NBER WORKING PAPER SERIES

INFLATION DYNAMICS: A STRUCTURAL ECONOMETRIC ANALYSIS

Jordi Galí
Mark Gertler

Working Paper 7551
http://www.nber.org/papers/w7551

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2000

The authors thank participants at the JME-SNB Gerzensee Conference on "The Return of the Phillips Curve, NBER Summer Institute and ME Meetings, and seminars at Lausanne, UPF, Delta, Chicago, Michigan, Princeton, Yale, Columbia, San Francisco Fed, BIS, IIES, and the ECB, for useful comments. Special thanks also to John Roberts and Mark Watson. Tomasso Monacelli and Fabio Natalucci provided excellent research assistance. Financial support from the C.V. Starr Center for Applied Economics, the National Science Foundation, and CREI is gratefully acknowledged. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Inflation Dynamics: A Structural Econometric Analysis
Jordi Galí and Mark Gertler
NBER Working Paper No. 7551
February 2000
JEL No. E31

ABSTRACT

We develop and estimate a structural model of inflation that allows for a fraction of firms that use a backward looking rule to set prices. The model nests the purely forward looking New Keynesian Phillips curve as a particular case. We use measures of arginal cost as the relevant determinant of inflation, as the theory suggests, instead of an ad-hoc output gap. Real marginal costs are a significant and quantitatively important determinant of inflation. Backward looking price setting, while statistically significant, is not quantitatively important. Thus, we conclude that the New Keynesian Phillips curve provides a good first approximation to the dynamics of inflation.

Jordi Galí
Department d'Economia
Universitat Pompeu Fabra
Ramon Trias Fargas 25
08005 Barcelona
Spain
and NBER, and New York University
jordi.gali@econ.upf.es

Mark Gertler
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
and NBER
gertlerm@fasecon.econ.nyu.edu

1 Introduction

Among the central issues in macroeconomics is the nature of short run inflation dynamics. This matter is also one of the most fiercely debated, with few definitive answers available after decades of investigation. At stake, among other things, is the nature of business cycles and what should be the appropriate conduct of monetary policy.[1]

In response to this challenge, important advances have emerged recently in the theoretical modeling of inflation dynamics.[2] This new literature builds on early work by Fischer (1977), Taylor (1980), Calvo (1983) and others that emphasized staggered nominal wage and price setting by forward looking individuals and firms. It extends this work by casting the price setting decision within an explicit individual optimization problem. Aggregating over individual behavior then leads, typically, to a relation that links inflation in the short run to some measure of overall real activity, in the spirit of the traditional Phillips curve. The explicit use of microfoundations, of course, places additional structure on the relation and also leads to some important differences in details.

Despite the advances in theoretical modeling, accompanying econometric analysis of the "new Phillips curve" has been rather limited, though with a few notable exceptions.[3] The work to date has generated some useful findings, but these findings have also raised some troubling questions about the existing theory. As we discuss below, it appears difficult for these models to capture the persistence in inflation without appealing either to some form of stickiness in inflation that is hard to motivate explicitly or to adaptive expectations, which also poses difficulty from a modeling standpoint. In addition, with quarterly data, it is often difficult to detect a statistically significant effect of real activity on inflation using the structural formulation implied by theory, when the measure of real activity is an output gap (i.e., real output relative to some measure of potential output). Failure to find a significant short run link between real activity and inflation is obviously unsettling for the basic story.

In this context, we develop and estimate a structural model of the Phillips curve. Our approach has three distinctive features. First, in the empirical implementation, we use a measure of real marginal cost in place of an ad hoc output gap, as the theory suggests. As will become apparent, a desirable feature of a marginal cost measure is that it directly accounts for the impact of productivity gains on inflation, a factor

[1] For recent work that explores how the appropriate course of monetary policy depends on the nature of short run inflation dynamics, see Svensson (1997a, 1997b), Clarida, Gali and Gertler (1997b), Rotemberg and Woodford (1997b), McCallum and Nelson (1998), King and Wolman (1998), and Erceg, Henderson and Levin (1998).

[2] See Goodfriend and King (1997) for a comprehensive survey.

[3] Examples of work that attempts to estimate the new Phillips curve include, Chada, Masson and Meredith (1992), Fuhrer and Moore (1995), Fuhrer (1997), and Roberts (1997, 1998). For discussions of the traditional empirical literature on the Phillips curve, see King and Watson (1994), Gordon (1996) and Lown and Rich (1997).

that simple output gap measures often miss. In this respect, our approach is complementary to Sbordone (1998), though she uses a different methodology to empirical assess the model than we do.[4] Second, we extend the baseline theory underlying the new Phillips curve to allow for a subset of firms that set prices according to a backward looking rule of thumb. Doing so allows us to directly estimate the degree of departure from a pure forward looking model needed to account for the observed inflation persistence. Third, we identify and estimate all the structural parameters of the model using conventional econometric methods. The coefficients in our structural inflation equation are "mongrel" functions of two key model primitives: the average duration that an individual price is fixed (i.e., the degree of price "stickiness") and the fraction of firms that use rule of thumb behavior (i.e., the degree of "backwardness").

As we show, several results stand out and appear to be quite robust: (a) Real marginal costs are indeed a statistically significant and quantitatively important determinant of inflation, as the theory predicts; (b) Forward looking behavior is very important: our model estimates suggest that roughly sixty to eighty percent of firms exhibit forward looking price setting behavior; (c) Backward looking behavior is statistically significant though, in our preferred specifications, is of limited quantitative importance. Thus, while the benchmark pure forward looking model is rejected on statistical grounds, it appears still to be a reasonable first approximation of reality; (d) The average duration a price is fixed is considerable, but the estimates are in line with survey evidence.

Taken as whole, our results are supportive of the new, theory-based Phillips curves. But they also raise a puzzle. Traditional explanations of inertia in inflation (and hence the costs of disinflations) rely on some form of "backwardness" in price setting. To the extent this backwardness is not quantitatively important, as we seem to find, the story needs to be re-examined. In our view, the "black box" to investigate is the link between aggregate activity and real marginal costs. To the extent they are reasonably characterized by unit labor costs, real marginal costs tend to lag output over the cycle rather than move contemporaneously, in contrast to the prediction of the standard sticky price macroeconomic framework.[5] In this respect, our analysis suggests that a potential source of inflation inertia may be sluggish adjustment of real marginal costs to movements in output. We elaborate on this possibility in the conclusion.

[4] Sbordone (1998) explores how well the model fits the data conditional on different choices of a parameter that governs the degree of price rigidity. Our approach is to directly estimate the structural parameters using an instrumental variables procedure that is based on the orthogonality conditions that evolve from the underlying theory. In addition, we develop a general model that nests the pure forward looking model as a special case. Doing so allows us to test directly the departure from the pure forward looking model that is required to explain the data. Despite the sharp differences in methodology, the main conclusions we draw are very similar to hers, as we discuss later.

[5] Christiano, Eichenbaum and Evans (1997) also stress that the standard sticky price framework does not seem to explain the cyclical behavior of marginal cost.

The paper proceeds as follows: Section 2 reviews the basic theory underlying the new Phillips curve and discusses the existing empirical literature. We make clear why specifications based on the output gap are likely to be unsuccessful. Section 3 then presents estimates of the new Phillips curve using a measure of real marginal cost, and shows that with this specification the theory does a reasonably good job of describing the data. To explore further the issue of how well the theory captures the inertia in inflation, section 4 extends the model to allow for a subset of firms that use rule of thumb behavior. It then presents estimates of the resulting augmented Phillips curve and a variety of robustness exercises. In addition, we construct a measure of "fundamental inflation" based on the solution to the estimated model that relates inflation to a discounted stream of expected future marginal costs, as well as lagged inflation. We in turn show that this measure does a good job of describing the actual path of inflation, including the recent period. Section 5 concludes.

2 The New Phillips Curve: Background Theory and Evidence

In this section we review the recent theory that generates an estimable Phillips curve relation. We then discuss some of the pitfalls involved in estimating this relation and how the literature has dealt with these issues to date. Finally, we describe our approach.

2.1 A Baseline Model

The typical starting point for the derivation of the new Phillips curve is an environment of monopolistically competitive firms that face some type of constraints on price adjustment. In the most common incarnations, the constraint is that the price adjustment rule is time dependent. For example, every period the fraction $\frac{1}{X}$ of firms set their prices for X periods. The scenario is in the spirit of Taylor's (1980) staggered contracts model. A key difference is that the pricing decision evolves explicitly from a monopolistic competitor's profit maximization problem, subject to the constraint of time dependent price adjustment.

