

European inflation dynamics

Jordi Galí^{a,b,c,*}, Mark Gertler^b, J. David López-Salido^c

^a*CREI, Universitat Pompeu Fabra, Ramon Trias Fargas 25,
08005 Barcelona, Spain*

^b*Department of Economics, New York University, 269 Mercer St., New York, NY 10003, USA*

^c*Banco de España, Alcalá 50, 28014 Madrid, Spain*

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. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: E31

Keywords: Inflation; Phillips curve; EMU

1. Introduction

Over the postwar, the pattern of inflation in the countries that now constitute the new Euro area has been broadly similar to that in many other industrialized

* Correspondence address: CREI, Universitat Pompeu Fabra, Ramon Trias Fargas 25, 08005 Barcelona, Spain. Tel.: + 34-93-542-27-54; fax: + 34-93-542-17-46.

E-mail addresses: jordi.gali@econ.upf.es (J. Galí), mark.gertler@nyu.edu (M. Gertler), davidl@bde.es (J.D. López-Salido).

nations, including the United States, England and Japan. The issue of Euro area inflation, however, 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 area is a rather immediate and central task.

In this spirit, we propose and estimate a simple theory-based Phillips curve for the new Euro area.¹ We make use of a newly constructed aggregate historical data set for this region. In addition, given our approach, we also necessarily confront the debate over whether recent structural models of inflation – loosely known as ‘new Phillips curves’ – can explain the data, particularly the high degree of persistence in inflation. At issue is the nature of short run inflation dynamics and the associated implications for monetary policy.²

The structural equation for inflation that we estimate for the Euro area is in the spirit of the new Phillips curve literature.³ It evolves explicitly from a model of staggered nominal price setting by monopolistically competitive firms. This formulation has inflation vary positively with real sector economic activity in the short run, similar in spirit to a traditional Phillips curve. One key difference is that, in its primitive form, the new Phillips curve relates inflation to movements in real marginal cost (averaged across firms). That is, real marginal cost is the theoretically appropriate measure of real sector inflationary pressures, as opposed to the cyclical measures used in traditional Phillips curve analysis, such as detrended output or unemployment.

Recently, Sbordone (1999) and Galí and Gertler (1999) (hereafter GG) have shown that this ‘marginal cost-based’ version of the new Phillips curve can provide a reasonable account of postwar inflation in the U.S.⁴ In this paper we show that the same is largely true for the Euro area. That is conditional on the path of real marginal cost, the structural equation captures the pattern of Euro area inflation, including the rise to double digit levels in the 1970s, the disinflation of the 1980s, as well as the current era of relative price stability. A virtue of the real marginal cost measure, which in our analysis corresponds to real unit labor costs, is that it directly accounts for the influence of both productivity and

¹ Coenen and Wieland (2000) also analyze the new Euro area inflation data, using a somewhat different approach.

² On this debate, among others, see Fuhrer and Moore (1995), Roberts (1997), Sbordone (1999), Galí and Gertler (1999) and Mankiw (2000).

³ See Goodfriend and King (1997) for a survey.

⁴ Specifically, these authors obtain sensible and similar estimates of the marginal cost-based new Phillips curve using different methodologies. Though GG reject the pure forward looking model in favor of a hybrid specification in the spirit of Fuhrer and Moore (1995) that allows for a fraction of firms to use rule of thumb pricing, they nonetheless find that the forward looking behavior suggested by the baseline theory remains predominant (see Section 4). Further, Sbordone (1999) finds that the pure baseline model does a good job of tracking the aggregate data, while GG find that a hybrid version with a modest amount of backward looking behavior does the job.

wage pressures on inflation. In this respect, we find that productivity, wages and inflation move together largely as the new Phillips curve theory suggests.

An auxiliary finding is that, as with the U.S. data, real marginal cost in the Euro area is not well approximated by detrended output. This finding is of some significance: It provides a potential explanation for the empirical failure of specifications of the new Phillips curve based on detrended output. Put differently, much of the recent criticism of the new Phillips curve applies to this formulation, and not to the marginal cost-based specification. Among other things, real marginal cost appears to move more sluggishly in the data relative to detrended output. This sluggishness in real marginal cost, in turn, appears to help the model account for the high degree of persistence in inflation.

In part, our results push the mystery of inflation back to understanding the factors that underlie the apparent inertia in the real marginal cost. Given the link between unit labor costs and marginal cost, wage rigidity arises as a possibility. We pursue that hypothesis by presenting a decomposition of the cyclical movement in real marginal cost. We find that for both Euro area and the U.S., wage rigidity was indeed a significant factor in accounting for sluggish cyclical movement in marginal cost. In addition, for the Euro area alone, steady real wage increases from the early 1970s through the early 1980s – possibly emanating from union pressures – placed consistent upward pressure on real marginal cost. This persistent supply shock (in conjunction with accommodating European central banks) likely played a key role in the double-digit inflation and general stagnation in Europe at this time.

In Section 2 we provide a background discussion of the debate over use of old versus new Phillips curves in the context of the Euro area. Section 3 develops the theoretical model used for estimation. In addition to the pure forward looking model, we also consider a hybrid model in the spirit of Fuhrer and Moore (1995) and Galí and Gertler (1999) that allows a fraction of firms to be backward looking. Section 4 discusses some econometric issues and then 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 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 likely have played an important role in the Euro area both at the medium and high frequencies in a way that is compatible with the anecdotal evidence. Concluding remarks are in Section 6.

2. Euro inflation and the Phillips curve debate

We first analyze European inflation from the perspective of the traditional Phillips curve, partly to provide some descriptive evidence and partly to motivate use of the new Phillips curve. We then describe in general terms the new Phillips curve, and briefly discuss the debate over this approach.

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

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

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 detrended 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 Eq. (1) with four lags of inflation fits well quarterly U.S. data over the period 1960–1999⁵. The output term enters significantly with a positive sign and the sum of the coefficients on lagged inflation does not differ significantly from unity.

