results from Table 1 are worth noting.

1. Excluding detrended output from the three-variable VAR—which is necessary in order to obtain the well-fitting inflation series shown in Figure 3—is strongly rejected on statistical grounds. Lags of detrended output receive statistically significant coefficients in the labor-share equation (see Column 4).
2. The improvement in fit for inflation that occurs when we use Woodford's bivariate system (described in Column 2) stems from the small, positive coefficients that lagged unit labor cost growth receives when detrended output is excluded from the labor share equation. However, these coefficients are not statistically significant. If we omit unit labor cost growth from the VAR—thereby using a

6. The estimates of γ that we obtain range from essentially 0 to 0.024. (As our focus is on the fit of the models, we do not report specific values of γ in the table; note also that, because the discounted sum is a generated regressor—and because β is treated as known—the OLS standard error for γ will not be correct.)

univariate regression to forecast future labor shares—the model's fitted inflation series has an R^2 of only 0.16 (see Column 1).

3. In general, the labor share variant of the new-Keynesian Phillips curve explains a relatively small fraction of the observed variation in inflation. Moreover, this poor fit obtains even for the VAR specifications summarized in Columns 4 through 7—all of which include unit labor cost growth—as well as for specifications whose forecasting equation for the labor share fits as well as or better than the equation from Woodford's bivariate VAR.

1.2 Results from VARs Including Inflation

While the preceding analysis excluded lagged inflation from the VARs used to forecast future values of the labor share, there is no good reason to do so. In fact, it turns out that if the labor share variant of the new-Keynesian Phillips curve is to successfully explain the observed dynamics of inflation, then it must be the case that lagged inflation is a very useful predictor of the labor income share.

To see why this is so, note that as an empirical matter, U.S. inflation dynamics are well represented by a reduced-form regression of the form

$$\pi_t = A(L)\pi_{t-1} + \alpha y_t, \quad (4)$$

where y_t is usually defined to be detrended output 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 1. Hence, if the labor share version of the new-Keynesian model (Equation 2) is the correct *structural* description of inflation dynamics, then it must be that the role played by lagged inflation in the empirical model (Equation 4) stems purely from its serving as a proxy for expected future values of the labor share (the true determinants of current inflation).

Table 2 summarizes the results obtained from using VAR specifications that include lagged inflation. Importantly, we find no evidence that inflation Granger causes the labor share: For the seven specifications reported here, the lowest p -value for an F -test of the hypothesis that lagged inflation can be excluded from the labor share equation equals 0.199.⁷ In terms of the fit of the new-Keynesian Phillips curve, the results in Table 2 are generally similar to those in Table 1: In most cases, the discounted sum explains only a tiny fraction of the observed variation in inflation.

One new result worth noting from Table 2 is the R^2 of 0.415 obtained by the fitted inflation series derived from the bivariate VAR in inflation and the labor share. Almost all of this fit comes from the fact that, by construction, the series used to measure the discounted sum of labor shares in this case places small

7. These results run somewhat counter to the arguments of Galí and Gertler (1999) and Goodfriend and King (2001), who decide in favor of the hypothesis that inflation leads the labor share. However, neither of these papers report Granger causality tests: Goodfriend and King report coefficients from a VAR in inflation and the labor share, but do not report standard errors, while Galí and Gertler (1999) only report correlations between various leads and lags of inflation and the labor share.

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 Share VAR Equation	0.849	0.849	0.862	0.863	0.863	0.864	0.864
D. Sum of Lagged π Coefficients in Inflation Equation (If Included)	0.780	0.800	0.897	0.887	0.869	0.861	0.883

NOTES: s_t $\equiv$ labor share, π_t $\equiv$ inflation, Δulc_t $\equiv$ unit labor cost growth, y_t $\equiv$ detrended output, 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.

positive weights on both lagged and contemporaneous inflation, the very variable we are attempting to explain. However, because the Granger causality tests indicate that there is no statistical reason to include inflation in this VAR, 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.

These results illustrate a general problem that can arise when the VAR approach to testing rational expectations models is used uncritically. In particular, it is hardly surprising that we can obtain a better-fitting series for predicted inflation in this case by using a linear combination of variables that includes terms that we know, a priori, to be highly correlated with inflation. However, without evidence that these variables also forecast the labor income share, this good fit should not be considered an endorsement of the underlying theoretical model. This insight helps to explain the rather wide variation in fit generated by our various VAR specifications; however, it also underpins our conclusion that the poor performance of the labor share version of the model can be considered a relatively robust finding.