In general, however, aggregation is cumbersome with deterministic time dependent pricing rules at the micro level: It is necessary to keep track of the price histories of firms. For this reason, it is common to employ an assumption due to Calvo (1983) that greatly simplifies the aggregation problem.[6] The idea is to assume that in any

[6]Examples of frameworks that employ the Calvo assumption include Yun (1996), King and Wolman (1995), Woodford (1996), Rotemberg and Woodford (1997a, 1997b), Clarida, Gali and Gertler (1997a), McCallum and Nelson (1998), and Bernanke, Gertler and Gilchrist (1998). For a general equilibrium sticky price model based on the Taylor (1980) formulation, see Chari, Kehoe and McGratten (1996) and Kiley (1997).

given period each firm has a fixed probability $1-\theta$ that it may adjust its price during that period and, hence, a probability θ that it must keep its price unchanged. This probability is independent of the time elapsed since the last price revision. Hence, the average time over which a price is fixed is given by $(1-\theta)\sum_{k=0}^{\infty} k\theta^{k-1} = \frac{1}{1-\theta}$. Thus, for example, with $\theta = .75$ in a quarterly model, prices are fixed on average for a year. Because the adjustment probabilities are independent of the firm's price history, the aggregation problem is greatly simplified.

We can derive the new Phillips curve by proceeding as follows[7]: Assume that firms are identical ex ante, except for the differentiated product they produce and for their pricing history. Assume also that each faces a conventional constant price elasticity of demand curve for its product. Then it is possible to show that the aggregate price level p_{t-1} evolves as a convex combination of the lagged price level p_t and the optimal reset price p_t^* (i.e. the price selected by firms that are able to change price at t), as follows:

$$p_t = \theta\ p_{t-1} + (1-\theta)\ p_t^* \tag{1}$$

where each variable is expressed as a percent deviation from a zero inflation steady state. Intuitively, the fraction $1-\theta$ of firms that set their price at t all choose the same price p_t^* since they are identical (except for the differentiated product they produce). By the law of large numbers, further, the index of prices for firms that do not adjust during the period is simply equal to the lagged price level.

Let mc_t^n be the firm's nominal marginal cost at t (as a percent deviation from the steady state) and let β denote a subjective discount factor. Then, for a firm that chooses price at t to maximize expected discounted profits subject to the time dependent pricing rules given by the Calvo formulation, the optimal reset price may be expressed as:

$$p_t^* = (1-\beta\theta)\ \sum_{k=0}^{\infty}(\beta\theta)^k\ E_t\{mc_{t+k}^n\} \tag{2}$$

In setting its price at t, the firm takes account of the expected future path of nominal marginal cost, given the likelihood that its price may remain fixed for multiple periods. Note that in the limiting case of perfect price flexibility ($\theta = 0$), the firm simply adjusts its price proportionately to movements in the current marginal cost. The future becomes relevant only when there is price rigidity (i.e., $\theta > 0$).

[7]For an explicit derivation, see, e.g., Goodfriend and King (1997), King and Wolman (1996), or Woodford (1996).

2.1.1 Inflation and Marginal Cost

The Calvo formulation leads to a Phillips curve with properties reasonably similar to the standard staggered price formulation, but at the same time it is more tractable.[8] From the standpoint of estimation, further, the parsimonious representation is highly advantageous.

Let $\pi_t \equiv p_t - p_{t-1}$ denote the inflation rate at t, and mc_t the percent deviation of the firm's real marginal cost from its steady state value. By combining equations (1) and (2) it is possible to derive an inflation equation of the form:

$$\pi_t = \lambda \; mc_t + \beta \; E_t\{\pi_{t+1}\} \tag{3}$$

where the coefficient $\lambda \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}$ depends on the frequency of price adjustment θ and the subjective discount factor β.

Intuitively, because firms' (a) mark up price over marginal costs, (b) are forward looking, and (c) must lock into a price for (possibly) multiple periods, they base their pricing decisions on the expected future behavior of marginal costs. Iterating equation (3) forward yields

$$\pi_t = \lambda \; \sum_{k=0}^{\infty} \beta^k \; E_t\{mc_{t+k}\} \tag{4}$$

The benchmark theory thus implies that inflation should equal a discounted stream of expected future marginal costs.

2.1.2 Marginal Cost and the Output Gap

Traditional empirical work on the Phillips curve emphasizes some output gap measure as the relevant indicator of real economic activity, as opposed to marginal cost. Under certain assumptions, however, there is an approximate log-linear relationship between the two variables. Let y_t denote the log of output; y_t^* the log of the "natural" level of output (the level that would arise if prices were perfectly flexible); and $x_t \equiv y_t - y_t^*$ the "output gap". Then, under certain conditions one can write:[9]

$$mc_t = \kappa \; x_t \tag{5}$$

where κ is the output elasticity of marginal cost.

[8] Roberts (1997) demonstrates that the Calvo (1983) and Taylor (1980) models have very similar implications for inflation dynamics. Kiley (1997), however, shows that differences can emerge if the elasticity of inflation with respect to real marginal cost is large. Our estimates below point to a relatively small elasticity. Nonetheless, extending our analysis to allow for alternative forms of price staggering would be a useful undertaking.

[9] In the standard sticky price framework without variable capital (e.g, Rotemberg and Woodford (1997)), there is an approximate proportionate relation between marginal cost and output. With variable capital the relation is no longer proportionate. Simulations suggest, however, that the relation remains very close to proportionate.

Combining the relation between marginal cost and the output gap with equation (3) yields a Phillips curve-like relationship:

$$\pi_t = \lambda\kappa \; x_t + \beta \; E_t\{\pi_{t+1}\} \tag{6}$$

As with the traditional Phillips curve, inflation depends positively on the output gap and a "cost push" term that reflects the influence of expected inflation. A key difference is that it is $E_t\{\pi_{t+1}\}$ as opposed to $E_{t-1}\{\pi_t\}$ (generally assumed to equal π_{t-1}) that matters. As a consequence, inflation depends exclusively on the discounted sequence of future output gaps. This can be seen by iterating equation (6) forward, which yields:

$$\pi_t = \lambda\kappa \sum_{k=0}^{\infty} \beta^k \; E_t\{x_{t+k}\} \tag{7}$$

2.2 Empirical Issues

Reconciling the new Phillips curve with the data, has not proved to be a simple task. In particular, equation (6) implies that current *change* in inflation should depend negatively on the lagged output gap. To see, lag equation (6) one period; and then assume $\beta \simeq 1$ to obtain

$$\pi_t = -\lambda\kappa \; x_{t-1} + \pi_{t-1} + \varepsilon_t \tag{8}$$

where $\varepsilon_t \equiv \pi_t - E_{t-1}\pi_t$. But estimating equation (8) with U.S. data, and using (quadratically) detrended log GDP as a measure of the output gap yields

$$\pi_t = 0.081 \; x_{t-1} + \pi_{t-1} + \varepsilon_t \tag{9}$$

i.e., the inflation rate depends positively on the lagged output gap rather than negatively: The estimated equation, unfortunately, resembles the old curve rather than the new!

The essential problem, as emphasized by Fuhrer and Moore (1995), is that the benchmark new Phillips curve implies that inflation should lead the output gap over the cycle, in the sense that a rise (decline) in current inflation should signal a subsequent rise (decline) in the output gap. Yet, exactly the opposite pattern can be found in the data. The top panel in Figure 1 presents the cross-correlation of the current output gap (measured by detrended log GDP) with leads and lags of inflation.[10] As the panel indicates clearly, the current output gap co-moves positively with future inflation and negatively with lagged inflation. This lead of the output gap over inflation explains why the lagged output gap enters with a positive coefficient in equation (9), consistent with the old Phillips curve theory but in direct contradiction of the new.

[10] The cross-correlations reported in figure 1 were computed on HP-detrended series over the period 1960:1-1997:4. We provide a more extensive discussion of Figure 1 in the conclusion.