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 Eq. (1) for quarterly Euro area data over the sample 1970:I–1998:II yield

$$\pi_t = 0.520 \pi_{t-1} + 0.233 \pi_{t-2} - 0.070 \pi_{t-3} + 0.256 \pi_{t-4} + 0.051 \hat{y}_{t-1} + \varepsilon_t.$$

(0.087) (0.073) (0.084) (0.086) (0.016)

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

$$\pi_t = 0.602 \pi_{t-1} + 0.041 \pi_{t-2} + 0.152 \pi_{t-3} + 0.155 \pi_{t-4} + 0.048 \hat{y}_{t-1} + \varepsilon_t.$$

(0.041) (0.153) (0.119) (0.055) (0.014)

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 25 years. That is, the

⁵ See 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.

stability of this equation across policy regimes is unclear, particularly since the coefficients on lagged inflation may very well embed expectations of future inflation. This issue is of particular concern in the Euro area, to the extent that EMU 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.⁶ 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.

2.2. *The new Phillips curve*

The new Phillips curve is based on staggered nominal price setting, in the spirit of Taylor's (1980) seminal work. A key difference is that price setting behavior is the product of optimization by monopolistically competitive firms subject to constraints on the frequency of price adjustment. A popular example is based on Calvo's model (1983) of staggered price setting, which has the virtue of 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.

The basic building block is the following equation that relates inflation π_t to anticipated future inflation and real 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, β is a subjective discount factor, and λ is a slope coefficient that depends on the primitive parameters of the model, particularly the parameter that governs the degree of price rigidity. Eq. (2) is a log-linear approximation of a relation obtained from aggregating across the pricing decisions of individual firms.⁷ This relation is what we referred to in the introduction as the 'primitive formulation' of the new Phillips curve; i.e., it is the formulation that arises

⁶ See, for example, the discussion in Gordon (1998) and Stock (1998).

⁷ As we discuss in Section 3, the new Phillips curve is obtained as log-linear approximation around a deterministic steady-state inflation rate. The implicit assumption is that monetary policy is aimed at obtaining this steady-state rate. Allowing for shifts in the steady-state inflation rate would give us more flexibility in fitting the data, but would raise the problem of trying to explain changes in the central bank's long run target inflation rate.

directly as a consequence of the frictions in the price adjustment process that are the key aspect of the theory.

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. Under certain restrictions on technology and labor market structure (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:

$$\widehat{mc}_t = \delta(y_t - y_t^*), \quad (3)$$

where y_t and y_t^* are the logarithms of real output and the natural level of real output, respectively. Combining (2) with (3) then yields the standard output gap-based formulation of the new Phillips curve:

$$\pi_t = \beta E_t \{ \pi_{t+1} \} + \kappa(y_t - y_t^*), \quad (4)$$

where $\kappa = \lambda\delta$.

It is Eq. (4) that has been the subject of considerable controversy. 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. Iterating Eq. (4) forward yields

$$\pi_t = \kappa \sum_{k=0}^{\infty} \beta^k E_t \{ (y_{t+k} - y_{t+k}^*) \}. \quad (5)$$

A striking implication is the absence of a tradeoff between inflation and output; to the extent a central bank can commit to stabilizing the output gap ($y_{t+k} - y_{t+k}^*$), it can achieve price stability. However, as emphasized by Fuhrer and Moore (1995), GG and others, Eq. (5) is at odds with the data. It suggests that inflation should anticipate movements in the output gap.⁸ Yet, as the estimates of the traditional Phillips curve suggest, the output gap (measured by detrended output) tends to lead inflation.⁹ While this result is widely known to hold for U.S. data, our Phillips curve estimates in the previous section suggest that it applies equally well to the Euro area. Overall, the output-gap based

⁸ Mankiw (2000) has recently emphasized that Eq. (5) predicts that inflation should respond quickly to monetary policy shocks (since it anticipates the response of output), which is counterfactual. Note, however, that this criticism does not extend to the marginal cost-based formulation of the new Phillips curve (Eq. (2)), to the extent marginal costs responds sluggishly to the policy shock, relative to output.

⁹ To see precisely the problem, note that assuming $\beta \approx 1$, Eq. (4) may be expressed as follows:

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

with $u_t = \pi_t - E_{t-1} \{ \pi_t \}$. Thus the theory implies that current inflation should be negatively related to the lagged output gap, in contrast to the evidence.

formulation of the new Phillips curve cannot account for the persistence of inflation either for the U.S. or for the Euro area.

As we noted in the introduction, however, Sbordone (1999) and GG find that the central aspect of the theory, the relation between inflation and real marginal cost given by Eq. (2) is roughly consistent with the data (see footnote 4). These results suggest that it is Eq. (3), the hypothesized link between real marginal cost and the output gap, that is at variance with the data. GG present some direct evidence for U.S. data to show that this is indeed the case. Real marginal cost tends to respond sluggishly and with a lag to movements in the output gap, much as inflation does. There are two possible explanations for this finding. One is that conventional measures of the output gap may be poor. To the extent that there are significant real shocks to the economy (e.g., shifts in technology growth, fiscal shocks, etc.), using detrended output as a proxy for y_t^* may not be appropriate. Whether this factor alone could account for the observed inertia in real marginal cost relative to detrended output is an open question, however.

A second, and perhaps more likely 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. In particular, as we discuss in Section 5, with frictions in the labor market, either, in the form of real or nominal wage rigidities, Eq. (3) is no longer valid. These labor market rigidities, further, can in principle offer a rationale for the inertial behavior of real marginal cost.¹⁰ 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 marginal cost-based Phillips curve

In this section we derive a structural relation between inflation and average 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 (1999). We choose this path because allowing marginal cost to vary across firms produces more plausible estimates of the degree of price rigidity in

¹⁰ As we discuss in Section 5, further, inertial behavior of marginal cost opens up the possibility of a short run tradeoff between inflation and output. See also Erceg et al. (2000).

the Euro area. Our baseline model, accordingly, is exactly the theoretical framework in Sbordone (1999). Our hybrid model is a generalization that extends GG to allow for increasing marginal cost. The appendix provides a detailed derivation.