1.3 The Role of Lagged Inflation

As a final way of illustrating the inability of the labor share version of the new-Keynesian Phillips curve to explain inflation dynamics, note that if the model were correct, then there should be little role for lagged inflation in an equation like

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t x_{t+k} + B(L) \pi_{t-1}. \quad (5)$$

In practice, however, this turns out not to be the case: Even the inclusion of the discounted sums 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—does little to reduce the estimated sum of the coefficients on lagged inflation in Equation (5). (These coefficients are reported in Panel D of Tables 1 and 2).⁸

Two additional results are also worth highlighting. First, even the best-fitting inflation series in Tables 1 and 2 fit far less well than simple regressions of inflation on its own lagged values (which typically receive an R^2 of 0.7 and above). Second, the coefficients on detrended output in reduced-form regressions like Equation (4) are also invariably statistically significant. This latter result underscores an important point, which is that the failure of the detrended output version of the new-Keynesian Phillips curve should not, on its own, be used as evidence that traditional proxies for the output gap represent poor measures of resource utilization. The fact that deterministically detrended output works well in empirical reduced-form inflation regressions suggests that it may still provide a useful—albeit somewhat crude—measure of the output gap, and further suggests that the failure of the detrended output variant of the new-Keynesian Phillips curve may merely reflect the inadequacy of the sticky-price framework that underpins the model.

2. COMPARISON WITH SBORDONE'S APPROACH

Sbordone (2002) has examined the same new-Keynesian pricing equation that we consider, but in a somewhat different manner. In contrast to our findings, she concludes that the labor share version of the new-Keynesian Phillips curve provides a good description of empirical inflation dynamics. In this section, we reconcile our findings with hers, and explain why her results do not contradict our own conclusions.

Sbordone begins by rewriting Equation (1) as

$$p_t - p_{t-1} = \beta E_t p_{t+1} - \beta p_t + \gamma n_t - \gamma p_t, \quad (6)$$

where n_t denotes *nominal* marginal cost. She then notes that this equation—which we have treated as a first-order difference equation in inflation with real marginal cost as the driving term—can instead be treated as a second-order difference equation in the price level with nominal marginal cost as the driving term. This price-level equation has a solution of the form:

8. It is instructive to compare these results with the findings of Rudd and Whelan (2001), who estimate Equation (5) using an instrumental variables approach in which the realized present value of x_t is related to a set of predetermined instruments. (This method has the advantage of not requiring us to explicitly specify a process for the driving term x_t .) In line with the conclusions presented here, Rudd and Whelan find that the coefficients on lagged inflation are little affected by including a discounted sum of the driving variable.

$$p_t = \lambda_1 p_{t-1} + (1 - \lambda_1) \left[(1 - \lambda_2) \sum_{i=0}^{\infty} \lambda_2^i E_t n_{t+i} \right], \quad (7)$$

where λ_1 and λ_2 are obtained from the roots of the characteristic equation of Equation (6). In her empirical implementation, Sbordone rearranges Equation (7) to obtain

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

and then constructs forecasts for Δn_t (the rate of change of nominal marginal cost) using a VAR that includes this variable.⁹ Sbordone assumes that the labor income share is the appropriate proxy for real marginal cost—implying that Δn_t corresponds to the growth rate of nominal unit labor costs—and finds that her empirical implementation of Equation (8) produces an inflation series that tracks observed inflation well.

Citing the negative results previously reported by other researchers for the detrended output version of the new-Keynesian pricing equation, Sbordone interprets her equation's ability to fit inflation as indicating that the new-Keynesian model works much better when the labor share is used as an output gap proxy. However, we believe that the results from this exercise should not be considered evidence in favor of the new-Keynesian Phillips curve, nor should they be considered strong evidence in favor of using the labor share measure of the output gap instead of the detrended output measure.

To see why, first observe that the price-level equation, Equation (8), implies an inflation equation of the form

$$\begin{aligned} \pi_t = & \lambda_1 \pi_{t-1} + (1 - \lambda_1) \Delta n_t \\ & + (1 - \lambda_1) \left[\sum_{i=1}^{\infty} \lambda_2^i E_t \Delta n_{t+i} - \sum_{i=1}^{\infty} \lambda_2^i E_{t-1} \Delta n_{t+i-1} \right], \end{aligned} \quad (9)$$

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 is hardly surprising that Sbordone's procedure yields a well-fitting inflation series. As we have noted already, lagged inflation is a highly useful predictor of current inflation—in practice, even a single lag explains a large fraction of the variation in the series. Moreover, contemporaneous unit labor cost growth also contains some incremental explanatory power for inflation. Thus, even if the rational price-setting posited by the new-Keynesian Phillips curve were entirely incorrect—for example, because agents had backward-looking inflation expectations—we would still expect this procedure to give us an empirically reasonable inflation series.