Another discomforting feature of the new Phillips curve as given by equation (6) is the stark prediction of no short run trade-off between output and inflation. Put differently, equation (7) implies that a disinflation of any size could be achieved costlessly and immediately by a central bank that could commit to setting the path of future output gaps equal to zero. The historical experience suggests, in contrast, that disinflations involve a substantial output loss [e.g., Ball (1994)]. It may be possible to appeal to imperfect credibility to reconcile the theory with the data. If, for example, the central bank cannot commit to stabilizing future output, then reduction of inflation may involve current output losses [e.g., Ball (1995)]. While this theory clearly warrants further investigation, there is currently, however, little direct evidence to support it. Further, countries with highly credible central banks (e.g., Germany) have experienced very costly disinflations [e.g., Clarida and Gertler (1997)].

The empirical limitations of the new Phillips curve have led a number of researchers to consider a hybrid version of the new and old:

$$\pi_t = \delta\ x_t + (1-\phi)\ E_t\{\pi_{t+1}\} + \phi\ \pi_{t-1} \tag{10}$$

with $0 < \phi < 1$. The idea is to let inflation depend on a convex combination of expected future inflation and lagged inflation. The addition of the lag term is designed to capture the inflation persistence that is unexplained in the baseline model.[11] A further implication of the lag term is that disinflations now involve costly output reduction.

The motivation for the hybrid approach is largely empirical. Fuhrer and Moore (1995) appeal to Buiter and Jewitt's (1985) relative wage hypothesis. While the story may be plausible, it does not evolve from an explicit optimization problem, in contrast to the benchmark formulation. Roberts (1997, 1998) instead appeals to adaptive expectations on the part of a subset of price setters. Under his formulation, some form of adaptive rule replaces lagged inflation.

Oddly enough, however, the hybrid Phillips curve has met with rather limited success. In particular, the relation does not seem to provide a good characterization of inflation dynamics at the quarterly frequency. Chadha, Masson, and Meredith (1992), for example, obtain reasonable parameter estimates of equation (10), but only with annual data. Roberts (1997, 1998) similarly works mainly with annual and semi-annual data. With quarterly data, he has difficulty obtaining significant estimates of the effect of the output gap on inflation. Fuhrer (1997) is able to obtain

[11] A special case of equation (10) with $\phi = 0.5$ is the widely used "sticky inflation" model of Buiter and Jewitt (1985) and Fuhrer and Moore (1995):

$$(\pi_t - \pi_{t-1}) = \frac{\delta}{0.5}\ x_t + \ (E_t\pi_{t+1} - \pi_t) \tag{11}$$

Under this formulation, the *change* in the inflation rate is related the expected path of the future output gaps.

a significant output gap coefficient with quarterly data, but only when the model is heavily restricted. In this instance the estimated model is consistent with the old Phillips curve: expected future inflation does not enter significantly in the inflation equation; lagged inflation enters with a coefficient near unity, as in the traditional framework.

2.3 Shortcomings

There are, however, several problems with this approach that could possibly account for the empirical shortcomings. First, conventional measures of the output gap x_t are likely to be ridden with error, primarily due to the unobservability of the natural rate of output y_t^*. [12] A typical approach (followed above) to measuring y_t^* is to use a fitted deterministic trend. Alternatives are to use the Congressional Budget Office (CBO) estimate or instead use a measure of capacity utilization as the gap variable. It is widely agreed that all these approaches involve considerable measurement error. To the extent there is significant high frequency variation in y_t^* (e.g., due to supply shocks) mismeasurement could distort the estimation of an inflation equation like (6) or (10).[13] Though, whether correcting for measurement error alone could reverse the lead-lag pattern between the output gap and inflation that is apparent from Figure 1 is problematic in our view.

A more fundamental issue, we believe, is that even if the output gap were observable the conditions under which it corresponds to marginal cost may not be satisfied. Our analysis of the data suggests that movements in our measure of real marginal cost (described below) tend to lag movements in output, in direct contrast to the identifying assumptions that imply a co-incident movement. This discrepancy, we will argue, is one important reason why structural estimation of Phillips curves based on the output gap have met with limited success, at best.

2.4 Our Approach

In light of the difficulties with using the output gap, we instead use in the empirical analysis below measures of real marginal cost, in a way consistent with the theory. In other words, we estimate (3) instead of (6). Since real marginal cost is not directly observable, we use restrictions from theory to derive a measure based on observables. Conditional on our measure of real marginal cost, we can then obtain estimates of the structural parameters in equation (3), including the frequency of price adjustment θ,

[12] This issue is currently of great practical importance in the U.S.: in recent years the measured output gap is well above trend, but inflation is well below trend. It thus appears that mismeasurement of the true output gap is confounding the ability of traditional Phillips curves to explain the data. See Lown and Rich (1997).

[13] For example, in the presence of nominal rigidities, supply shocks are likely to move detrended output and the true output gap in opposite directions [Gali (1999)]. In addition, unobserved supply shocks could potentially account for some of the explanatory power of lagged inflation.

the parameter that governs the degree of price stickiness (i.e., the average period a price remains fixed.)

We also derive an econometric specification that permits us to assess the degree to which the new Phillips curve can account for the inertia in inflation. In particular, we derive a "hybrid Phillips" curve that nests the new Phillips curve as a special case, but allows for a subset of firms use a backward looking rule of thumb to set prices. The advantage of proceeding this way is that the coefficients of our hybrid Phillips curve will be functions of two key parameters: the frequency of price adjustment and the fraction of backward looking price setters. Note that the latter parameter provides a direct measure of the departure from a pure forward looking model needed to account for the persistence in inflation.

In the next section we present estimates of the new Phillips curve, and in the subsequent one we present estimates of our hybrid Phillips curve.

3 New Estimates of the New Phillips Curve

We first describe our econometric specification of the new Phillips curve, along with our general estimation procedure. We then present both reduced form and structural estimates of the model.

3.1 Econometric Specification

We begin by describing how we obtain a measure of real marginal cost. For simplicity, we restrict ourselves to the simplest measure of marginal cost available, one based on the assumption of a Cobb-Douglas technology. Let A_t denote technology, K_t capital, and N_t labor. Then output Y_t is given by

$$Y_t = A_t \; {K_t}^{\alpha_k} \; {N_t}^{\alpha_n} \tag{12}$$

Real marginal cost is then given by the ratio of the wage rate to the marginal product of labor, i.e., $MC_t = \frac{W_t}{P_t} \; \frac{1}{\partial Y_t / \partial N_t}$. Hence, given equation (12) we have:

$$MC_t = \frac{S_t}{\alpha_n} \tag{13}$$

where $S_t \equiv \frac{W_t N_t}{P_t Y_t}$ is the labor income share (equivalently, real unit labor costs).[14]. Letting lower case letters denote percent deviations from the steady state we have:

$$mc_t = s_t \tag{14}$$

[14]Interestingly, Lown and Rich (1997) show that augmenting the growth of a traditional Phillips curve with the growth rate of nominal unit labor costs greatly improves the fit. We also stress the role of unit labor costs, except that in our approach, (the log level) of real unit labor costs enters as the relevant gap variable, as the theory suggests.

Combining equations (14) and (3) yields the inflation equation:

$$\pi_t = \lambda \; s_t + \beta \; E_t\{\pi_{t+1}\} \tag{15}$$

where the coefficient λ is given by

$$\lambda = \frac{(1-\theta)(1-\beta\theta)}{\theta} \tag{16}$$

Since under rational expectations the error in the forecast of π_{t+1} is uncorrelated with information dated t and earlier, it follows from equation (14) that

$$E_t\{(\pi_t - \lambda \; s_t - \beta \; \pi_{t+1}) \; \mathbf{z}_t\} = 0 \tag{17}$$

where $\mathbf{z}_t$ is a vector of variables dated t and earlier (and, thus, orthogonal to the inflation surprise in period $t+1$). The orthogonality condition given by equation (17) then forms the basis for estimating the model via Generalized Method of Moments (GMM).

The data we use is quarterly for U.S. over the period 1960:1 to 1997:4. We use the (log) labor income share in the non-farm business sector for s_t. Our inflation measure is the percent change in the GDP deflator. We use the overall deflator rather than the non-farm deflator for most of our analysis because we are interested evaluating how well our model accounts for the movement in a standard broad measure of inflation. We show, however, that our results are robust to using the non-farm deflator. Finally, our instrument set includes four lags of inflation, the labor income share, output gap, the long-short interest rate spread, wage inflation, and commodity price inflation.