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 $Y_t(j) = (P_t(j)/P_t)^{-\varepsilon} Y_t$, 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 $1/(1 - \theta)$. 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 of 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} :

$$p_t = (1 - \theta)p_t^* + \theta p_{t-1}. \quad (6)$$

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:

$$p_t^* = \log \mu + (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \{ mc_{t,t+k}^n \}, \quad (7)$$

where β is a subjective discount factor, $mc_{t,t+k}^n$ is the logarithm of nominal marginal cost in period $t + k$ of a firm that last reset its price in period t , and $\mu \equiv \varepsilon/(\varepsilon - 1)$ 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 ($\theta = 0$), $p_t^* = \log \mu + mc_t^n$: 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:¹¹

$$MC_t \equiv \frac{(W_t/P_t)}{(1 - \alpha)(Y_t/N_t)}. \tag{8}$$

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 log-linear relation between $MC_{t,t+k}$ and MC_t :

$$\widehat{mc}_{t,t+k} = \widehat{mc}_{t+k} - \frac{\varepsilon\alpha}{1 - \alpha}(p_t^* - p_{t+k}), \tag{9}$$

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 Eqs. (6), (7), and (9):

$$\pi_t = \beta E_t \{ \pi_{t+1} \} + \lambda \widehat{mc}_t \tag{10}$$

¹¹ Note that this measure allows for supply shocks (entering through A_t in the production). An adverse supply shock, for example, results in a decline in average labor productivity, Y_t/N_t . Also, the specification is robust to the addition of other variable factors (e.g., imported imports), so long as the elasticity of output with respect to labor is constant, firms take wages as given, and there are no labor adjustment costs.

with

$$\lambda \equiv \frac{(1 - \theta)(1 - \beta\theta)(1 - \alpha)}{\theta[1 + \alpha(\varepsilon - 1)]}. \quad (11)$$

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 Eq. (10) can be expressed completely in terms of observables, since (8) implies that average real marginal costs correspond to real unit labor costs (or, equivalently, to the labor income share).¹² 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

Eq. (10) is the baseline relation for inflation that we estimate. An alternative to Eq. (10) 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.

¹²In an earlier version of GG we showed that the results are robust to some alternative measures of marginal cost. See also Sbordone (1999).

Note that though the rule is not optimization based, it converges to the optimal rule in the steady state.¹³

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:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \{ \pi_{t+1} \} + \tilde{\lambda} \widehat{mc}_t \tag{12}$$

with

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

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. Evidence

We next present estimates of both the baseline model (Eq. (10)) and the hybrid model (Eq. (12)) 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. Fig. 1 plots that variable, as well as detrended GDP. Our measure of average real marginal cost is the log of real unit labor costs, consistent with the theory presented on Section 3.1.¹⁴ Accordingly, we use the log deviation of real unit labor costs from its mean as a measure of \widehat{mc}_t .

Fig. 2 displays our measure of real marginal cost together with inflation for the Euro area. Both variables move closely together, at least at medium frequencies.

¹³ Note 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.

¹⁴ Our data for the Euro area are from Fagan et al. (2001). Real unit labor costs are constructed as the ratio of compensation to employees (WIN) to GDP (YER). Inflation is measured as the quarterly percent change in the GDP deflator (YED). The data for the U.S. are described in GG. In particular, real unit labor costs are for the non-farm business sector.

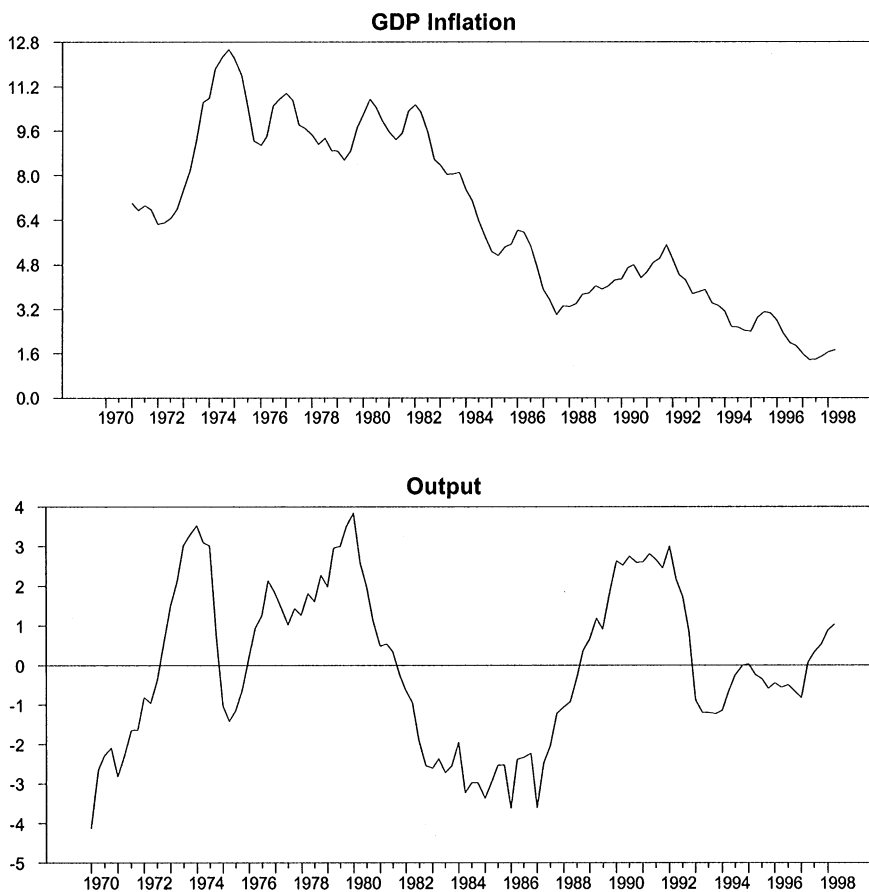


Fig. 1. Inflation and output in the Euro area.

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 (iii) the current period of low inflation.^{15, 16} This informal evidence

¹⁵ Blanchard (1997) and Blanchard and Wolfers (2000) have also drawn attention to the rise and fall in the labor share in Europe over this time, which they interpret as reflecting shifts in the aggregate demand for labor. Also, Blanchard and Muet (1992) draw the connection between movements in the labor share and inflation for the French economy. We pursue this observation of strong co-movement of the labor share with inflation as a central implication of new Phillips curve theories.

¹⁶ One possibility, emphasized by Benabou (1992), is that inflation may be influencing movements in the labor share by affecting firms' desired markup. Our instrument variables procedure controls for this possibility of reverse causality in principle, though it is an issue we plan to investigate further. In the meantime, we observe that much of the movement in the labor share is associated with the wage markup as opposed to the price markup (see Section 5). Accordingly, the issue is whether inflation affects workers' desired markup.




