

European Inflation Dynamics*

Jordi Galí[†] Mark Gertler[‡] J. David López-Salido[§]

May 2000

Preliminary and Incomplete
(first draft: March 2000)

Abstract

We provide evidence on the fit of the New Phillips Curve (NPC) for the Euro area over the period 1970-1998, and use it as a tool to compare the characteristics of European inflation dynamics with those observed in the U.S. We also analyze the factors underlying inflation inertia by examining the cyclical behavior of marginal costs, as well as that of its two main components, namely, labor productivity and real wages. Some of the findings can be summarized as follows: (a) the NPC fits Euro area data very well, possibly better than U.S. data, (b) the degree of price stickiness implied by the estimates is substantial, but in line with survey evidence and U.S. estimates, (c) inflation dynamics in the Euro area appear to have a stronger forward-looking component (i.e., less inertia) than in the U.S., (d) labor market frictions, as manifested in the behavior of the wage markup, appear to have played a key role in shaping the behavior of marginal costs and, consequently, inflation. in Europe.

JEL Classification no.: E31

Keywords: Inflation, Phillips Curve, EMU

*Under preparation for the International Seminar on Macroeconomics, to be held in Helsinki, June 16-17, 2000. We have benefited from comments by seminar participants at the Bank of Spain, ECB, UCL, LBS, Rochester, CEMFI, and CFS. The views expressed here are those of the authors and do not represent the view of the Bank of Spain.

[†]New York University, Universitat Pompeu Fabra, and Bank of Spain

[‡]New York University

[§]Bank of Spain

1 Introduction

The past three decades have witnessed large and persistent fluctuations in the rate of inflation in most industrialized economies. After an era of relative price stability, inflation began to rise in the late sixties, reached two-digit levels in the mid-seventies, before receding gradually in the early eighties. Since then it has remained relatively subdued, though inflationary pressures (late eighties, today) have alternated with prospects of deflation (late nineties). This characterization of inflation applies equally well to the new Euro area. While in broad respects, the inflationary experience of the Euro area is not very different from that of North America or Japan, the issue of European inflation is of distinct interest given the formation of the new European Central Bank (ECB). The explicit mission of the ECB is the preservation of price stability. To this end, analysis of the sources and nature of inflation in the Euro is a rather immediate and central task.

In this spirit, our paper proposes a simple model of inflation dynamics for the Euro area and assesses its empirical performance.¹ The model builds on recent theoretical work aimed at developing a structural model of inflation. In particular, the framework is based on staggered nominal price setting by monopolistically competitive firms. Aggregating across the price-setting decisions of individual firms leads to an aggregate relation between inflation and real sector activity, referred to as “the new Phillips curve.” A key aspect of this new theory is the attempt to explain the dynamics of inflation without appealing to arbitrary lags, as in traditional Phillips curve analysis.

As we discuss below, in its primitive form, the new Phillips curve relates inflation to the movements in average real marginal cost. That is, the real marginal cost is the theoretically appropriate measure of real sector inflationary pressures; as opposed to cyclical measures used in traditional Phillips curve analysis, such as detrended output or unemployment. Work by Galí and Gertler (1999, henceforth GG) and Sbordone (1998) shows that the new Phillips based on real marginal cost does a reasonably good job of accounting for post-war inflation in the U.S. A virtue of real marginal cost - which in the most basic case corresponds to real unit labor costs - is that it directly accounts for the influence of productivity on inflation, as well as wage pressures. In this paper we show that the marginal-cost based Phillips curve also works well for the Euro area.

Another relevant finding from GG is that real marginal cost is not well approximated by detrended output. It is for this reason that empirical specifications of the new Phillips curve based on detrended output do not fit the data. In particular, detrended output and other such measures fail to account adequately for supply shocks and other factors, such as labor market frictions, that affect firms marginal cost. As we show, a similar phenomenon is true for Europe. Indeed, we show that for the Euro area labor market frictions have likely been a very important factor in the dynamics

¹A recent attempt to model inflation dynamics, though with a different perspective, in the Euro area is Coenen and Wieland (2000).

of marginal cost.

In section 2 we provide a background discussion of the use of old versus new Phillips curve in the context of the Euro area. Section 3 develops the theoretical model used for estimation. Section 4 presents empirical results for the Euro area, and draws a comparison for the U.S.. Among other things, we show that the estimated baseline model tracks the actual Euro inflation very well. In section 5, we present a simple decomposition of real marginal cost in order to understand the forces that have driven this variable. We show that labor market frictions have played an important role in the Euro area both at the medium and high frequencies, in a way that is highly compatible with the anecdotal evidence. Section 6 concludes.

2 Old vs. New Phillips Curves: Some Background

We first analyze European inflation from the perspective of the traditional Phillips and then motivate the use of new Phillips curve. We then describe in general terms the use of the new Phillips curve, and justify the use of real marginal cost instead of detrended output the appropriate variable driving inflation.

2.1 The Traditional Phillips Curve

The traditional Phillips curve relates inflation to some cyclical indicator plus lagged values of inflation. For example, let π_t denote inflation and $\hat{y}_t$ the log deviation of real GDP from its long run trend. A common specification of the traditional Phillips curve is:

$$\pi_t = \sum_{i=1}^h \varphi_i \pi_{t-i} + \delta \hat{y}_{t-1} + \epsilon_t \quad (1)$$

where ϵ_t is a random disturbance. Often the restriction is imposed that the sum of the weights on lagged inflation is unity, so that the model implies no long run trade-off between output and inflation. Sometimes the equation includes additional lags of the output. Alternative specifications may use different cyclical indicators (e.g., the unemployment rate, capacity utilization, etc.)

Despite considerable criticism, however, the traditional Phillips curve does a reasonable job of characterizing post war inflation in the U.S. For example, Rudebusch and Svensson (1999, henceforth RS) show that a variant of equation (1) with four lags of inflation fits well quarterly U.S. data over the period 1960-1999.² The output term enters significantly with a positive sign and the sum of the coefficients on lagged inflation does not differ significantly from unity.

²See Stock and Watson (1999) for a more general analysis. In particular, the authors show that many real activity variables suggested in traditional Phillips curve analysis remain helpful in forecasting inflation.

Here we show that the traditional Phillips curve similarly appears to provide a reasonable description of inflation in the Euro area, over the available sample. To measure inflation we use the log difference of the GDP deflator. The output term is the log of real GDP, detrended with a fitted quadratic function of time. Estimates of the RS specification of equation (1) for quarterly Euro area data over the sample 1970:1-1999:4 yield:

$$\pi_t = \underset{(0.087)}{0.520} \pi_{t-1} + \underset{(0.073)}{0.233} \pi_{t-2} - \underset{(0.084)}{0.070} \pi_{t-3} + \underset{(0.086)}{0.256} \pi_{t-4} + \underset{(0.065)}{0.205} \hat{y}_{t-1} + \epsilon_t$$

For comparison, estimates of the model for U.S. data over the same sample yield:

$$\pi_t = \underset{(0.041)}{0.602} \pi_{t-1} + \underset{(0.153)}{0.041} \pi_{t-2} + \underset{(0.119)}{0.152} \pi_{t-3} + \underset{(0.055)}{0.155} \pi_{t-4} + \underset{(0.055)}{0.192} \hat{y}_{t-1} + \epsilon_t$$

Not only does the RS specification appear to work well for the Euro area, the estimated coefficients are quite similar to those obtained for U.S. data.