3.2 Reduced Form Evidence

We first report our estimate of equation (17). We refer to this evidence as "reduced form" since it contains an estimate of the overall slope coefficient on marginal cost, λ, but not of the structural parameter θ (the measure of price rigidity) that underlies λ (see equation 16). The resulting estimated equation is given byλ

$$\pi_t = \underset{(0.012)}{0.023} \; s_t + \underset{(0.045)}{0.942} \; E_t\{\pi_{t+1}\}$$

Overall, the estimated new Phillips curve is quite sensible. The slope coefficient λ on real marginal cost is positive and significant, as is consistent with the a priori theory. The estimate of the coefficient on expected inflation, the subjective discount factor β, is also reasonable, particularly after accounting for the sampling error implied by the estimated standard deviation.[15] Thus, at first pass, it appears that the new Phillips curve provides a reasonable description of inflation.

[15] In particular, the estimate of β is within two standard deviations of typical values for this parameter that are used in the literature (e.g., 0.99).

To highlight the virtues of using real marginal cost as the relevant real sector driving variable in the new Phillips curve, we reestimate equation (6), using detrended log GDP as a proxy for the output gap x_t:

$$\pi_t = \underset{(0.005)}{-0.016}\ x_t + \underset{(0.030)}{0.988}\ E_t\{\pi_{t+1}\}$$

The model clearly doesn't work in this case: the coefficient associated with the output gap is negative and significant, which is at odds with the prediction of the theory. This finding, of course, is completely consistent with our earlier result that, when the model is reversed and estimated in the form of the old Phillips curve, the coefficient on the lagged output gap is positive (see equation (9)). Thus, it is the use of real marginal cost over the output gap, and not the estimation strategy, that accounts for the econometric success of the new Phillips curve.

3.3 Structural Estimates

We now redo the exercise in a way that allows us to obtain direct estimates of the structural parameter θ. In particular, we substitute the relation for λ, equation (16), into equation (17) to obtain an econometric specification that is nonlinear in the structural parameters θ and β.

One econometric issue we must confront is that, in small samples, nonlinear estimation using GMM is sometimes sensitive to the way the orthogonality conditions are normalized. For this reason, we use two alternative specifications of the orthogonality conditions as the basis for our GMM estimation procedure.[16] The first specification takes the form

$$E_t\{(\theta\ \pi_t - (1-\theta)(1-\beta\theta)\ s_t - \theta\beta\ \pi_{t+1})\ \mathbf{z}_t\} = 0 \tag{18}$$

while the second is given by:

$$E_t\{(\pi_t - \theta^{-1}(1-\theta)(1-\beta\theta)\ s_t - \beta\ \pi_{t+1})\ \mathbf{z}_t\} = 0 \tag{19}$$

We estimate the structural parameters θ and β using a nonlinear instrumental variables estimator, with the set of instruments the same as is in the previous case. For robustness, we consider two alternatives to the benchmark case. In the first alternative we restrict the estimate of the discount factor β to unity. In the second, we use the non-farm GDP deflator as opposed to the overall deflator. Finally, we estimate each specification using the two different normalizations, as given by equations (18) and (19).

The results are reported in Table 1. The first two columns give the estimates of θ and β. The third then gives the implied estimate of λ, the reduced form slope

[16] Among the possible normalizations we have chosen the two which we view as most natural. The first one appears to minimize the non-linearities, while the second normalizes the inflation coefficient to unity. See, e.g., Fuhrer, Moore, and Schuh (1995) for further discussion of the normalization issue.

coefficient on real marginal cost. In general, the structural estimates tell the same overall story as the reduce form estimates. The implied estimate of λ is always positive and is highly significant in every case but one (restricted β, normalization (2)). The estimate of β in the unrestricted case is somewhat low, but not unreasonably so, given the sampling uncertainty.

The estimate of the structural parameter θ is somewhat large and also somewhat sensitive to the normalization in the GMM estimation. Using method (1), we estimate θ to be around 0.83 with a small standard error, which implies that prices are fixed for between roughly five and six quarters on average. That period length is close to the average price duration found in survey evidence, though perhaps on the high side.[17] Method (2) yields a slightly higher estimate of θ, around 0.88. Since λ is decreasing in θ (greater price rigidity implies that inflation is less sensitive to movements in real marginal cost), the higher estimates of θ implies a lower estimate of λ for method (2).

For several reasons, however, our estimates of the degree of price rigidity are likely to be upward biased. First, it is likely that the labor share does not provide an exact measure of real marginal cost. In this instance, the estimate of the slope coefficient λ is likely to be biased towards zero. This translates into upward bias of θ, given the inverse link between the two parameters. Second, the underlying theory that is used to identify θ from estimates of λ, assumes a constant markup of price over marginal cost in the absence of prices rigidities. If the markup in the frictionless benchmark model were countercyclical, as much recent theory has argued, the implied estimate of θ would be lower.[18] With a countercyclical markup, desired price setting is less sensitive to movements in marginal cost, which could help account for low overall sensitivity of inflation to the labor share.

The model also works well in the sense that we do not reject the overidentifying restrictions. Though we do not report the results here, the p-values for the null hypothesis that the error term is uncorrelated with the instruments are all in the range of 0.9 or higher. This kind of test has low power, however, since it is not applied against any specific alternative hypothesis. In the next section we develop a more refined test to measure how well the model accounts for inflation dynamics.

4 A New Hybrid Phillips Curve

We now explicitly address the issue of how well the new Phillips curve captures the apparent inertia in inflation. To do so, we extend the basic Calvo model to allow for a subset of firms that use a backward looking rule of thumb to set prices. Our formulation allows us to estimate the fraction of firms that lies in this subset. By doing so we obtain a measure of the residual inertia in inflation that the baseline new

[17] See Taylor (1998) for an overview of that evidence.

[18] See, e.g., Kimball (1995) for an illustration of a countercycical desired markup in the context of a sticky price model.

Phillips curve leaves unexplained.[19]

4.1 Theoretical Formulation

We continue to assume, as in Calvo's model, that each firm is able to adjust its price in any given period with a fixed probability $1-\theta$ that is independent of the time the price has been fixed. We depart from Calvo by having two types of firms co-exist. A fraction $1-\omega$ of the firms, which we refer to as "forward looking," behave like the firms in Calvo's model: they set prices optimally, given the constraints on the timing of adjustments and using all the available information in order to forecast future marginal costs. The remaining firms, of measure ω, which we refer to as "backward looking," instead use a simple rule of thumb that is based on the recent history of aggregate price behavior.

The aggregate price level now evolves according to:

$$p_t = \theta \; p_{t-1} + (1-\theta) \; \overline{p}_t^* \tag{20}$$

where $\overline{p}_t^*$ is an index for the prices newly set in period t. Let p_t^f denote the price set by a forward looking firm at t and p_t^b the price set by a backward looking firm. Then the index for newly set prices may be expressed as

$$\overline{p}_t^* = (1-\omega) \; p_t^f + \omega \; p_t^b \tag{21}$$

Forward looking firms behave exactly as in the baseline Calvo model described above. Accordingly, p_t^f may be expressed as

$$p_t^f = (1-\beta\theta) \; \sum_{k=0}^{\infty} (\beta\theta)^k \; E_t\{mc_{t+k}^n\} \tag{22}$$

We assume that backward looking firms obey a rule of thumb that has the following two features: (a) no persistent deviations between the rule and optimal behavior; i.e., in a steady state equilibrium the rule is consistent with optimal behavior; (b) the price in period t given by the rule depends only on information dated $t-1$ or earlier. We also assume that the firm is unable to tell whether any individual competitor is backward looking or forward looking.

These considerations lead us to a rule that is based on the recent pricing behavior of the firm's competitors, as follows:

$$p_t^b = p_{t-1}^* + \pi_{t-1} \tag{23}$$

In other words, a backward looking firm at t sets its price equal to the average price set in the most recent round of price adjustments, p_{t-1}^*, with a correction for inflation.

[19] Thus, by adding rule-of-thumb price setters, we measure the departure from the baseline forward looking model similar to the way that Campbell and Mankiw (1989) used rule-of-thumb consumers to test the life-cycle/ permanent income hypothesis.

Importantly, the correction is based on the lagged inflation rate, i.e., lagged inflation is used in a simple way to forecast current inflation.