Fig. 2. Inflation (dashed line) and marginal cost (continuous line) in the Euro area.

provides some encouragement that inflation is related to movements in marginal costs along the lines that the theory suggests.¹⁷

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, while the baseline can be formally rejected against a hybrid model with some mild backward-lookingness, it still does a good job at accounting for the dynamics of inflation in the Euro area.

4.1. Baseline model estimates

We begin by presenting estimates of the coefficients in Eq. (10). 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.

¹⁷ We emphasize that the theory suggests that real marginal cost is effectively a measure of capacity utilization. Accordingly, underlying the persistent high inflation in the 1970s is overly accommodative central bank behavior. One possibility is that European central banks did not properly take into account reductions in potential output stemming from high wage increases. We expand on this in Section 5.

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, Eq. (10) 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).

We use instruments dated $t - 1$ or earlier for two reasons: First, there is likely to be considerable error in our measure of marginal cost. Assuming this error is uncorrelated with past information, it is appropriate to use lagged instruments. Second, not all current information may be available to the public at the time they form expectations.

For the Euro area estimates, our vector of instruments \mathbf{z}_t includes five lags of inflation, and two lags of the real marginal cost, detrended output, and wage inflation. We choose a relatively small number of lags for instruments other than inflation in order to minimize the potential estimation bias that is known to arise in small samples when there are too many overidentifying restrictions. We based the lag length for inflation on reduced form forecasting evidence. For the U.S. estimates, the instrument set is the same, except that we only use four lags of inflation, again based on the reduced form evidence.¹⁸

The estimated inflation equation for the Euro area is given by

$$\pi_t = 0.914 E_t\{\pi_{t+1}\} + 0.088 \widehat{mc}_t, \quad (13)$$

(0.040) (0.041)

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

$$\pi_t = 0.924 E_t\{\pi_{t+1}\} + 0.250 \widehat{mc}_t. \quad (14)$$

(0.029) (0.118)

In each instance, the standard errors are modified, using a Newey–West correction, given evidence of serial correlation in the error term, as we discuss below.

We performed a number of diagnostic tests to evaluate these regressions. We begin with the results for the Euro area. To check for potential weakness of the instruments, we perform an F -test applied to the first-stage regression; the results clearly suggest that the instruments used are relevant (F statistic = 61.8, with a p -value = 0.00).¹⁹ Next we test the model's overidentifying restrictions. Based

¹⁸ Adding a fifth lag of inflation to the instrument set does not affect the results.

¹⁹ Recently, Staiger and Stock (1997) point out the importance of examining this statistic, as conventional asymptotic results may break down under weak correlation between the instruments and the endogenous regressor. This is clearly not the case in our estimated equation.

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

4.1.2. Structural estimates

We next estimate the structural parameter θ , which measures the extent of price rigidity. As Eq. (11) 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 Eq. (10), 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 .

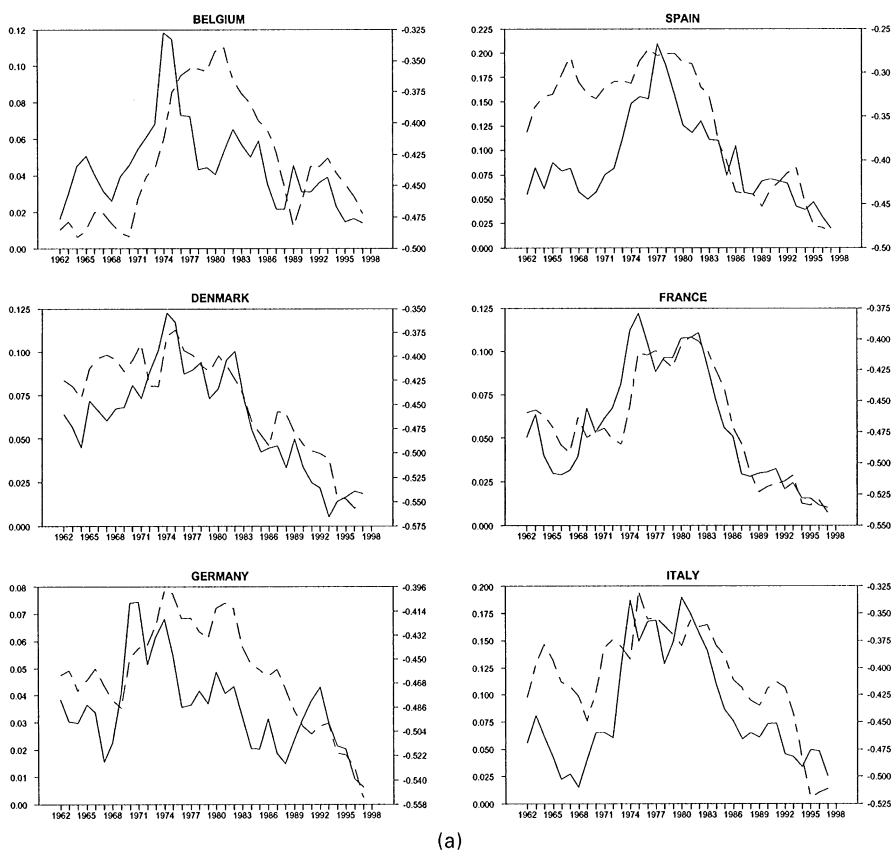


Fig. 3. Inflation (continuous line) and marginal cost (dashed line) in selected OECD countries.