Despite the apparent empirical success of the traditional Phillips curve, however, there are two basic concerns: The first, of course, is that the Lucas critique remains an issue, as it has been for the past the past twenty-five years. A central banker deciding the course of policy would ideally like to use equation (1) to calculate the co-movement of inflation and output under different monetary rules. The reliability of this calculation will depend on whether the estimated coefficients are likely to remain stable as the policy regime varies. Instability is a strong possibility on a priori grounds, as we now understand, given that the coefficients on lagged inflation may very well embed expectations of future inflation. This issue is potentially of particular concern in the Euro area, to the extent the ECB signifies a brand new policy regime.

The second basic concern involves the ability of the traditional Phillips curve to explain recent data. This concern is related to the first in the sense that it involves the stability of the relationship over time. In particular, in both the U.S. and Europe, inflation has been low despite high GDP levels relative to trend, owing to robust growth. As a result, traditional Phillips curve relations have been over-predicting inflation. Some observers have simply pronounced the death of the Phillips curve. Others have noted that by making some ex post adjustments (e.g., changing the measure of potential output, adjusting for certain types of supply shocks) it is possible to resurrect the basic relation. In either case, the lesson remains that an empirically based Phillips curve that does a reasonable job of accounting for the past, need not continue to do well in the future.

All this suggests that structural modeling of inflation is desirable, in the same way it is desirable for all other aspects of a macroeconomic framework. A literature has developed in recent years with this objective in mind, as we discuss next.

2.2 The New Phillips Curve

A common aspect of recent theoretical work on inflation dynamics is the derivation of a Phillips curve-like relation explicitly from individual optimization. The approach is based on staggered nominal price setting, in the spirit of Taylor's (1980) seminal work. A key difference is that price setting behavior evolves directly from optimization by firms. In particular, firms choose prices optimally subject to constraints on the frequency of price adjustment. Aggregating across the decision rules of firms then leads to an aggregate Phillips curve equation that relates inflation to cyclical activity. The measure of cyclical activity, further, is tied explicitly to the underlying theory.

A popular example is based on Calvo's model (1983) of staggered price setting. The virtue of the Calvo framework is its parsimony. Here we outline the key aspects, and defer some of the details relevant for an explicit derivation of an estimable relation to Section 3.1 below.

In the standard setting, firms are monopolistic competitors who set prices in staggered fashion based on the expected path of future marginal costs. Aggregating across the price setting behavior of individual firms yields the following relation between inflation, expected future inflation and marginal cost.

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \lambda \widehat{mc}_t \quad (2)$$

where $\widehat{mc}_t$ is average real marginal cost, in percent deviation from its steady state level.³ The slope coefficient λ depends on primitive parameters of the model, particularly the parameter that governs the degree of price rigidity. β is a subjective discount factor.

We refer to equation (2) as the “primitive formulation” of the new Phillips curve. We choose this terminology because the expression makes clear that the movement of real marginal cost relative to trend is the primitive measure of cyclical activity that underlies inflationary pressure. More precisely, the measure of inflation pressure stemming from real sector activity is a discounted stream of expected real marginal costs. To see explicitly, iterating equation (2) forward to obtain

$$\pi_t = \lambda \sum_{k=0}^{\infty} \beta^k E_t\{\widehat{mc}_{t+k}\} \quad (3)$$

Because their prices may be locked in for a period time, firms set prices based on the expected future path of marginal costs. Because not all firms are changing price at the same time, the movement in marginal costs affects the rate of inflation, as opposed to the price level.

³Under the assumption of an isoelastic demand and monopolistic competition, the level of real marginal costs in the (zero-inflation) steady state will coincide with the (constant) real marginal cost that would obtain in an economy with fully flexible prices (and which corresponds to the reciprocal of the “frictionless” markup).

What is most often seen in the literature, however, is the “standard formulation” of the new Phillips curve that instead relates inflation to an *output gap* variable. The latter is defined as the percent deviation of output from its “natural” level, i.e., the level that would obtain under fully flexible prices. This specification arises by combining the primitive formulation given by equation (2) with an equation that links the output gap with real marginal cost. In particular, under certain restrictions on technology (see, e.g., Rotemberg and Woodford (1997)), within a local neighborhood of the steady state real marginal costs are proportionately related to the output gap as follows,

$$\widehat{mc}_t = \delta (y_t - y_t^*) \quad (4)$$

where y_t and y_t^* are the logarithms of real output and the natural level of real output, respectively. In this instance, accordingly, the movement in the conventional output gap exactly mirrors the movement in the theoretically appropriate measure of cyclical inflationary pressure, given by real marginal cost.

Combining (2) with (4) then yields the standard formulation of the new Phillips curve.

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \kappa (y_t - y_t^*) \quad (5)$$

where $\kappa = \lambda\delta$. As with the traditional Phillips curve, inflation varies positively with the output gap. In contrast to the traditional Phillips curve, however, inflation is an entirely forward looking phenomenon, without any appeal to arbitrary lags of inflation. Iterating equation (5) forward yields:

$$\pi_t = \kappa \sum_{k=0}^{\infty} \beta^k E_t\{(y_{t+k} - y_{t+k}^*)\} \quad (6)$$

Inflation depends on the expected path of the output gap. Future movements in the output gap signal future movements in real marginal costs, given equation (4).

2.3 Marginal Cost vs Output Gap: Empirical Issues

As we noted in the introduction, a number of authors have emphasized that the standard formulation of the new Phillips curve, given by equation (5), is vastly at odds with U.S. postwar data. As stressed by GG, however, the empirical breakdown appears to involve mainly the assumed proportionate link between real marginal costs and the output gap. In contrast, the evidence in both Sbordone (1998) and GG (1999) suggests that the central aspect of the theory – the forward looking staggered price setting mechanism inherent in the primitive formulation of the new Phillips curve – is largely consistent with the U.S. evidence. In other words, for U.S. data, inflation appears to respond to anticipated movements in real marginal cost, as implied by equation (3).

We now show that the standard formulation based on the output gap is similarly at odds with European data. Beyond being useful for understanding the gap between

the standard theory and the European data, the exercise justifies our subsequent focus on the marginal cost based formulation of the new Phillips curve.

The essential difficulty with the standard output gap-based formulation is that it implies inflation should anticipate future movements in the output gap (see equation (6)). In the data, however, the reverse seems to be true: appears to lead inflation—as the traditional Phillips curve suggest—, at least when detrended GDP is used as a proxy for the output gap. To see precisely the problem, note that assuming $\beta \approx 1$, equation (5) may be expressed as follows:

$$\pi_t = \pi_{t-1} - \kappa (y_{t-1} - y_{t-1}^*) + \varepsilon_t \quad (7)$$

with $\varepsilon_t = \pi_t - E_{t-1}\{\pi_t\}$. Thus the theory implies that current inflation should be negatively related to the lagged output gap, implying a simple testable implication. We know from the estimates of the traditional Phillips curve that this implication is generally rejected, since lagged output typically enters with a positive coefficient (see, e.g. the estimates of the RS model in the previous section.)

To test that implication directly, we estimate equation (7), imposing the restriction that lagged inflation enters with a coefficient of unity and approximating y_t^* with a quadratic trend. Given that the forecast error ε_t is orthogonal to information at $t-1$, least squares yields a consistent estimate of the coefficient of the output gap. The estimated equation for the Euro area is:

$$\pi_t = \pi_{t-1} + \underset{(0.068)}{0.070} (y_{t-1} - y_{t-1}^*) + \varepsilon_t$$

while for the U.S. we have:

$$\pi_t = \pi_{t-1} + \underset{(0.048)}{0.099} (y_{t-1} - y_{t-1}^*) + \varepsilon_t$$

In neither case is the estimated coefficient on the lagged output gap significantly negative. Indeed, the opposite is true: The lagged output gap enters with a positive sign in each. Again, the results are wholly consistent with the evidence for the traditional Phillips, described in the previous section.⁴

We thus conclude that for the Euro area, as for the U.S., the new Phillips curve based on the output gap cannot explain the data. Based on the U.S. evidence, however, it is reasonable to believe a priori that the problem is likely involves the use an output gap measure to proxy real marginal costs, rather than the basic theory underlying the model of inflation.