Though admittedly ad hoc, the rule has several appealing features. First, as long as inflation is stationary, the rule converges to optimal behavior over time.[20] Second, the rule implicitly incorporates information about the future in a useful way, since the price index p^*_{t-1} is partly determined by forward looking price setters. Thus, to the extent the percent difference between the forward and backward price is not large, the loss to a firm from rule of thumb behavior will be second order, for the usual arguments due to the envelope theorem. This is more likely to be the case if backward looking price setters are a relatively small fraction of the population.[21]

We obtain our hybrid Phillips curve by combining equations (20), (21), (22), and (23):

$$\pi_t = \lambda \; mc_t + \gamma_f \; E_t\{\pi_{t+1}\} + \gamma_b \; \pi_{t-1} \tag{24}$$

where

$$\begin{aligned} \lambda &\equiv (1-\omega)(1-\theta)(1-\beta\theta)\phi^{-1} \\ \gamma_f &\equiv \beta\theta\phi^{-1} \\ \gamma_b &\equiv \omega\phi^{-1} \end{aligned} \tag{25}$$

with $\phi \equiv \theta + \omega \; [1 - \theta(1-\beta)]$.

This specification differs from the hybrid model used in recent empirical research (discussed in the previous section) in two fundamental ways. First, real marginal cost as opposed to the output gap is the forcing variable. Second, all the coefficients are explicit functions of three model parameters: θ, which measures the degree of price stickiness; ω, the degree of "backwardness" in price setting, and the discount factor β.

Two special cases provide useful benchmarks: First, when $\omega = 0$, all firms are forward looking and the model converges to the benchmark new Phillips curve introduced in the previous section. Second, when $\beta = 1$, then $\gamma_f \; + \gamma_b = 1$, which implies that the model takes the form of hybrid equation discussed earlier (except that marginal cost and not the output gap appears now as the driving force).

[20] More precisely, as long as inflation is stationary, there are no persistent deviations between the rule and optimal behavior; this can be seen by noting that $p^b_t - p_t = \theta(1-\theta)^{-1} \; \pi_t$

[21] When backward looking price-setters are a relatively small fraction of the population, the index of newly set prices p^*_t is dominated by forward looking price setters. Given that p^b_t closely tracks p^*_{t-1}, the backward looking price will be close on average to the forward looking price. We have conducted simulations of a complete model that bear out this logic.

4.2 Estimation and Results

In this section we present estimates of the previous structural model and also evaluate its overall performance vis-a-vis the data. As in the previous section we use the labor share to measure real marginal cost. The empirical version of our hybrid Phillips curve is accordingly given by:

$$\pi_t = \lambda\ s_t + \gamma_f\ E_t\{\pi_{t+1}\} + \gamma_b\ \pi_{t-1} \tag{26}$$

together with equation (25), which describes the relation between the reduced form and structural parameters.

We estimate the structural parameters β, θ, and ω using a non-linear instrumental variables (GMM) estimator. The instrument set is the same as we used in the previous exercises. To address the small sample normalization problem with GMM that we discussed earlier, we again use two alternative specifications of the orthogonality conditions, one which does not normalize the coefficient on inflation to be unity (method 1) and one which does (method 2):

$$E_t\{(\phi\ \pi_t - (1-\omega)(1-\theta)(1-\beta\theta)\ s_t - \theta\beta\ \pi_{t+1})\ \mathbf{z}_t\} = 0 \tag{27}$$

$$E_t\{(\pi_t - (1-\omega)(1-\theta)(1-\beta\theta)\phi^{-1}\ s_t - \theta\beta\phi^{-1}\ \pi_{t+1})\ \mathbf{z}_t\} = 0 \tag{28}$$

Table 2 presents the estimates of equation (26) As in the previous section, we consider three cases: the baseline model; the model with β restricted to unity; and the non-farm deflator substituted for the overall GDP deflator. The first three columns give the estimated structural parameters. The next three give the implied values of the reduced form coefficients (see equation 25).

Overall, the estimates are consistent with the underlying theory. The results, further, are reasonably consistent across specifications, though the precise estimate of the fraction of backward looking price-setters is somewhat sensitive to the use of method (1) versus method (2).

We begin with the baseline case. With method 1, the parameter θ is estimated to be about 0.81 with standard error 0.02, which implies that prices are fixed for roughly five quarters on average.[22] That period length may seem a bit long, but is not far off from survey evidence which suggests three to four quarters.[23] Method 2 yields an estimate that is not statistically different.

We now turn to the estimate of the fraction of backward looking price setters. With method 1, the parameter ω is estimated to be 0.26 with a standard error 0.06, implying that roughly quarter of price setters are backward looking. Thus,

[22] Interestingly, Sbordone (1998) finds that the value of the price adjustment parameter that maximizes the goodness of fit of the data by Watson's (1993) criterion also corresponds to an average of five quarters between adjustments. Thus, despite the difference in methodology, our results line up very closely with hers.

[23] For a discussion of the survey evidence, see Rotemberg and Woodford (1997a). Our sub-sample estimates (reported shortly) yield numbers directly in line with this evidence.

the pure forward looking model is rejected by the data. However, the quantitative importance of backward looking behavior for inflation dynamics is not large. The implied estimates for the reduced form coefficients on lagged versus expected future inflation are 0.25 (for γ^b) and 0.68 (for γ^f). Method 2 yields a higher estimate of ω, 0.49, implying that nearly half of price setters are backward looking. However, forward looking behavior remains predominant: The implied estimate of γ^f is 0.59 versus 0.38 for γ^b.

Thus, while the results suggest some imprecision in the estimate of the degree of backwardness, the central conclusions do not change across methods (1) and (2): In accounting for inflation dynamics, forward looking behavior is more important than backward looking behavior. In either case the estimate of the coefficient on expected future inflation in equation (26) lies well above the coefficient on lagged inflation. It is also true in either case that the estimates of the primitive parameters yield an estimate of the slope coefficient on the labor share λ that is positive and significant.[24] Thus, we are able to identify (in a robust manner) a significant impact of marginal costs on inflation.

It is also the case the model estimated using method (1) does a better job of tracking actual inflation the model based on method (2) estimates. (In section 4.4 we make precise the sense in how we evaluate the ability of the model to track the data.) To the extent that this provides a ground for preferring method (1), we can conclude that not only is forward looking behavior predominant but, given the small estimate of the degree of backwardness, the pure forward looking model may do a reasonably good job of describing the data.

The estimate of β is reasonably similar across the two methods, but somewhat on the low side at roughly 0.90. We thus next explore the implications of restricting β equal to unity, as implied in the standard hybrid case. Interestingly, there is little impact on the estimates of the other primitive parameters. Thus, restricting β to a plausible range does not affect the results in any significant way.

Finally, we consider the use of the non-farm deflator. Interestingly, there is no significant impact on the estimate of the degree of price rigidity. However, the estimate of the degree of backwardness drops. Indeed, with method (1), the estimate of ω is only 0.07. Though somewhat larger with method 2, it is still just 0.239. In either case, backward looking behavior is not quantitatively important. Overall, the pure forward looking model may provide a reasonably good description of inflation, as measured by the non-farm deflator.

[24] We note that the link between inflation and marginal cost is related to Benabou's (1992) finding using retail trade data that inflation is inversely related to the markup (which he measured as the inverse of the labor share). He interpreted the findings as evidence that the markup may depend on inflation, whereas in our model, causation runs from marginal cost to inflation. Sorting out possible simultaneity is an interesting topic for future research. We note, however, that our model has the additional implication that inflation should be related to a discounted stream of future marginal costs, and we shortly demonstrate that this appears to be the case.

4.3 Robustness Analysis

We now consider two robustness exercises.[25] The first allows extra lags of inflation to enter the right hand side of the equation for inflation. The second explores sub-sample stability.

We next add three additional lags of inflation to the baseline case (equation (26)). Here the idea is to explore whether our estimated importance of forward looking behavior may reflect not allowing for sufficient lagged dependence. Put differently, since we use four lags of inflation in our instrument set, we may be inadvertently biasing our "horse race" between expected future inflation and one quarter lagged inflation in favor of the former. The way to address this issue is to add the three additional lags of inflation to the right hand side, and then determine wether they have any predictive power for current inflation, π_t, beyond the signalling power they have for expected future inflation, $E_t\{\pi_{t+1}\}$.

Table 3 report the results. The parameter ψ denotes the sum of the coefficients on the three additional inflation lags. Since the estimates do not change much across method (1) and (2), we only report results for the former case. The overall effect of the additional lags is quite small, especially when the GDP deflator is used as the measure of inflation. The estimate of ψ is only 0.09 in the baseline case, and not significantly different from zero when β is restricted to unity. When the non-farm deflator is used the estimate of ψ rise to 0.21 with a standard error of 0.06. However, in this instance the first lag of inflation is not significantly different, so that the overall effect of lagged inflation is minimal. Thus, even though a total of four lags of inflation enters the right hand side, forward looking behavior still predominates. It thus appears that we account for inflation inertia with minimal reliance on arbitrary lags.