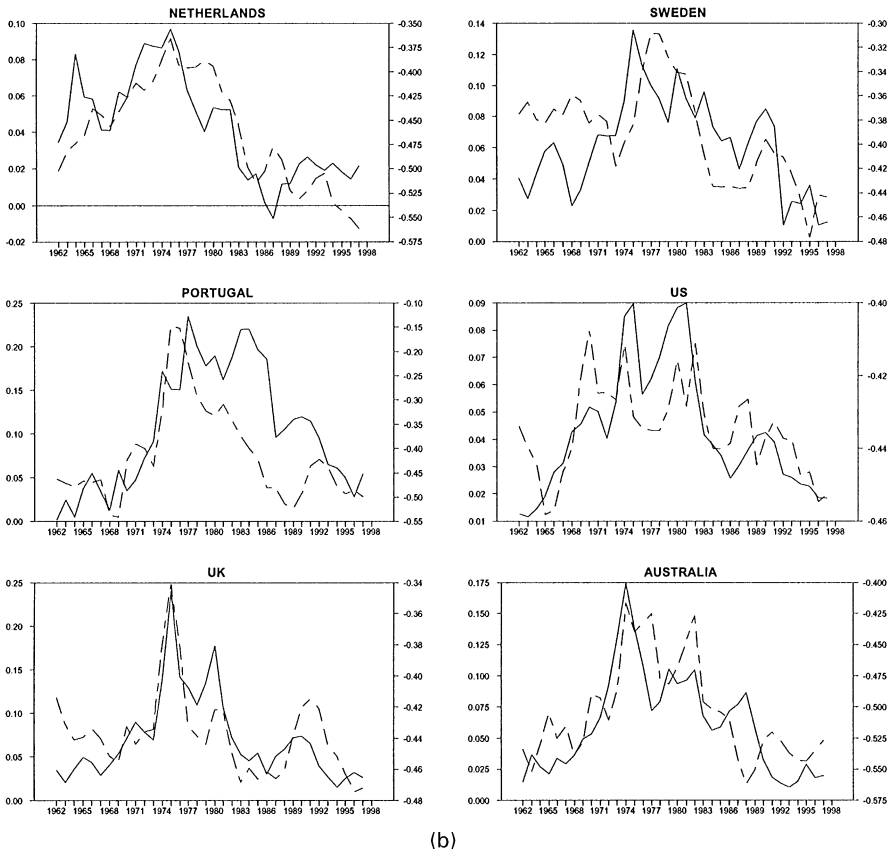


Fig. 3. (continued)

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 thus 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 $\frac{3}{4}$; for the U.S. it is approximately $\frac{2}{3}$.²³ Unfortunately there is more controversy over the size of the average markup μ . Our baseline results are based on an average markup of 1.1.²⁴

We next define the constant $\xi \equiv (1 - \alpha)/(1 + \alpha(\varepsilon - 1)) \in (0, 1)$, 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 Eq. (10), as the following function of ξ :

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

In our baseline estimates below we treat ξ as known with certainty (conditional on the average labor income share and markup) which permits us to identify β and θ . In addition we also report estimates under the assumption of constant returns to scale, which corresponds to $\xi = 1$. In the latter case identification of θ does not require the calibration of any parameter.

Before proceeding, note that the restrictions we impose to identify θ are highly nonlinear (see Eq. (11)). As is well known, nonlinear estimation using GMM is sometimes sensitive to the way the orthogonality conditions are imposed.²⁵ For this reason, and following GG, we consider two alternative specifications of the orthogonality conditions, which we refer to, respectively, as specifications 1 and 2:

$$\begin{aligned} E_t\{(\theta\pi_t - \theta\beta\pi_{t+1} - (1 - \theta)(1 - \beta\theta)\xi\widehat{mc}_t)z_t\} &= 0, \\ E_t\{(\pi_t - \beta\pi_{t+1} - \theta^{-1}(1 - \theta)(1 - \beta\theta)\xi\widehat{mc}_t)z_t\} &= 0. \end{aligned}$$

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 ξ , 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. Next we report the average duration of a price (in quarters), corresponding to the estimate of θ . Standard errors (with a Newey–West correction) for all the parameter estimates are reported in brackets. The final column displays Hansen's J statistic of the overidentifying restrictions, together with the associated p -values (in brackets).

²³ Average labor shares for the Euro area were drawn from European Economy (1999). The value for the U.S. was taken from Cooley and Prescott (1995).

²⁴ An earlier version of the paper considered alternative values within the interval (1.1, 1.4), a range of plausible estimates from the literature (e.g., Rotemberg and Woodford, 1995; Basu and Fernald, 1997). None of the results were affected by that choice.

²⁵ See, e.g., Fuhrer et al. (1995) for a discussion.

Table 1
Structural estimates^a

	Parameters				Test
	θ	β	λ	D	J
<i>Euro area</i>					
$\mu = 1.1, \alpha = 0.32$					
(1)	0.699 (0.026)	0.805 (0.051)	0.228 (0.052)	3.0 (0.08)	9.081 (0.430)
(2)	0.771 (0.043)	0.914 (0.040)	0.088 (0.041)	4.4 (0.19)	8.213 (0.513)
$\xi = 1$					
(1)	0.904 (0.011)	0.886 (0.042)	0.021 (0.007)	10.4 (0.12)	8.506 (0.484)
(2)	0.918 (0.015)	0.914 (0.040)	0.014 (0.006)	12.2 (0.18)	8.214 (0.513)
<i>United States</i>					
$\mu = 1.1, \alpha = 0.40$					
(1)	0.475 (0.060)	0.837 (0.053)	0.665 (0.238)	2.0 (0.11)	7.681 (0.465)
(2)	0.627 (0.067)	0.924 (0.029)	0.250 (0.114)	2.7 (0.18)	5.759 (0.674)
$\xi = 1$					
(1)	0.845 (0.026)	0.910 (0.031)	0.042 (0.015)	6.4 (0.17)	5.845 (0.665)
(2)	0.867 (0.030)	0.924 (0.029)	0.031 (0.014)	7.5 (0.23)	5.760 (0.674)

^aParameter 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 $\frac{2}{3}$ for the U.S. and $\frac{3}{4}$ for the Euro area. Sample Period: 1970–1998. Column D corresponds to the associated sticky prices duration, and J to the Hansen test of the overidentifying restrictions (below in brackets we report the p -value). Instruments for Euro area: inflation $t - 1$ to $t - 5$, output gap, labor income share and wage inflation: $t - 1$ to $t - 2$. Instruments for the U.S.: the same expects inflation from $t - 1$ to $t - 4$.

The first two rows of Table 1 report the baseline estimates using Euro area data. All of them have the right sign and plausible size, and reasonably robust across the two normalizations. The estimated average duration of prices lies somewhere around three to four quarters. 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 for both normalizations. 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). As we suggested earlier, imposing the assumption of constant returns to labor yields an implausibly high estimate of the stickiness parameter and its implied duration.

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 baseline case, versus six to seven quarters in the case of constant returns to labor. It is interesting to notice that our estimates of the degree of price rigidity in the baseline case are very similar to Sbordone (1999), even though the estimation procedure is quite different.