In addition, it is crucial to note that Equations (2) and (9) represent two different ways of describing the same theoretical relationship. *If the new-Keynesian pricing*

9. Technically, because the term inside the square bracket in Equation (8) starts at $i = 1$, we measure this discounted sum using $A(I - \beta A)^{-1} Z_t$ instead of $(I - \beta A)^{-1} Z_t$.

theory were correct, then Equations (2) and (9) would provide exactly the same characterization of the determination of prices. Given that our previous results demonstrate that Equation (2) generally does very poorly as a model of price inflation no matter which proxy for the output gap is used, the model should be viewed as a poor one, irrespective of the fits generated by empirical implementations of Equation (9).¹⁰

Finally, we note that, contrary to previous interpretations, only a very small part of the good fit for inflation obtained under Sbordone's method comes from her use of the labor income share as a proxy for real marginal cost. To illustrate this, we replicate the results from her estimation procedure, and compare them to the results from a parallel exercise in which we use detrended output as the real marginal cost proxy. For the labor share version of the model, nominal marginal cost n_t equals unit labor costs, and so we can use the same three-variable VAR that we employed in the previous section in order to generate forecasts for Δn_t (recall that this VAR included unit labor cost growth as one of the variables in the system). For the variant of the model that assumes real marginal cost to be proportional to detrended output, the corresponding measure of nominal marginal cost equals nominal detrended output (defined as detrended log real GDP plus the log of the price level). Hence, to estimate this version of the model, we remove unit labor cost growth from the VAR system and replace it with the first difference of this alternative n_t measure. 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 (8) that yield the best-fitting series for inflation.¹¹

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. Specifically, for the labor's share version of the model (the upper panel of Figure 4), we obtain an R^2 for the fitted inflation series of 0.80 (with $\lambda_1 = 0.77$ and $\lambda_2 = 0.72$). Likewise, for the model that uses detrended GDP (the lower panel), we obtain an R^2 for the fitted inflation series of 0.73 (with $\lambda_1 = 0.92$ and $\lambda_2 = 0.94$). While the labor share version of the inflation model fits slightly better than the GDP-gap version, the principal message of these figures is clearly that both series fit well.¹² Again, we emphasize that these results in no way contradict the

10. One possible critique of this position would be to argue that the empirical implementation of Equation (2)—which requires using a VAR to forecast future labor shares—is somehow inferior to implementations of Equation (9). However, it appears that the opposite is the case: The R^2 for the labor-share equations in the VARs reported in the previous section are significantly higher than those for nominal unit labor cost inflation, which is the variable being forecasted when Equation (8) or (9) is implemented using the labor share as the output gap proxy.

11. This differs slightly from Sbordone (2002), 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 procedure to be somewhat more natural in the context we are discussing, our point—that the fit for inflation under either marginal cost proxy is good when this method is used—holds just as well if we use her approach to estimate λ_1 and λ_2 .

12. 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.

A Price-Level Equation Using Expected ULC Growth

B Price-Level Equation Using Expected Nominal GDP Gap Growth

FIG. 4. Actual and Predicted Inflation, Alternative Method (VAR models include GDP gap, labor's share, and unit labor cost or nominal GDP gap growth.) Actual inflation given by solid line; predicted inflation given by dashed line

conclusion of Section 1 that the model does a poor job in capturing the observed dynamics of inflation. Rather, they illustrate that Sbordone's framework is not well designed for revealing the underlying weakness of the new-Keynesian Phillips curve.

3. CONCLUSIONS

In this paper, we have 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) are used. We find that the robust conclusion that emerges is that *neither* variable allows the new-Keynesian model to fit well.

Our relatively negative assessment is closely related to Fuhrer and Moore's (1995) critique of standard sticky-price models, which highlighted 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 argue that lags of inflation are proxying 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: The labor share of income does not appear to drive inflation. Hence, there is little reason to view the labor income share as providing a good measure of the output gap (particularly given that this would imply that every postwar U.S. recession has actually been a *boom* relative to the prevailing level of potential output). Similarly, a compelling case cannot be made for replacing conventional output gap measures with the labor share in a monetary policy rule.

Finally, we note that our conclusions should not be interpreted as implying that forward-looking inflation models based on real marginal cost *cannot* work, since it may be that both of the driving variables considered here are actually 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.

LITERATURE CITED

- Fuhrer, Jeffrey C., and George R. Moore (1995). "Inflation Persistence." *Quarterly Journal of Economics* 110, 127–159.
- Galí, Jordi, and Mark Gertler (1999). "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44, 195–222.
- Goodfriend, Marvin, and Robert King (2001). "The Case for Price Stability." NBER Working Paper No. 8423.
- Rotemberg, Julio, and Michael Woodford (1999). "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics*, Vol. 1B, edited by John Taylor, and Michael Woodford, pp. 1051–1135. Amsterdam: North-Holland.
- 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.
- Sargent, Thomas J. (1987). *Macroeconomic Theory*, 2nd edition. Boston, MA: Academic Press.
- Sbordone, Argia (2001). "An Optimizing Model of U.S. Wage and Price Dynamics." March draft, Rutgers University.
- Sbordone, Argia (2002). "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49, 265–292.
- Woodford, Michael (2001). "The Taylor Rule and Optimal Monetary Policy." *American Economic Review* 91, 232–237.

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