Finally, we consider sub-sample stability. Table 4 reports estimates over the intervals *60:1-79:4, 70:1-89:4,* and *80:1-97:4.* Again, since the conclusions we draw are unaffected by the normalization used, we restrict attention to method (1).

Overall, the broad picture remains unchanged. Marginal costs have a significant impact on short run inflation dynamics of roughly the same quantitative magnitudes as suggested by the full sample estimates. Forward looking behavior is always important. For the GDP deflator, in the first two sub-periods, the estimate of ω is close to the full sample estimate; i.e. roughly 0.25. Interestingly, though, in the last sub-period the estimate of ω drops in half to about 0.12. The pattern is the opposite for the non-farm deflator: estimates of ω near zero for the first two sub-samples (which correspond to the full-sample estimates), but rising slightly to 0.22 in the last sub-sample.

Another interesting result with the GDP deflator is that the estimate of θ for the

[25] In an earlier version of the paper we also allowed for increasing returns (in the form of overhead labor) in constructing the measure of marginal cost. Since this modification does not affect the results, we do not report the exercise here.

first two sub-samples drops from the full sample estimate of 0.8 to the range $0.75 - 0.77$. The important implication is that pre-1990, the estimated average duration a price is fixed is around four quarters, which is directly in line with the survey evidence. For the last sample, 1980:1 -1997:4, the estimate of θ rises to roughly 0.85, implying duration of six quarters. The longer duration might reflect the fact that inflation was lower over the last sub-sample. As a consequence, the average length between price adjustments may have increased (as, for example, a model of state-dependent pricing might imply.)

4.4 Actual vs. Fundamental Inflation

Our econometric Phillips curve, as given by equation (26), takes the form of a difference equation for inflation, with expected real marginal costs as the forcing variable. The solution for inflation implied by the model will depend on a discounted stream of expected future marginal costs, as well as lagged inflation. As a way to assess the model's goodness-of-fit, we consider how well the solution to the difference equation lines up against the actual data. We term our model-based measure of inflation "fundamental" inflation because it is analogous to Campbell and Shiller's (1987) construct of fundamental stock prices in terms of forecasts of discounted future dividends.

Our baseline estimates of γ_b and γ_f imply the existence of one stable and one unstable root associated with the stationary solution to the difference equation for inflation given by (26). Let $\delta_1 \leq 1$ denote the stable root and $\delta_2 \geq 1$ denote the unstable root. The model's solution is then given by:

$$\pi_t = \delta_1 \pi_{t-1} + \left(\frac{\lambda}{\delta_2 \gamma_f}\right) \sum_{k=0}^{\infty} \left(\frac{1}{\delta_2}\right)^k E_t \{s_{t+k}\} \tag{29}$$

The lagged term in equation (29) arises from the presence of backward looking price setters. In the benchmark case with pure forward looking behavior, the lagged term disappears (i.e., $\delta_1 = 0$).

Let $I_t = \{\pi_t, \pi_{t-1,...}, z_t, z_{t-1}, ...\}$ where z_t is a vector of variable other than inflation observed as of time t. Taking expectations conditional on I_t on both sides of (29):

$$\pi_t = \delta_1 \pi_{t-1} + \left(\frac{\lambda}{\delta_2 \gamma_f}\right) \sum_{k=0}^{\infty} \left(\frac{1}{\delta_2}\right)^k E[\ s_{t+k} \mid I_t] \equiv \pi_t^* \tag{30}$$

We construct our measure of fundamental inflation π_t^* using equation (30) based on $I_t = \{\pi_t,\ \pi_{t-1,...},\ s_t,\ s_{t-1}, ...\}$. Figure 2 plots fundamental inflation π_t^* versus actual inflation π_t. [26]

[26] In experimentation, we found that the model estimates based on method (1) do better in terms of tracking inflation than those based on method (2). Specifically, we found that the sum of squares of deviations between actual and fundamental inflation is lowest with method (1). We thus report only method (1) estimates in performing the exercise.

Overall fundamental inflation tracks the behavior of actual inflation very well.[27] It is particularly interesting to observe that it does a good job of explaining the recent behavior of inflation. During the past several years, of course, inflation has been below trend. Output growth has been above trend, on the other hand, making standard measures of the output gap highly positive. As a consequence, traditional Phillips curve equations have been overpredicting recent inflation.[28] However, because, real unit labor costs have been quite moderate recently despite rapid output growth, our model of fundamental inflation is close to target.

5 Conclusions

Our results suggest that, conditional on the path of real marginal costs, the baseline new Phillips curve with forward looking behavior may provide a reasonably good description of inflation dynamics. When tested explicitly against an alternative that allows for a fraction of price setters to be backward looking, the structural estimates suggest that this fraction, while statistically significant, is not quantitatively important. One qualification, however, is that there is some imprecision in our estimates of the importance of backward looking behavior. Yet, across all specifications forward looking behavior remains dominant. In the estimated hybrid Phillips curve, the weight on inflation lagged one quarter is generally small. Further, additional lags of inflation beyond one quarter do not appear to matter much at all. Taken as a whole, accordingly, the results suggest that it is worth searching for explanations of inflation inertia beyond the traditional ones that rely heavily on arbitrary lags.

One important avenue to investigate, we think, involves the cyclical behavior of real marginal cost. Figure 1 presents sets of cross-correlations that help frame the issue. The data are quarterly from 1960:1-1997:4 and HP-detrended. The top panel, discussed earlier, displays the cross-correlation of inflation (the percent change in the GDP deflator) with the output gap (i.e., detrended log GDP). The middle one compares the output gap and the labor income share (our measure of real marginal costs), while the last one looks at the labor share and inflation.

Among other things, the figure makes clear why real unit labor costs outperforms the output gap in the estimation of the new Phillips curve. As the top panel indicates, the output gap leads inflation, rather than vice-versa, in direct contradiction of the theory (see equation (7)). In contrast, as the third panel indicates, real unit labor costs exhibit a strong contemporaneous correlation with inflation. Further, lagged inflation is positively correlated with current unit labor costs, consistent with the

[27] Sbordone (1998) similarly finds that inflation is well explained by a discounted stream of future real marginal costs, though using a quite different methodology to parametrize the model.

[28] An exception is Lown and Rich (1997). Because they augment a traditional Phillips curve with the growth in nominal unit labor costs, their equation fares much better than the standard formulation. Though the way unit labor costs enters our formulation is quite different, it is similarly the sluggish behavior of unit labor costs that helps the model explain recent inflation.

theory. Thus, (with the benefit of this hindsight), it is perhaps not surprising why real unit labor costs enters the structural inflation equation significantly and with the right sign. The middle panel completes the picture: the labor income share lags the output gap in much the same way as does inflation. The lag in the response of real unit labor costs explains why the output gap performs poorly in estimates of the new Phillips curve.

It is also true that the sluggish behavior of real marginal cost might help account for the slow response of inflation to output and thus (possibly) why disinflations may entail costly output reductions.[29] For this reason, modifying existing theories to account for the rigidities in marginal costs suggested by Figure 1 could offer important insights for inflation dynamics.[30] Given the link between unit labor costs and marginal costs, a candidate source for the necessary friction is wage rigidity. Indeed, a likely reason for the strong counterfactual contemporaneous positive correlation between output and real marginal cost in the standard sticky price framework is the absence of any type of labor market frictions [see, e.g., the discussion in Christiano, Eichenbaum and Evans (1997)]. At this stage, one cannot rule out whether it is nominal or real wage rigidities that can provide the answer. Both seem worth exploring.

[29] Interestingly, Blanchard and Muet (1992) find that disinflations in France have been associated with declines in real unit labor costs. In this respect it seems worth exploring data from other countries.

[30] The existing literature on business cycle models that features sticky prices has long emphasized the need to incorporate real rigidities (see, e.g., Blanchard and Fischer (1989) and Ball and Romer (1990)). Typically, however, the discussion is in terms of trying to explain a large response of output to monetary policy: Real rigidities help flatten the short run marginal cost curve. However, it is also the case, as we have been arguing, that real rigidities may be needed to account for inflation dynamics, and in particular the sluggish response of inflation to movements in output.