Again, the model's overidentifying restrictions are not rejected under any specification. However, this test is likely to have low power since it does not consider a specific alternative. We next report estimates for the hybrid model introduced above, 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 (12). 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 to calibrate ξ . Now this allows us to identify ω , as well as the price rigidity parameter θ .

Again we consider two alternative specifications of the orthogonality conditions. They are given by

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

and

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

where parameter ξ is the same known function of α and ε used in the estimation of the pure forward looking model, and where $\phi \equiv \theta + \omega[1 - \theta(1 - \beta)]$.

The first three columns of Table 2 report estimates of 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 . Again, we report the implied average duration of price rigidity, and the overidentifying restriction test.

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 that forward looking behavior remained highly important for the U.S. If anything, however, backward looking behavior is 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 somewhere between $\frac{1}{4}$ and $\frac{1}{3}$. The

Table 2
Hybrid model^a

	Parameters						Test	
	ω	θ	β	γ_b	γ_r	λ	D	J
<i>Euro area</i>								
$\mu = 1.1, \alpha = 0.32$								
(1)	0.030 (0.083)	0.668 (0.029)	0.804 (0.056)	0.043 (0.115)	0.773 (0.064)	0.214 (0.079)	3.0 (0.09)	8.983 (0.344)
(2)	0.287 (0.126)	0.787 (0.089)	0.925 (0.069)	0.272 (0.072)	0.689 (0.047)	0.039 (0.049)	4.7 (0.42)	7.484 (0.380)
$\xi = 1$								
(1)	0.024 (0.122)	0.907 (0.015)	0.897 (0.053)	0.025 (0.127)	0.877 (0.045)	0.018 (0.012)	10.0 (0.14)	8.428 (0.393)
(2)	0.335 (0.129)	0.922 (0.031)	0.920 (0.074)	0.272 (0.072)	0.689 (0.044)	0.006 (0.007)	12.8 (0.40)	7.485 (0.380)
<i>United States</i>								
$\mu = 1.1, \alpha = 0.40$								
(1)	0.255 (0.054)	0.498 (0.072)	0.863 (0.056)	0.347 (0.045)	0.584 (0.054)	0.291 (0.139)	2.0 (0.14)	4.993 (0.661)
(2)	0.317 (0.065)	0.569 (0.080)	0.916 (0.042)	0.364 (0.042)	0.599 (0.041)	0.162 (0.093)	2.3 (0.19)	4.216 (0.754)
$\xi = 1$								
(1)	0.400 (0.074)	0.818 (0.038)	0.878 (0.052)	0.339 (0.0046)	0.610 (0.034)	0.026 (0.013)	5.5 (0.21)	4.332 (0.741)
(2)	0.451 (0.075)	0.827 (0.042)	0.898 (0.052)	0.364 (0.042)	0.599 (0.032)	0.020 (0.011)	5.8 (0.24)	4.216 (0.755)

^a See notes to Table 1 for details.

estimates are statistically significant, but still quantitatively small, suggesting that forward looking behavior is dominant in shaping the dynamics of inflation. The estimates of the other structural parameters, β and θ are plausible and very close to their values for the forward looking case. Again, after accounting for standard errors, the estimates appear reasonably robust across the two different specifications of the orthogonality conditions.²⁶

Once again, the U.S. estimates look broadly similar to those for the Euro area, with prices appearing to be more flexible (i.e., the average duration of price rigidity is shorter) in the former. Backward looking behavior is statistically

²⁶ We also detected serial correlation of the error term in the hybrid model, and accordingly adjusted the standard errors. Note that the hybrid model does not necessarily predict a serial uncorrelated error term, since some of the error could be due to backward looking price setting (i.e., the error term in this case is not just a forecast error.)

Table 3
Hybrid model: further inflation lags^a

	Parameters								Test
	ω	θ	β	γ_b	γ_f	λ	ψ	D	J
<i>Euro area</i>									
(1)	0.105 (0.074)	0.669 (0.083)	0.847 (0.059)	0.138 (0.084)	0.742 (0.049)	0.168 (0.095)	-0.037 (0.093)	3.0 (0.25)	6.566 (0.087)
(2)	0.183 (0.101)	0.811 (0.137)	0.863 (0.069)	0.188 (0.083)	0.719 (0.043)	0.048 (0.062)	0.049 (0.077)	5.3 (0.72)	5.928 (0.115)
<i>United States</i>									
(1)	0.265 (0.094)	0.563 (0.111)	0.870 (0.125)	0.328 (0.089)	0.606 (0.066)	0.203 (0.135)	0.044 (0.066)	2.2 (0.25)	2.011 (0.570)
(2)	0.290 (0.103)	0.598 (0.122)	0.899 (0.128)	0.333 (0.088)	0.617 (0.059)	0.151 (0.116)	0.036 (0.065)	2.5 (0.30)	1.566 (0.815)

^aEstimates reported for the case of decreasing returns to labor. See notes to Table 1 for details.

significant, though quantitatively small: the estimates of ω , which range from $\frac{1}{4}$ to $\frac{1}{2}$ are slightly higher than in the Euro area. Notice that allowing for decreasing returns to labor yields lower estimates of both the degree of price rigidity and the fraction of backward looking price setters than those obtained under the constant returns assumption (corresponding to $\xi = 1$).

We have thus far tested our forward looking model against the hypothesis that inflation lagged one quarter also matters.²⁷ One possibility, accordingly, is that we may have biased our test against finding backwardness by not letting additional lags of inflation directly enter our Phillips curve. To examine this possibility, we added several lags of inflation to the hybrid model. Table 3 presents the results for the baseline model with $\mu = 1.1$. Parameter ψ denotes the sum of the coefficients on the additional lags. Note that for both the Euro area and the U.S., this sum is small and not statistically significant.²⁸ This result holds across all specifications. Thus, it appears that the structural marginal cost based model can account for the inflation dynamics with relatively little reliance on arbitrary lags of inflation, as compared to the traditional Phillips curve (see Section 2).

²⁷ Recall that due to the form of backward looking price setting we permit, price setters look back just one period to adjust current prices.

²⁸ For the Euro area, some of the individual lag coefficients were statistically significant, though not large quantitatively.