Essentially there exist two basic reasons why the real marginal costs may not move proportionately with conventional measures of the output gap. First, detrended log

⁴Attempts to rescue the standard output gap formulation of the new Phillips curve by adding in lagged inflation along with expected future inflation have not been particularly successful. Typically lagged inflation is quantitatively far more important than expected future inflation. See, e.g., Fuhrer (1997).

GDP may be a poor measure of the output gap. The implicit assumption in using detrended output is that the natural level of output $\{y_t^*\}$ can be represented as a smooth, deterministic function of time. To the extent that there are significant real shocks to the economy (e.g. shifts in technology growth, supply shocks, etc.), this assumption is not reasonable. In this instance, appropriately measured real marginal costs may provide a more accurate measure of the true output gap. A second possibility, is that even if the output gap is correctly measured, it may not be the case that real marginal cost moves proportionately, as assumed. As we discuss later, for example, this relation is unlikely to hold if there are labor frictions, either in the form of real or nominal wage rigidities. Indeed, in section 5 we provide evidence that labor market frictions were an important factor in the dynamics of marginal cost for both the Euro area and the U.S., though with some important differences across the two regions.

3 A Structural Phillips Curve Based on Real Marginal Cost

In this section we derive a structural relation between inflation and real marginal cost across firms that we estimate in the subsequent section. As in GG, we first present a baseline model. We then derive a hybrid model that allows for a fraction of firms to set prices using a backward looking rule of thumb. Here the idea is to test the baseline model explicitly against the alternative that arbitrary lags of inflation are required to explain inflation, as in the traditional Phillips curve analysis. One difference from GG is that we relax the assumption that firms face identical constant marginal costs (which greatly simplifies aggregation), and instead allow for increasing real marginal cost, following Woodford (1996) and Sbordone (1998). We choose this path because allowing marginal cost to vary across firms produces more plausible estimates of the degree of price rigidity in the Euro.

3.1 The Baseline Model

We assume a continuum of firms indexed by $j \in [0, 1]$. Each firm is a monopolistic competitor and produces a differentiated good $Y_t(j)$, that it sells at nominal price $P_t(j)$. Firm j faces an isoelastic demand curve for its product, given by $Y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} Y_t$, where Y_t and P_t are aggregate output and the aggregate price level, respectively. Suppose also that the production function for firm j is given by $Y_t(j) = A_t N_t(j)^{1-\alpha}$, where $N_t(j)$ is employment and A_t is a common technological factor.

Firms set nominal prices on a staggered basis, following the approach in Calvo (1983): Each firm resets its price only with probability $1 - \theta$ each period, independently of the time elapsed since the last adjustment. Thus, each period a measure $1 - \theta$ of producers reset their prices, while a fraction θ keep their prices unchanged. Accordingly, the expected time a price remains fixed is $\frac{1}{1-\theta}$. Thus, the parameter

θ provides a measure of the degree of price rigidity. It is one of the key structural parameters we seek to estimate.

After appealing to the law large numbers and log-linearizing the price index around a zero inflation steady state, we obtain the following expression for the evolution of the (log) price level p_t as function of (the log of) the newly set price p_t^* and the lagged (log) price p_{t-1} .

$$p_t = (1 - \theta) p_t^* + \theta p_{t-1} \quad (8)$$

Because there are no firm-specific state variables, all firms that change price in period t choose the same value of p_t^* . A firm that is able to reset in t chooses price to maximize expected discounted profits given technology, factor prices and the constraint on price adjustment (defined by the reset probability $1 - \theta$). It is straightforward to show that an optimizing firm will set p_t^* according to the following (approximate) log-linear rule:

$$p_t^* = \log \mu + (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \{ \log MC_{t,t+k}^n \} \quad (9)$$

where β is a subjective discount factor, $MC_{t,t+k}^n$ is the nominal marginal cost in period $t + k$ of a firm that last reset its price in period t , and $\mu \equiv \frac{\varepsilon}{\varepsilon-1}$ is the firm's desired gross markup. Intuitively, the firm sets price as a markup over a discounted stream of expected future nominal marginal cost. Note that in the limiting case of perfect price flexibility ($\theta = 0$), $p_t^* = \log \mu + \log MC_t^n$: price is just a fixed markup over current marginal cost. As the degree of price rigidity (measured by θ) increases, so does the expect time the price is likely to remain fixed. As a consequence, the firm places more weight on expected future marginal costs in choosing current price.

The goal now is to find an expression for inflation in terms of an observable measure of aggregate marginal cost. Cost minimization implies that the firm's real marginal cost will equal the real wage divided by the marginal product of labor. Given the Cobb-Douglas technology, the real marginal cost in $t + k$ for a firm that optimally sets price in t , $MC_{t,t+k}$, is given by

$$MC_{t,t+k} = \frac{(W_{t+k}/P_{t+k})}{(1 - \alpha) (Y_{t,t+k}/N_{t,t+k})}$$

where $Y_{t,t+k}$ and $N_{t,t+k}$ are output and employment for a firm that has set price in t at the optimal value P_t^* . Individual firm marginal cost, of course, is not observable in the absence of firm level data. Accordingly it is helpful to define the observable variable “average” marginal cost, which depends only on aggregates, as follows

$$MC_t \equiv \frac{(W_t/P_t)}{(1 - \alpha)(Y_t/N_t)} \quad (10)$$

Following Woodford (1996) and Sbordone (1999), we exploit the assumptions of a Cobb-Douglas production technology and the isoelastic demand curve introduced to obtain the following loglinear relation between $MC_{t,t+k}$ and MC_t :

$$\widehat{mc}_{t,t+k} = \widehat{mc}_{t+k} - \frac{\varepsilon\alpha}{1-\alpha} (p_t^* - p_{t+k}) \quad (11)$$

where $\widehat{mc}_{t,t+k}$ and $\widehat{mc}_{t+k}$ are the log deviations of $MC_{t,t+k}$ and MC_{t+k} from their respective steady state values. Intuitively, given the concave production function, firms that maintain a high relative price will face a lower marginal cost than the norm. In the limiting case of a linear technology ($\alpha = 0$), all firms will be facing a common marginal cost.

We obtain the primitive formulation of the new Phillips curve that relates inflation to real marginal cost by combining equations (8), (9), and (11),

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \lambda \widehat{mc}_t \quad (12)$$

with

$$\lambda \equiv \frac{(1-\theta)(1-\beta\theta)(1-\alpha)}{\theta [1+\alpha(\varepsilon-1)]} \quad (13)$$

Note that the slope coefficient λ depends on the primitive parameters of the model. In particular, λ is decreasing in the degree of price rigidity, as measured by θ , the fraction of firms that keep their prices constant. A smaller fraction of firms adjusting prices implies that inflation will be less sensitive to movements in marginal cost. Second, λ is also decreasing in the curvature of the production function, as measured by α , and in the elasticity of demand ε . The larger α and ε , the more sensitive is the marginal cost of an individual firm to deviations of its price from the average price level: everything else equal, a smaller adjustment in price is desirable in order to offset expected movements in average marginal costs.