REFERENCES

Ball, Laurence (1994): "What Determines the Sacrifice Ratio ?," in G. Mankiw, ed., *Monetary Policy,* Chicago University Press.

Ball, Laurence (1995), "Disinflation with Imperfect Credibility," *Journal of Monetary Economics* 35, 5-24.

Ball, Laurence and David Romer (1990), "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies,* 183-204.

Blanchard, Olivier and Pierre Alain Muet (1992), "Competitiveness Through Disinflation: An Assessment of the French Macro-Strategy," mimeo, M.I.T.

Blanchard, Olivier and Stanley Fischer (1989), *Lectures on Macroeconomics,* M.I.T. Press.

Benabou, Roland, (1992), "Inflation and Markups: Theories and Evidence from the Retail Trade Sector, *European Economic Review,* 36, 566-574.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist, (1998), "The Financial Accelerator in a Quantitative Business Cycle Framework," in the *Handbook of Macroeconomics*, John Taylor and Michael Woodford eds.

Buiter Willem and Ian Jewitt, "Staggered Wage Setting with Real Wage Relativities: Variations on a Theme of Taylor," in *Macroeconomic Theory and Stabilization Policy (*University of Michigan Press, Ann Arbor), 183-199.

Calvo, Guillermo A. (1983): "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, 12, 383-398.

Campbell, John Y. and N. Gregory Mankiw (1989), "Consumption, Income and Interest Rates: Re-interpreting the Times Series Evidence," *NBER Macroeconomics Annual,* edited by Olivier Blanchard and Stanley Fischer (M.I.T. Press).

Campbell, John Y. and Robert Shiller (1987), "Cointegration and Tests of the Present Value Relation," *Journal of Political Economy* 95, 1062-1088.

Chadha, Bankim, Paul R. Masson and Guy Meredith (1992), "Models of Inflation and the Costs of Disinflation," *IMF Staff Papers,* vol.39, no. 2, 395-431.

Chari, V.V., Patrick Kehoe and Ellen McGratten, (1996), "Sticky Price Models of the Business Cycle: The Persistence Problem," manuscript, University of Minnesota.

Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (1997), "Sticky Price and Limited Participation Models: A Comparison," *European Economic Review,* 41, 1201-1249.

Clarida, Richard and Mark Gertler(1997), "How the Bundesbank Conducts Monetary Policy," in *Reducing Inflation: Motivation and Strategies,* Christina and David Romer editors.

Clarida, Richard, Jordi Gali, and Mark Gertler, (1997a), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics,* forthcoming.

Clarida, Richard, Jordi Gali, and Mark Gertler (1997b): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature,* forthcoming.

Erceg, Christopher J., Dale W. Henderson and Andrew Levin (1998), "Output-Gap and Price Inflation Volatilities: Reaffirming Tradeoffs in an Optimizing Model," mimeo, Federal Reserve Board.

Fischer, Stanley, (1997): "Long Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy* 85, 163-190.

Fuhrer, Jeffrey C., and George R. Moore (1995): "Inflation Persistence", *Quarterly Journal of Economics*, No. 440, February, pp 127-159.

Fuhrer, Jeffrey C. (1997): "The (Un)Importance of Forward-Looking Behavior in Price Specifications," *Journal of Money, Credit, and Banking* 29, no. 3, 338-350.

Fuhrer, Jeffrey C., Geoffrey Moore, and Scott Schuh (1995): "Estimating the Linear-Quadratic Inventory: Maximum Likelihood versus Generalized Method of Moments," *Journal of Monetary Economics* 35, 115-157.

Galí, Jordi (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, March

Ghezzi, Piero, (1997), "Backward Looking Indexation, Credibility and Inflation Persistence," John Hopkins University, mimeo.

Goodfriend, Marvin, and Robert King (1997): "The New Neoclassical Synthesis and the Role of Monetary Policy" *NBER Macroeconomics Annual*, forthcoming.

Gordon, Robert J.(1996), "Time-Varying NAIRU and its Implications for Economic Policy, *NBER Working Paper No. 5735.*

Jeanne, Olivier, (1998), "Generating Real Persistent Effects of Monetary Shocks: How Much Nominal Rigidity Do We Really Need?" *European Economic Review* no. 6, 1009-1032.

Kiley, Michael T.(1997), "Staggered Price Setting and Real Rigidities," Federal Reserve Board FEDS working paper no. 1997-46.

Kimball, Miles S. (1995): "The Quantitative analytics of the Basic Neomonetarist Model," Journal of Money, Credit, and Banking, vol. 27, no. 4, 1241-1278.

King, Robert G., and Alexander L. Wolman, (1996): "Inflation Targeting in a St. Louis Model of the 21st Century" *NBER Working Paper No. 5507.*

King, Robert G. and Alexander L. Wolman, (1998): "What Should The Monetary Authority Do When Prices Are Sticky?" forthcoming in *Monetary Policy Rules,* J. Taylor, ed, The University of Chicago Press.

King Robert G. and Mark Watson (1994), "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," *Carnegie-Rochester Conference on Public Policy, 157-219.*

Lown Cara S. and Robert W. Rich (1997), "Is There An Inflation Puzzle," *Federal Reserve Bank of New York Quarterly Review, 51-69.*

McCallum, Bennett T., and Edward Nelson (1998): "Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model," forthcoming in *Monetary Policy Rules,* J. Taylor, ed, The University of Chicago Press.

Roberts, John M. (1997): "Is Inflation Sticky? ", *Journal of Monetary Economics*, No. 39, pp 173-196.

Roberts, John M. (1998): "Inflation Expectations and the Transmission of Monetary Policy," mimeo, Federal Reserve Board.

Rotemberg, Julio, and Michael Woodford (1997a): "An Optimization-Based Econometric Framework for the Evolution of Monetary Policy," mimeo

Rotemberg, Julio and Michael Woodford (1997b): "Interest Rate Rules in a Estimated Sticky Price Model," mimeo, Princeton University.

Sbordone, Argia, M., "Prices and Unit Labor Costs: Testing Models of Pricing," mimeo, Princeton University.

Svensson, Lars E. O., (1997a), "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets", *European Economic Review* 41, June, pp. 1111-47.

Svensson, Lars E. O., (1997b), "Inflation Targeting: Some Extensions", *NBER Working Paper No. 5962*, March.

Taylor, John B. (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, 1-23.

Watson, Mark, (1993). "Measures of Fit For Calibrated Models," *Journal of Political Economy,* 101-141.

Woodford, Michael (1996): "Control of the Public Debt: A Requirement for Price Stability?", *NBER Working Paper No. 5684*, July .

Yun, Tack (1996): "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles", *Journal of Monetary Economics*, No. 37, 1996, pp 345-370.

Table 1
Estimates of the New Phillips Curve

	θ	β	λ
GDP Deflator			
(1)	0.829	0.926	0.047
	(0.013)	(0.024)	(0.008)
(2)	0.884	0.941	0.021
	(0.020)	(0.018)	(0.007)
Restricted β			
(1)	0.829	1.000	0.035
	(0.016)		(0.007)
(2)	0.915	1.000	0.007
	(0.035)		(0.006)
NFB Deflator			
(1)	0.836	0.957	0.038
	(0.015)	(0.018)	(0.008)
(2)	0.884	0.967	0.018
	(0.023)	(0.016)	(0.008)

Note: Table 1 reports GMM estimates of the structural parameters of equation (15). Rows (1) and (2) correspond to the two specifications of the orthogonality conditions found in equations (18) and (19) in the text, respectively. Estimates are based on quarterly data and cover the sample period 1960:1-1997:4. Instruments used include four lags of inflation, labor income share, long-short interest rate spread, output gap, wage inflation, and commodity price inflation. A 12 lag Newey-West estimate of the covariance matrix was used. Standard errors are shown in brackets.