4.3. *Actual versus fundamental inflation*

Next we propose, following GG and Sbordone (1999), 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.²⁹ We consider only the pure forward looking baseline model given by Eq. (10), 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 Eq. (10):

$$\pi_t = \lambda \sum_{k=0}^{\infty} \beta^k E_t \{ \widehat{mc}_{t+k} \} \equiv \pi_t^* \tag{17}$$

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 (17) 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 (17) that

$$\pi_t^* = \lambda \sum_{k=0}^{\infty} \beta^k E_t \{ \widehat{mc}_{t+k} | \mathbf{z}_t \}. \tag{18}$$

Let \mathbf{A} denote the companion matrix of the VAR(1) representation for \mathbf{z}_t . Accordingly, $E_t \{ \widehat{mc}_{t+k} | \mathbf{z}_t \} = \mathbf{e}'_1 \mathbf{A}^k \mathbf{z}_t$, 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 (17) is the true model of inflation. Of course, we cannot realistically expect (17) 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?

²⁹The test is in the spirit of Campbell and Shiller (1987).




Fig. 4. Fundamental inflation in the Euro area.

Fig. 4 displays our measure of fundamental inflation for the Euro area together with actual inflation. The measure of fundamental inflation is constructed using the estimated reduced form equation for the Euro area, given by Eq. (13). Virtually identical results obtain from using either of the estimated structural equations (specification (1) and (2)) in Table 1. Overall, fundamental inflation tracks the behavior of actual inflation quite well, especially at medium frequencies.³⁰ In particular, it seems to succeed in accounting for the rise of inflation in the mid-1970s 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

³⁰Galí and Gertler (1999) obtain a similar finding for the US, using the estimated hybrid model. Sbordone (1999) 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.

different fashion across the two regions. In this vein, the results suggest that labor market frictions may help explain inflation persistence in both cases.³¹

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:

$$\frac{W_t}{P_t} = - \frac{U_{N,t}}{U_{C,t}} \mu_t^w, \tag{19}$$

where $-U_{N,t}/U_{C,t}$ is the marginal rate of substitution between consumption and labor. Because that variable is the marginal cost to the household of supplying additional labor (in consumption units), 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 W_t/P_t exceeds the marginal cost $-U_{N,t}/U_{C,t}$, we have $\mu_t^w \geq 1$.

Conditional on measures of W_t/P_t and $-U_{N,t}/U_{C,t}$, Eq. (19) 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, nominal wage rigidities, distortionary taxes on labor income, etc.

Using Eq. (19) to eliminate the real wage in the measure of real marginal cost yields the following decomposition:

$$MC_t = \frac{(W_t/P_t)}{(1 - \alpha)(Y_t/N_t)} = - \frac{U_{N,t}/U_{C,t}}{(1 - \alpha)Y_t/N_t} \mu_t^w. \tag{20}$$

According to Eq. (20), 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, $(-U_{N,t}/U_{C,t})/((1 - \alpha)Y_t/N_t)$. 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

³¹ Christiano et al. (1997) also emphasize the need to consider labor market frictions in this kind of framework. Here we provide some direct evidence in favor of this conjecture.

wedge is unity when output is at potential, and declines monotonically with the ratio of output to potential.³² 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 $U(C_t, N_t) = \log C_t - (1/(1 + \varphi))N_t^{1+\varphi}$, implying $U_{C,t} = 1/C_t$ and $U_{N,t} = -N_t^\varphi$. Log-linearizing Eq. (20) and ignoring constants, yields an expression for marginal cost and its components that is linear in observable variables:

$$mc_t = \log \mu_t^w + [(c_t + \varphi n_t) - (y_t - n_t)] \quad (21)$$

with

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

where lower case variables are used to denote logarithms. The expression $[(c_t + \varphi n_t) - (y_t - n_t)]$ is the log-linearized inefficiency wedge, with $(c_t + \varphi n_t)$ 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

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

where y_t^* is now the level of output that would obtain with flexible prices and wages, $\delta \equiv 1 + \varphi$, 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 Eq. (21) that in this case real marginal cost is given by

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

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 Eq. (10) to obtain

$$\pi_t = \beta E_t \{\pi_{t+1}\} + \lambda \hat{\mu}_t^w + \kappa (y_t - y_t^*) \quad (22)$$

with $\kappa = \lambda \delta$. Eq. (22) 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., $\hat{\mu}_t^w = 0$).

³² To 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 inefficiency wedge is less than unity.

To see the impact on inflation dynamics, iterate Eq. (22) 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 et al., 2000). 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 real marginal costs and 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 Eq. (21), 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 $\varphi = 1$. The results are robust to variations in labor supply elasticities within a reasonable neighborhood of unity, and also to allowing for nonseparability of preferences over consumption and leisure.

Figs. 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).