Finally, we observe that equation (12) can be expressed completely in terms of observables, since (10) implies that average real marginal costs correspond to real unit labor costs (or, equivalently, to the labor income share). In the end, accordingly, the model suggests that inflation should equal a discounted stream of expected future real unit labor costs.

3.2 The Hybrid Model

Equation (12) is the baseline relation for inflation that we estimate. An alternative to equation (12) is that inflation is principally a backward looking phenomenon, as suggested by the strong lagged dependence of this variable in traditional Phillips curve analysis. As a way to test the model against this alternative, we follow GG by considering a hybrid model that allows a fraction of firms to use a backward looking rule of thumb. Accordingly, a measure of the departure of the pure forward looking

model from the data in favor of the traditional approach is the estimate of the fraction of firms that are backward looking.

All firms continue to reset price with probability $1 - \theta$. However, only a fraction $1 - \omega$ resets price optimally, as in the baseline Calvo model. The remaining fraction ω choose the (log) price p_t^b according to the simple backward looking rule of thumb:

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

where p_{t-1}^* is the average reset price in $t - 1$ (across both backward and forward looking firms). Backward looking firms see how firms set price last period and then make a correction for inflation, using lagged inflation as the predictor. Note that though the rule is not optimization based, it converges to the optimal rule in the steady state.⁵

In analogy to the baseline case, the only difference here from GG is that we relax the assumption of constant marginal cost across firms. We defer the details of the derivation to an appendix and simply report the resulting hybrid version of the marginal cost based Phillips curve:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t\{\pi_{t+1}\} + \lambda \widehat{mc}_t \quad (14)$$

with

$$\lambda \equiv \frac{(1 - \omega)(1 - \theta)(1 - \beta\theta)(1 - \alpha)}{\phi [1 + \alpha(\varepsilon - 1)]} \quad ; \quad \gamma_b \equiv \omega\phi^{-1} \quad ; \quad \gamma_f \equiv \beta\theta\phi^{-1}$$

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

As in the pure forward looking baseline case, relaxing the assumption of constant marginal cost affects only the slope coefficient on average marginal cost. The coefficients γ_b and γ_f are the same as in the hybrid model of GG. In this regard, note that the hybrid model nests the baseline model in the limiting case of no backward looking firms (i.e., $\omega = 0$). Accordingly, if the baseline model is true, ω should not differ significantly from zero.

4 Estimates for the Euro Area

We next present estimates of both the baseline model (equation (12)) and the hybrid model (equation (14)) for the Euro area. For comparison, we also present results for the U.S. over the same sample period.

All data are quarterly time series over the period 1970:I-1998:II. To measure inflation we use the GDP deflator. Figure 1 plots that variable, as well as detrended

⁵Note also that backward looking firms free ride off of optimizing firms to the extent that p_{t-1}^* is influenced by the behavior of forward looking firms. In this regard, the welfare losses from following the rule need not be large, if the fraction of backward looking firms is not too dominant.

GDP. Our measure of average real marginal cost is the log of real unit labor costs, consistent the theory presented on section 3.1. Accordingly, we use the log deviation of real unit labor costs from its mean as a measure of $\widehat{mc}_t$. Figure 2 displays our measure of real marginal cost together with that of inflation for the Euro area. Both variables move closely together, at least at medium frequencies. The relation appears to hold throughout the three key phases of the sample: (i) the high inflation of the 1970s and early 1980s; (ii) the disinflation of the early 1980s; and the current period of low inflation. This informal evidence provides some encouragement that inflation is related to movements in marginal costs along the lines that the theory suggests.

We now proceed to provide formal evidence of this conjecture. First, we present estimates of the model, including estimates of the key structural parameters. We then show that the pure baseline not only survives against the hybrid model but does a good job overall for accounting for the dynamics of inflation in the Euro area.

4.1 Baseline Model Estimates

We begin by presenting estimates of the coefficients in equation (12). We refer to these estimates as “reduced form” since we do not try to identify the primitive parameters that underlie the slope coefficient λ . We then proceed to the structural version of the model and, in particular, obtain an estimate of the key underlying primitive parameter θ , which governs the degree of price rigidity.

4.1.1 Reduced Form Estimates

Our econometric procedure is relatively straightforward. Let $\mathbf{z}_t$ denote a vector of variables observed at time t . Then, under rational expectations, equation (12) defines the set of orthogonality conditions:

$$E_t\{(\pi_t - \beta \pi_{t+1} - \lambda \widehat{mc}_t) \mathbf{z}_t\} = 0$$

Given these orthogonality conditions, we can estimate the model using generalized method of moments (GMM). Our vector of instruments $\mathbf{z}_t$ includes four lags of inflation, the real marginal cost (i.e., real unit labor costs), the output gap, and wage inflation. The estimated inflation equation for the Euro area is given by:

$$\pi_t = \underset{(0.033)}{0.917} E_t\{\pi_{t+1}\} + \underset{(0.006)}{0.014} \widehat{mc}_t$$

where standard errors are shown in parentheses. The corresponding equation for the U.S. is:

$$\pi_t = \underset{(0.023)}{0.923} E_t\{\pi_{t+1}\} + \underset{(0.011)}{0.024} \widehat{mc}_t$$

Overall, empirical model works well in both cases. The slope coefficient on marginal cost is positive in each case, as implied by the model. The standard errors

suggest some imprecision in the point estimate, but the coefficient in each case are significantly different from zero. The estimate of the discount factor for the Euro area is a bit low, but is within the realm of reason, especially after taking into account the standard error. The same is true for the U.S.

To illustrate that the connection between inflation and real marginal cost is not simply a product of some kind of aggregation bias, we present evidence from country level annual data. Figure 3 plots GDP inflation versus marginal cost (again measured by the log labor share) for a number of OECD countries, including the member Euro countries, as well the UK, Australia and the U.S. In virtually every case, there is a close movement between inflation and marginal cost, as the theory suggests.

By way of contrast, when we estimate the model using detrended log GDP (as a proxy for the output gap, following other authors), the slope coefficient becomes the wrong sign:

$$\pi_t = \underset{(0.024)}{0.991} E_t\{\pi_{t+1}\} - \underset{(0.009)}{0.004} \hat{y}_t$$

and the corresponding equation for the U.S. yields the same conclusion:

$$\pi_t = \underset{(0.031)}{0.994} E_t\{\pi_{t+1}\} - \underset{(0.008)}{0.022} \hat{y}_t$$

Thus, our focus on real marginal cost in favor of the output gap appears justified.

4.1.2 Structural Estimates

We next estimate the structural parameter θ , which measures the extent of price rigidity. As equation (13) indicates, the reduced form coefficient λ is a function not only of θ and β , but also of the technology curvature parameter α and the elasticity of demand ϵ . The model's restrictions allow us to identify only two primitive parameters: β , the slope coefficient on expected inflation in equation (12), as well as one other parameter among θ , α , and ϵ . Our strategy is to estimate θ and β , conditional on a set of plausible values for α and ϵ .