Table 2
Estimates of the New Hybrid Phillips Curve

	ω	θ	β	γ_b	γ_f	λ
GDP Deflator						
(1)	0.265	0.808	0.885	0.252	0.682	0.037
	(0.031)	(0.015)	(0.030)	(0.023)	(0.020)	(0.007)
(2)	0.486	0.834	0.909	0.378	0.591	0.015
	(0.040)	(0.020)	(0.031)	(0.020)	(0.016)	(0.004)
Restricted β						
(1)	0.244	0.803	1.000	0.233	0.766	0.027
	(0.030)	(0.017)		(0.023)	(0.015)	(0.005)
(2)	0.522	0.838	1.000	0.383	0.616	0.009
	(0.043)	(0.027)		(0.020)	(0.016)	(0.003)
NFB Deflator						
(1)	0.077	0.830	0.949	0.085	0.871	0.036
	(0.030)	(0.016)	(0.019)	(0.031)	(0.018)	(0.008)
(2)	0.239	0.866	0.957	0.218	0.755	0.015
	(0.043)	(0.025)	(0.021)	(0.031)	(0.016)	(0.006)

Note: Table 2 reports GMM estimates of parameters of equation (26). Rows (1) and (2) correspond to the two specifications of the orthogonality conditions found in equations (27) and (28) in the text, respectively. Estimates are based on quarterly data and cover the sample period 1960:1-1997:4. Instruments used include four lags of inflation, labor income share, long-short interest rate spread, output gap, wage inflation, and commodity price inflation. A 12 lag Newey-West estimate of the covariance matrix was used. Standard errors are shown in brackets.

Table 3

Robustness Analysis: Extra Inflation Lags

	ω	θ	β	ψ	γ_b	γ_f	λ
GDP Deflator	0.244 (0.062)	0.860 (0.025)	0.772 (0.054)	0.090 (0.040)	0.231 (0.050)	0.628 (0.033)	0.033 (0.007)
Restricted β	0.291 (0.039)	0.787 (0.023)	1.000	-0.025 (0.014)	0.270 (0.028)	0.729 (0.021)	0.029 (0.006)
NFB Deflator	0.018 (0.041)	0.922 (0.023)	0.779 (0.050)	0.208 (0.058)	0.019 (0.043)	0.767 (0.046)	0.022 (0.007)
					-		

Note: Table 3 reports GMM estimates of a version of equation (26) with three extra lags of inflation added. ψ represents the sum of the coefficients of the extra lags. using the specification of the orthogonality conditions found in equation (27) in the text. Estimates are based on quarterly data and cover the sample period 1960:1-1997:4. Instruments used include four lags of inflation, labor income share, long-short interest rate spread, output gap, wage inflation, and commodity price inflation. A 12 lag Newey-West estimate of the covariance matrix was used. Standard errors are shown in brackets.

Table 4

Robustness Analysis: Subsample Stability

	ω	θ	β	γ_b	γ_f	λ
GDP Deflator						
60:1-79:4	0.244	0.770	0.892	0.245	0.691	0.054
	(0.027)	(0.017)	(0.027)	(0.022)	(0.016)	(0.007)
70:1-89:4	0.222	0.756	0.820	0.234	0.653	0.07
	(0.026)	(0.014)	(0.036)	(0.024)	(0.028)	(0.010)
80:1-97:4	0.116	0.843	0.773	0.123	0.696	0.051
	(0.022)	(0.007)	(0.030)	(0.021)	(0.027)	(0.006)
Restricted β						
60:1-79:4	0.233	0.753	1.000	0.236	0.763	0.047
	(0.022)	(0.016)		(0.018)	(0.015)	(0.006)
70:1-89:4	0.196	0.734	1.000	0.211	0.788	0.060
	(0.026)	(0.017)		(0.025)	(0.018)	(0.007)
80:1-97:4	0.116	0.843	1.000	0.339	0.539	0.020
	(0.022)	(0.007)		(0.020)	(0.026)	(0.004)
NFB Deflator						
60:1-79:4	-0.043	0.799	0.948	-0.057	1.001	0.066
	(0.022)	(0.012)	(0.012)	(0.030)	(0.015)	(0.008)
70:1-87:4	0.066	0.785	0.913	0.078	0.846	0.913
	(0.018)	(0.013)	(0.020)	(0.020)	(0.017)	(0.020)
80:1-97:4	0.219	0.823	0.778	0.219	0.638	0.049
	(0.023)	(0.008)	(0.046)	(0.023)	(0.034)	(0.005)

Note: Table 4 reports GMM estimates of parameters of equation (26) for alternative sample periods, using the specification of the orthogonality conditions found in equation (27) in the text. Estimates are based on quarterly data. Instruments used include four lags of inflation, labor income share, long-short interest rate spread, output gap, wage inflation, and commodity price inflation. A 12 lag Newey-West estimate of the covariance matrix was used. Standard errors are shown in brackets.

Figure 1: Dynamic Cross-Correlations

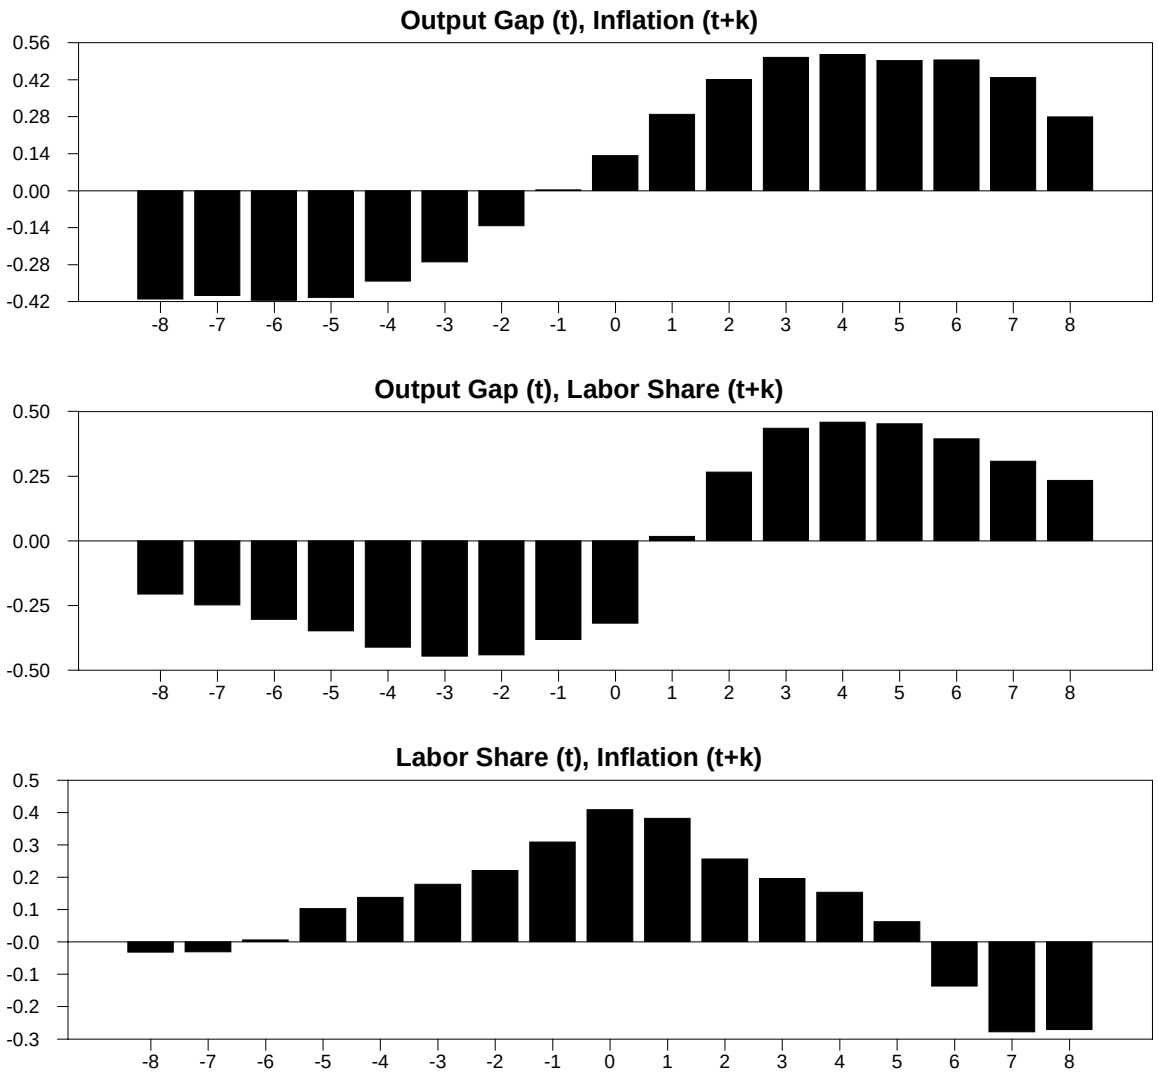

Figure 2. Inflation: Actual vs. Fundamental