We obtain measures of α and ϵ , based on information about the steady values of the average markup of price over marginal cost, μ_t and of the labor income share $S_t \equiv W_t N_t / P_t Y_t$. By definition, the average markup equals the inverse of average real marginal cost (i.e., $\mu_t = 1/MC_t$) it follows from our assumptions about technology that:

$$\alpha = 1 - \frac{S_t}{\mu_t}$$

We can accordingly pin down α using estimates of steady state (sample mean) values of the labor income share and the markup. Given an estimate of the steady state markup μ we can obtain a value for ϵ by observing that, given our assumptions,

the steady state markup should correspond to the desired or frictionless markup, implying the relationship which allows us to identify ε .

$$\varepsilon = \frac{\mu}{\mu - 1}$$

We can now feed values of S and μ in the two equations above to obtain measures of α and ε . For the Euro area the average labor share is approximately 3/4; for the U.S. it is approximately 2/3. Unfortunately there is more controversy over the size of the average markup μ . We therefore consider two alternative values which cover the range of plausible estimates: 1.1 and 1.4.⁶

We next define the constant $\gamma \equiv \frac{1-\alpha}{1+\alpha(\varepsilon-1)} \in (0, 1)$, which is conditional on the calibrated values for α and ε . Given this definition, we can express the slope coefficient on real marginal cost, λ in equation (12), as the following function of γ :

$$\lambda \equiv \theta^{-1}(1 - \theta)(1 - \beta\theta) \gamma.$$

In the estimation that follows we treat γ as known with certainty, which permits us to identify β and θ .

Before proceeding, note that the restrictions we impose to identify θ are highly nonlinear (see equation (13)). As is well known, nonlinear estimation using GMM is sometimes sensitive to the way the orthogonality conditions are imposed.⁷ For this reason, following GG, we consider two alternative specifications of the orthogonality conditions, which we refer to, respectively, as specifications 1 and 2:

$$E_t\{(\theta \pi_t - \theta\beta \pi_{t+1} - (1 - \theta)(1 - \beta\theta)\gamma \widehat{mc}_t) \mathbf{z}_t\} = 0$$

$$E_t\{(\pi_t - \beta \pi_{t+1} - \theta^{-1}(1 - \theta)(1 - \beta\theta)\gamma \widehat{mc}_t) \mathbf{z}_t\} = 0$$

Table 1 reports estimates of the baseline model for the Euro area, as well as the U.S. For each region, we report estimates conditional on two different values of the markup, as discussed above. Further, in each instance we report estimates based on the two different specifications of the orthogonality conditions. The first two columns report the estimates of the two primitive parameters, θ and β . The third column reports the implied estimate for λ , the reduced form slope coefficient on real marginal cost. The final column, labeled duration, reports the average number of quarters a price is fixed, corresponding to the estimate of θ . Throughout, standard errors are reported in brackets.

Overall, the parameter estimates of the baseline model using Euro area data are plausible and reasonably robust across the different specifications. The estimated duration of price rigidity lies somewhere around three to four quarters in the low

⁶See, e.g., Rotemberg and Woodford (1995)

⁷See, e.g., Fuhrer, Moore, and Schuh (1995) for a discussion.

markup case, four to five quarters in the high markup case. The estimate of the discount factor β is again a bit low, but not terribly so. Importantly, the implied value of λ is positive and significant across all specifications. Thus, the results suggest that real marginal cost is indeed a significant determinant of inflation, as the theory suggests. Finally, the estimates are fairly similar across specifications (1) and (2), though (1) tends to generate a somewhat lower estimate of the degree of price rigidity (and hence a higher estimate of the slope coefficient λ).

The estimates for the U.S are similar. If anything, they suggest that prices are less rigid. The implied average duration of price rigidity is roughly two to three quarters in the low markup case, versus three to four quarters in the high markup case. We add parenthetically that these numbers are smaller than the roughly five quarter estimates in GG. The difference, of course, is that here we relax the restriction on constant marginal cost. As a consequence, our estimates of the degree of price rigidity are virtually the same as in Sbordone (1999), even though the estimation procedure is quite different.

Though we do not report the associated statistics here, we cannot reject the model's overidentifying restrictions. This finding lends some additional support for the baseline specification. However, the test is of low power since it does not consider a specific alternative. We next consider the hybrid model, which allows us to test directly against the hypothesis of backward looking inflation inertia.

4.2 Hybrid Model Estimates

We extend the approach described in the previous section to the estimation of the hybrid model (14). We continue to use real unit labor costs to measure the real marginal cost (up to a multiplicative factor). In this instance, we estimate an additional parameter: ω , the fraction of backward looking price setters. As in the previous case, we use calibrated values of α and ϵ , which in this case allow us to identify ω , as well as the price rigidity parameter θ .

As in the baseline case, we consider two alternative specifications of the orthogonality conditions. They are given by specification (1):

$$E_t\{(\phi\pi_t - \phi\omega\pi_{t-1} - \phi\beta\theta\pi_{t+1} - (1-\omega)(1-\theta)(1-\beta\theta)\gamma\widehat{mc}_t)\mathbf{z}_t\} = 0$$

and specification (2):

$$E_t\{(\pi_t - \omega\pi_{t-1} - \beta\theta\pi_{t+1} - \phi^{-1}(1-\omega)(1-\theta)(1-\beta\theta)\gamma\widehat{mc}_t)\mathbf{z}_t\} = 0$$

The parameter γ is the same known function of α and ϵ used in the estimation of the baseline model.

Table 2 reports the resulting estimates, for both specifications and conditional on two alternative calibrations of the markup. In this case, the first three columns report the estimates for the primitive parameters ω , θ and β . The next three give the

implied values of the reduced form parameters, γ_b , γ_f and λ , while the last column again gives the implied average duration of price rigidity.

The estimates imply that backward looking price setting, measured by the size of ω , has been a relatively unimportant factor behind the dynamics of Euro area inflation. This is consistent with GG's evidence for the U.S. If anything, however, backward looking behavior is even less important in the Euro area. Under, specification 1, the estimate of ω , the fraction of backward looking price-setters does not differ significantly from zero. Under specification 2, the fraction rises to about a quarter: statistically significant, but still quantitatively small. The estimates of the other structural parameters, β and θ are plausible and very close to their values for the baseline case. Again, after accounting for standard errors, the estimates appear reasonably robust across the two different specifications of the orthogonality conditions.

Once again, the U.S. estimates look broadly similar to those for the Euro area. It appears, however, that while prices are more flexible (i.e., the average duration of price rigidity is shorter), backward looking behavior is statistically significant, though quantitatively small: the estimates of ω , which range from 0.20 to 0.38 still suggest that forward looking behavior is dominant. The overall results are consistent with GG. The main difference is that allowing for non-constant marginal costs yields a lower estimate of the degree of price rigidity than GG obtained.

We have thus far tested our forward looking model against the hypothesis that inflation lagged one quarter also matters. (Due to the form of backward looking price setting we permit - i.e., price setters look back just one period to adjust current prices.) The reduced form model of Rudebusch-Svensson suggests that additional lags of inflation may matter. One possibility, accordingly, is that we may have biased our test against finding backwardness by not letting the additional lags directly affect inflation.

To check that our results in favor of forward looking behavior are robust, we added three additional lags of inflation to the hybrid model. Table 3 presents the results. Parameter ψ denotes the sum of the coefficients on the three additional lags. Note that for both the Euro area and the U.S., this sum is small and not statistically significant. This result holds across all specifications. Thus, it appears that the structural marginal cost based model can account for the inflation dynamics with little reliance on arbitrary lags of inflation.

4.3 Actual versus Fundamental Inflation

Next we propose, following GG, an informal, but intuitive, way to assess the extent to which our model constitutes a good approximation to the dynamics of inflation in the Euro area.⁸ We consider only the pure forward looking baseline model given by

⁸The test is in the spirit of Campbell and Shiller (1987).

equation (12), since the hybrid model does not yield estimates that are appreciably different.

We next define the concept fundamental inflation π_t^* , which we obtain by iterating equation (12):

$$\pi_t = \lambda \sum_{k=0}^{\infty} \beta^k E_t\{\widehat{mc}_{t+k}\} \equiv \pi_t^* \quad (15)$$

Fundamental inflation π_t^* is a discounted stream of expected future real marginal costs, in analogy to the way a fundamental stock price is a discounted stream of expected future dividends. To the extent our baseline model is correct, fundamental inflation should closely mirror the dynamics of actual inflation.

Since expectations of future marginal costs are not observable we cannot construct a direct measure of π_t^* . Yet, under the maintained hypothesis that the model holds, we can construct an estimate of the right hand side of (15) as follows. Let

$$\mathbf{z}_t = [\widehat{mc}_t, \widehat{mc}_{t-1}, \dots, \widehat{mc}_{t-q}, \pi_t, \pi_{t-1}, \dots, \pi_{t-q}]'$$

for some finite q represent a restricted information set observable to the econometrician. Given that $\pi_t \in \mathbf{z}_t$ it follows from (15) is that:

$$\pi_t^* = \lambda \sum_{k=0}^{\infty} \beta^k E_t\{\widehat{mc}_{t+k} \mid \mathbf{z}_t\} \quad (16)$$

Let $\mathbf{A}$ denote the companion matrix of the VAR(1) representation for $\mathbf{z}_t$. Accordingly, $E_t\{\widehat{mc}_{t+k} \mid \mathbf{z}_t\} = \mathbf{e}_1' \mathbf{A}^k \mathbf{z}_t$, where $\mathbf{e}_1$ is a vector with a 1 in its first position and zeros elsewhere. If the model is correct we have

$$\pi_t^* = \lambda \mathbf{e}_1' (\mathbf{I} - \beta \mathbf{A})^{-1} \mathbf{z}_t$$

Hence, we can construct a measure of fundamental inflation using estimates of λ , β , as well as an estimate of $\mathbf{A}$. Strictly speaking, this constructed measure should coincide with actual inflation (except for sampling error) if (15) is the true model of inflation. Of course, we cannot realistically expect (15) to hold exactly since it is, at best, a good first approximation to reality. The question is then: to what extent observed fluctuations in inflation can be accounted for by our measure of fundamental inflation, i.e., how far is our model from reality?

Figure 4 displays our measure of fundamental inflation for the Euro area together with actual inflation. The measure of fundamental inflation is constructed using the Euro area estimates corresponding to specification (1) in Table 2. Overall, fundamental inflation tracks the behavior of actual inflation quite well, especially at medium frequencies.⁹ In particular, it seems to succeed in accounting for the rise of inflation

⁹Galí and Gertler (1999) obtain a similar finding for the US. Sbordone (1998) also finds that inflation is well explained by a discounted stream of future real marginal costs, though using a quite different methodology to parameterize the model.

in the mid 70s and the subsequent disinflation in the mid 1980s, as well as the current environment of low inflation in spite of high growth.

5 The Cyclical Behavior of Real Marginal Cost: The Role of Labor Market Frictions

In this section we present a simple decomposition of the movement in real marginal cost in order to isolate the factors that drive this variable. Our results suggest that labor market frictions likely played a key role in the evolution of real marginal cost in both the Euro area and the U.S., though in a somewhat different fashion across the two regions. In this vein, the results suggest that labor market frictions may help explain inflation persistence in both cases.

Our decomposition requires some restrictions from theory. Suppose the representative household has preferences given by $\sum_{t=0}^{\infty} \beta^t U(C_t, N_t)$, where $U(C_t, N_t)$ is separable in consumption C_t and labor N_t , and where usual properties are assumed to hold. Without taking a stand on the nature of the labor market (e.g. competitive versus non-competitive, etc.), we can without loss of generality express the link between the real wage and household preferences as follows:

$$\frac{W_t}{P_t} = -\frac{U_{N,t}}{U_{C,t}} \mu_t^w \quad (17)$$

where $-\frac{U_{N,t}}{U_{C,t}}$ is the marginal rate of substitution between consumption and labor. Because that variable is the marginal cost to the household in consumption units of supplying additional labor, the variable μ_t^w is interpretable as the gross wage markup (in analogy to the gross price markup over marginal cost, μ_t). Assuming that the household cannot be forced to supply labor to the point where the marginal benefit $\frac{W_t}{P_t}$ exceeds the marginal cost $-\frac{U_{N,t}}{U_{C,t}}$, we have $\mu_t^w \geq 1$.

Conditional on measures of $\frac{W_t}{P_t}$ and $-\frac{U_{N,t}}{U_{C,t}}$, equation (17) provides a simple way to identify the role of labor market frictions in the wage component of marginal cost. If the labor market were perfectly competitive and frictionless (and there were no measurement problems), then we should expect to observe $\mu_t^w = 1$, i.e., the real wage adjusts to equal the household's true marginal cost of supplying labor. With labor market frictions present, we should expect to see $\mu_t^w > 1$ and also possibly varying over time. Situations that could produce this outcome include: households' having some form of monopoly power in the labor market, staggered long term nominal wage contracting, distortionary taxes, and informational frictions that generate efficiency wage payments.

Using equation (17) to eliminate the real wage in the measure of real marginal cost yields the following decomposition:

$$MC_t = \frac{(W_t/P_t)}{(1-\alpha)(Y_t/N_t)} = -\frac{U_{N,t}/U_{C,t}}{(1-\alpha)Y_t/N_t} \mu_t^w \quad (18)$$

According to equation 18, real marginal cost is the product of two components (i) the wage markup μ_t^w and (ii) the ratio of the household's marginal cost of labor supply to the marginal product of labor, $\frac{-U_{N,t}/U_{C,t}}{(1-\alpha)Y_t/N_t}$. We refer to this latter component as the “inefficiency wedge,” since it is a proportionate measure of output relative to the efficient level of output, .i.e., the one corresponding to the frictionless competitive equilibrium. In general, the inefficiency wedge is unity when output is at potential, and declines monotonically with the ratio of output to potential.¹⁰ For our purposes, the key point is that absent frictions in the labor market, real marginal cost equals the inefficiency wedge, and thus varies positively with output relative to potential. With labor market frictions, however, marginal cost also depends on the wage markup, opening up a possible source of inertia.

Assume that $U(C_t, N_t) = \log C_t - \frac{1}{1+\phi} N_t^{1+\phi}$, implying $U_{C,t} = \frac{1}{C_t}$ and $U_{N,t} = -N_t^\phi$. Log-linearizing equation (18) and ignoring constants, yields an expression for marginal cost and its components that is linear in observable variables:

$$mc_t = \log \mu_t^w + [(c_t + \phi n_t) - (y_t - n_t)] \quad (19)$$

with

$$\log \mu_t^w = (w_t - p_t) - (c_t + \phi n_t)$$

where, as before, lower case variables denote logarithms. The expression $[(c_t + \phi n_t) - (y_t - n_t)]$ is the log linearized inefficiency wedge, with $(c_t + \phi n_t)$ being the marginal cost of labor supply. The parameter, ϕ , further, is the inverse of the elasticity of labor supply.

Before proceeding with the decomposition, it is useful to make precise the implications of the wage markup for inflation dynamics. For simplicity, consider an economy with just consumption goods, so that $c_t = y_t$. In this instance, the inefficiency wedge is related to the output gap according to:

$$(c_t + \phi n_t) - (y_t - n_t) = -\Theta + \delta (y_t - y_t^*)$$

where y_t^* is the now the level of output that would obtain with flexible prices and wages, and $\Theta \equiv \log \mu^w + \log \mu$ is an index of the steady state distortion associated with the existence of market power in both labor and goods markets. It follows from equation (19) that in this case real marginal cost is given by:

$$\widehat{mc}_t = \widehat{\mu}_t^w + \delta (y_t - y_t^*)$$

¹⁰To see, note that when output equals potential, marginal product of labor equals the marginal cost of labor supply, implying that the efficiency wedge is unity. Output below potential means $(1-\alpha)Y_t/N_t > -U_{N,t}/U_{C,t}$, implying that the efficiency wedge is less than unity.

where $\hat{\mu}_t^w \equiv \log(\mu_t^w/\mu^w)$ is the percent deviation of the wage markup from its steady state level. We can combine this expression for real marginal cost with the new Phillips curve given by equation (12) to obtain

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \lambda \hat{\mu}_t^w + \kappa (y_t - y_t^*) \quad (20)$$

with $\kappa = \lambda\delta$. Equation (20) makes clear that the standard formulation of the new Phillips curve based on the output gap is correct only under the assumption of constant wage markups (i.e., $\hat{\mu}_t^w = 0$).

To see the impact on inflation dynamics, iterate equation (20) forward to obtain

$$\pi_t = \sum_{k=0}^{\infty} \beta^k E_t\{\lambda \hat{\mu}_{t+k}^w + \kappa (y_{t+k} - y_{t+k}^*)\}$$

In this instance, inflation depends not only on the expected path of the output gap, but also on the fluctuations in the wage markup. Suppose for example that real wages are sticky, either due to some form of real rigidity, or nominal wage rigidity in conjunction with nominal price rigidities (as in Erceg, Henderson and Levin, (1999)). Suppose further that there is a decline in the output gap, possibly expected to persist for some time. The real wage rigidity will produce a persistent rise in the wage markup, since the output gap (and hence the inefficiency wedge $(c_t + \varphi n_t) - (y_t - n_t)$) decline relative to the wage. As a consequence, the expected path of real marginal cost and thus inflation decline less than they would relative to case of a frictionless labor market. In this way, labor market frictions may help account for the observed inertia in inflation.

We now proceed to decompose (log) real marginal cost into the sum of the (log) wage markup and (log) inefficiency wedge. As is apparent from equation (19), to identify the two components we need information on non-durable consumption per household, c_t , and employment per household n_t , as well as two variables we used earlier: the real wage $(w_t - p_t)$ and average labor productivity $(y_t - n_t)$. For the Euro area, only total consumption is available; however, experimenting with U.S. data suggest that the results are reasonably robust to using total consumption instead of just nondurable. To measure employment per household, we use the log difference between employment and the labor force. Hours are not available, but experimentation with the U.S. data suggests that the results are robust also to this modification. Finally, take as unity our benchmark measure of labor supply elasticity, implying $\varsigma = 1$. The results are robust to variations in labor supply elasticities within a reasonable neighborhood of unity.

Figures 5 and 6 present the decompositions for the Euro area and for the U.S., respectively. The top panel in each case illustrates the behavior of the (log) inefficiency wedge relative (log) real marginal cost and the bottom panel does the same for the (log) wage markup.

For the Euro area, perhaps most striking is the apparent secular upward drift in the wage markup from 1970 to early 1982. This behavior seems consistent with the

popular notion that labor union pressures produced a steady rise in the real wage over this era. The impact of this labor market distortion is mirrored in the steady decline in the inefficiency wedge over the entire period, which is especially apparent from comparing the pre-1982 and post-1982 behavior of this variable. This decline is most likely associated with rising employment (i.e. rising unemployment reduces our measure of n_t , which everything else equal, reduces $(c_t + \varphi n_t)$, the numerator in the inefficiency wedge.)

At the medium run frequency, accordingly, the evolution of marginal cost (our metric for inflationary pressures) in Europe goes as follows: In the early 1970s the economy is operating near full capacity, as measured by the high inefficiency wedge. Inflationary pressures are low, however, due to a low wage markup. Over the period, however, the steady rise in the wage markup produces an overall rise in marginal cost. In the latter half of the sample, however, the wage markup moderates, but a persistent decline in the inefficiency wedge associated with employment stagnation leads to low overall marginal cost, and thus low inflationary pressures. We stress, though, that our sample ends in 1998. Since this time there has been a decline in unemployment and a rise in output growth in the Euro area, without any corresponding rise in inflation. In the context of our analysis, either a declining wage markup or rising productivity (the new economy reaches in Europe?), or some combination of the two could be at work. We look forward to sorting this out in future research.

To be sure, it is likely that cyclical as well as secular forces influenced the joint dynamics of the wage markup and the inefficiency wedge in the Euro area. The sharp drop in the inefficiency wedge during the 1980s is likely a result of the severe recession in Europe at this time. The corresponding sharp rise in the wage markup during the severe downturn of the early 1980s is best explained by wage rigidity. The rise in the wage markup over this period accounts why marginal cost (and hence inflation), responded sluggishly to the recession.

Finally, for the U.S. it appears that mainly cyclical forces have been at work. The inefficiency wedge is closely correlated with the business cycle. The wage markup appears to move inversely with the inefficiency wedge, again suggesting the likelihood of temporary wage rigidity. Accordingly, for the U.S, temporary wage rigidities may provide a way to explain the sluggish response of marginal cost and inflation to cyclical output movements.

One somewhat surprising result is that our decomposition suggests that the moderate behavior of real marginal cost in recent years has been mainly the result of a declining wage markup. Indeed the decline in the wage markup has more than offset a sharp rise in the inefficiency wedge. Indeed, the latter has risen in recent years, despite the rise in labor productivity. Rapid growth in nondurable consumption and labor force participation in the U.S. appears responsible. (i.e. $(c_t + \varphi n_t) - (y_t - n_t)$ has risen despite the rise in $y_t - n_t$ since c_t as well as n_t has risen rapidly.) One possibility is that our simple measure of the households' marginal cost of supplying labor, $(c_t + \varphi n_t)$, is suspect. Beyond the issue of parametric assumptions, there may be ag-

gregation problems. To the extent, for example, it has been concentrated among the wealthy and or retirees, the recent rapid growth in nondurable consumption may not be a good proxy for the movement in a representative worker's marginal utility. Also, our measure of labor force participation does not adjust for demographic factors, as recently emphasized by Shimer (1999). On the other hand, the anecdotal evidence does suggest an easing of wage pressures in the U.S., so the notion of a decline in the wage markup is not unreasonable. In future work we plan to explore these measurement issues in more detail, as well as alternative parametric assumptions.

6 Summary and Conclusions

[to be completed]

Our results suggest that a marginal cost - based new Phillips curve provides a good description of Euro area inflation over the period 1970-1998. The empirical model appears to capture the high inflation of the 1970s, the disinflation of the 1980s, as well as the current environment of low inflation. Further, the model does not require arbitrary lags of inflation in order to do so.

As with the U.S., the sluggish movement in inflation relative to output appears related to sluggish movement in marginal cost. Our decomposition of marginal cost suggests that labor market frictions, as manifested in the behavior of the wage markup, may be critical to dynamics of this variable. In both the Euro area and the U.S. there is a countercyclical element to the behavior of the wage markup, consistent with the presence of wage rigidities. A distinctive feature of the Euro area, however, is an upward drift of the wage markup in the 1970s, consistent with the anecdotal evidence for wage pressures in Europe.

Understanding the determinants of the wage markup appears to be the critical next step. It is possible that the staggered nominal wage (and price) contracting model of Erceg, Henderson and Levin (1999) might account for the high frequency behavior of this markup. Under this approach, the ex post wage markup adjusts countercyclically for essentially the same reason the the baseline sticky price model produces an countercyclical price markup (given a constant desired markup). The sticky nominal wage model, however, is unlikely to provide a full explanation for the Euro area data since it would have difficulty accounting for medium term dynamics of the wage markup, particularly the rise in the 1970s. Here a model of real rigidities (e.g. union pressures, etc.) that accounts for variation in the desired wage markup would seem more appropriate.

REFERENCES [incomplete]

- Calvo, Guillermo A. (1983): "Staggered Prices in a Utility Maximizing Framework", *Journal of Monetary Economics*, 12, 383-398.
- Coenen, Gunter and Volker Wieland (2000): "A Small Estimated Euro-Area Model with Rational Expectations and Nominal Rigidities", mimeo ECB, Frankfurt.
- Campbell, John Y. and Robert Shiller (1987), "Cointegration and Tests of the Present Value Relation", *Journal of Political Economy* 95, 1062-1088.
- Fuhrer, Jeffrey, 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, J. and Mark Gertler (1999), "Inflation Dynamics: A Structural Econometric Analysis", *Journal of Monetary Economics*, 44, 195-222.
- Rotemberg, Julio, and Michael Woodford (1997), "An optimization-based Econometric Framework for the Evaluation of Monetary Policy", in *NBER Macroeconomic Annual 1997*, Ben Bernanke and Julio Rotemberg, eds., 297-346, MIT Press.
- Rotemberg, Julio, and Michael Woodford (1995), "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets", in *Frontiers of Business Cycle Research*, Thomas F. Cooley, ed., 243-293, Princeton University Press, Princeton New Jersey.
- Sbordone, Argia, M. (1998), "Prices and Unit Labor Costs: A New Test of Price Stickiness", IIES Seminar Paper 653, Stockholm University, October.
- Stock, James and Mark Watson (1999), "Forecasting Inflation," *Journal of Monetary Economics*
- Svensson, Lars E.O. and Glenn D. Rudebusch (1998): "Policy Rules for Inflation Targeting", in *Monetary Policy Rules*, John Taylor (ed.), 203-246, The University of Chicago Press.
- Taylor, John B. (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, 1-23.
- Woodford, Michael (1996): "Control of the Public Debt: A Requirement for Price Stability?", *NBER Working Paper No. 5684*, July .

Table 1
New Phillips Curve: Structural Estimates

	θ	β	λ	<i>Duration</i>
<i>Euro Area</i>				
$\mu = 1.1, \alpha = 0.32$				
(1)	0.670 (0.025)	0.819 (0.036)	0.222 (0.007)	3.0 (0.08)
(2)	0.771 (0.040)	0.917 (0.033)	0.086 (0.006)	4.3 (0.17)
$\mu = 1.4, \alpha = 0.46$				
(1)	0.747 (0.022)	0.844 (0.034)	0.120 (0.007)	4.0 (0.09)
(2)	0.815 (0.033)	0.917 (0.033)	0.056 (0.006)	5.4 (0.18)
<i>United States</i>				
$\mu = 1.1, \alpha = 0.40$				
(1)	0.407 (0.046)	0.769 (0.064)	1.001 (0.033)	1.7 (0.08)
(2)	0.660 (0.062)	0.923 (0.023)	0.201 (0.011)	3.0 (0.18)
$\mu = 1.4, \alpha = 0.52$				
(1)	0.556 (0.045)	0.829 (0.053)	0.431 (0.027)	2.2 (0.10)
(2)	0.731 (0.053)	0.923 (0.023)	0.119 (0.011)	3.7 (0.18)

Note: The parameter α was calibrated so $(1-\alpha)$ is equal to the average labor income share divided by the chosen markup (μ). The average labor income shares are taken to be equal to 2/3 for the US (source: Cooley and Prescott (1995)) and 3/4 for the Euro Area (source: European Commission (1999)). Sample Period: 1970-1998.

Table 2
Hybrid Model: Structural Estimates

	ω	θ	β	γ_b	γ_f	λ	<i>Duration</i>
<i>Euro Area</i>							
$\mu = 1.1, \alpha = 0.32$							
(1)	0.044 (0.036)	0.663 (0.026)	0.805 (0.041)	0.063 (0.049)	0.761 (0.050)	0.213 (0.008)	3.0 (0.08)
(2)	0.244 (0.071)	0.769 (0.059)	0.909 (0.051)	0.245 (0.049)	0.702 (0.040)	0.053 (0.006)	4.3 (0.25)
$\mu = 1.4, \alpha = 0.46$							
(1)	0.045 (0.043)	0.743 (0.022)	0.832 (0.038)	0.057 (0.052)	0.790 (0.043)	0.120 (0.007)	3.9 (0.09)
(2)	0.258 (0.073)	0.813 (0.049)	0.907 (0.052)	0.245 (0.049)	0.702 (0.039)	0.035 (0.006)	5.3 (0.26)
<i>United States</i>							
$\mu = 1.1, \alpha = 0.40$							
(1)	0.199 (0.034)	0.382 (0.042)	0.734 (0.067)	0.355 (0.044)	0.500 (0.074)	0.635 (0.020)	1.6 (0.07)
(2)	0.339 (0.055)	0.546 (0.076)	0.891 (0.045)	0.392 (0.039)	0.563 (0.044)	0.178 (0.011)	2.2 (0.17)
$\mu = 1.4, \alpha = 0.52$							
(1)	0.254 (0.041)	0.500 (0.044)	0.759 (0.063)	0.351 (0.043)	0.525 (0.059)	0.320 (0.019)	2.0 (0.09)
(2)	0.384 (0.057)	0.624 (0.068)	0.884 (0.048)	0.392 (0.039)	0.563 (0.042)	0.106 (0.011)	2.7 (0.18)

Table 3
Hybrid Model: Further Inflation Lags

	ω	θ	β	γ_b	γ_f	λ	ψ	<i>Duration</i>
<i>Euro Area</i>								
$\mu = 1.1$, $\alpha = 0.32$	0.057 (0.065)	0.696 (0.060)	0.809 (0.046)	0.077 (0.082)	0.755 (0.049)	0.167 (0.011)	0.024 (0.065)	3.2 (0.20)
$\mu = 1.4$, $\alpha = 0.46$	0.062 (0.071)	0.786 (0.056)	0.812 (0.045)	0.074 (0.081)	0.761 (0.042)	0.087 (0.010)	0.052 (0.063)	4.7 (0.26)
<i>United States</i>								
$\mu = 1.1$, $\alpha = 0.40$	0.197 (0.047)	0.399 (0.047)	0.666 (0.105)	0.345 (0.069)	0.467 (0.097)	0.620 (0.021)	0.081 (0.060)	3.0 (0.31)
$\mu = 1.4$, $\alpha = 0.52$	0.248 (0.058)	0.519 (0.059)	0.682 (0.103)	0.341 (0.069)	0.488 (0.078)	0.322 (0.020)	0.079 (0.059)	3.1 (0.32)

Figure 1. Inflation and Output in the Euro area




Figure 2. Inflation and Marginal Cost in the Euro area




Figure 3a. Inflation and Marginal Cost in OECD countries

Inflation (continuous line) and Marginal Cost (dotted line)




Figure 3b. Inflation and Marginal Cost in OECD countries

Inflation (continuous line) and Marginal Cost (dotted line)




Figure 4. Fundamental Inflation in the Euro area




Figure 5.Components of the Marginal Cost in the Euro area





Figure 6. Components of the Marginal Cost in the U.S.