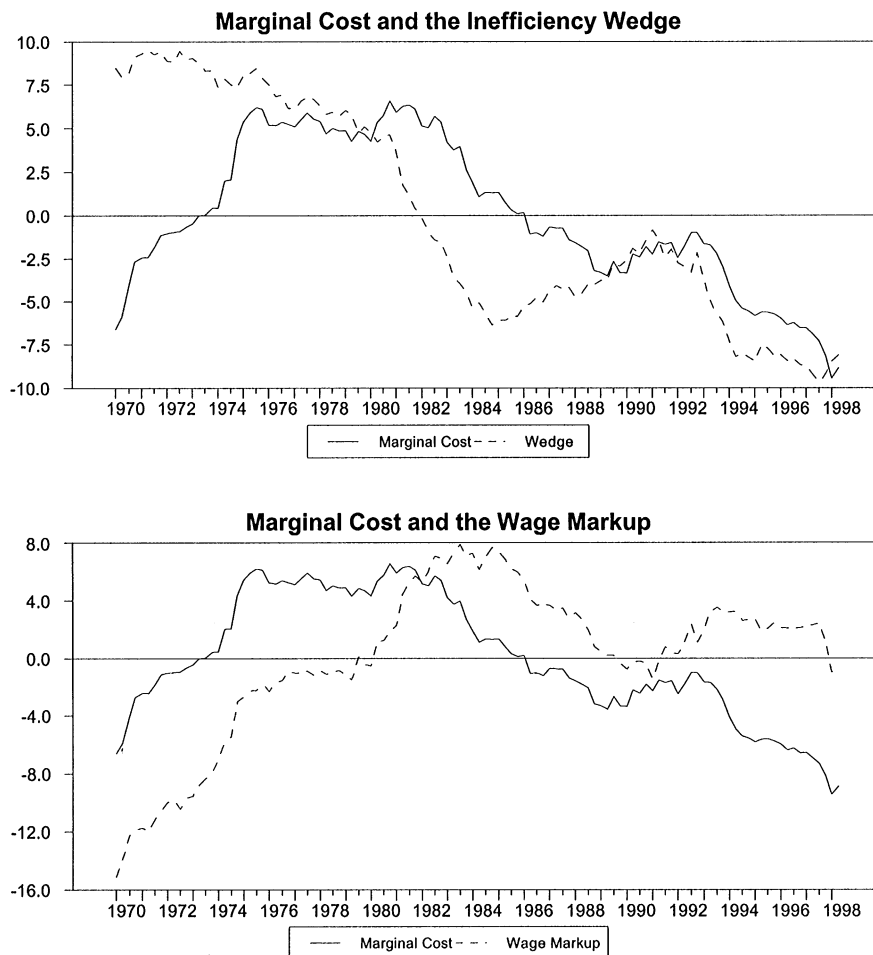


Fig. 5. Components of the marginal cost in the Euro area.

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.³³ Inflationary pressures are low, however, due to a low wage markup. Over the period, however, the steady rise in the wage markup produces an

³³ We stress that the inefficiency wedge is a measure of capacity utilization and not capacity output, i.e., Fig. 5 simply suggests that capacity utilization was high in the 1970s. Indeed, supply shocks in the 1970s, including wage pressures as well as oil shocks, likely had an adverse effect on capacity output. A likely scenario is that European central banks did not properly adjust monetary policy to account for the contraction in capacity output resulting from these shocks.

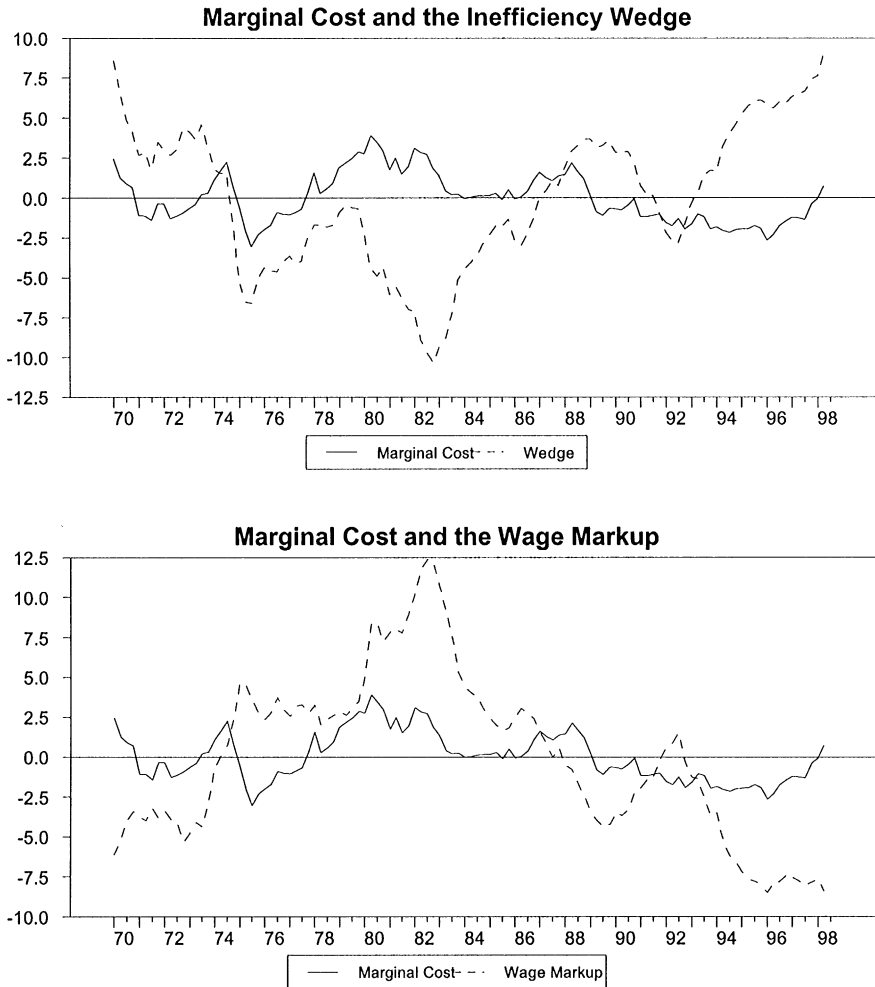


Fig. 6. Components of the marginal cost in the U.S.

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 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 for the U.S. 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 aggregation 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 (1998). 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. Conclusions

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.

As with the U.S., sluggish movement in marginal cost appears to be an important factor accounting for the observed high degree of persistence in inflation. Our decomposition of marginal cost suggests that labor market frictions, as

manifested in the behavior of the wage markup, may be critical to the 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. For one reason or another, European central banks at this time did not properly adjust for the impact of the rise in the wage markup (and other adverse supply shocks) on the natural level of output, which helps account for the persistent high inflation of this era.

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 et al. (2000) 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 baseline sticky price model produces a 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.

Acknowledgements

This paper was prepared for the *International Seminar on Macroeconomics*, hosted by the Bank of Finland in Helsinki, June 16–17, 2000. We thank the comments of our discussants, Francesco Giavazzi, and Jeremy Nalewaik, three anonymous referees, the editor, Harald Uhlig, as well as those provided by Oliver Blanchard, Jeff Fuhrer, Bob King, Argia Sbordone and Jim Stock. We have also benefited from comments of seminar participants at the Banco de Portugal *Conference on Monetary Economics*, Guimaraes, June 12–14, 2000, the Conference on *Monetary Policy Challenges in the 21st Century – A Transatlantic Perspective*, October 2000 Georgetown University, The Bank of Spain, ECB, UCL, LBS, Rochester, CEMFI, and CFS. We thank Ricardo Mestre for providing us with the Euro area data. Galí and Gertler acknowledge financial support from the National Science Foundation, the C.V. Starr Center for Applied Economics, and CREI. The views expressed here are those of the authors and do not represent the view of the Bank of Spain.

Appendix A. Derivation of the hybrid Phillips curve with increasing marginal cost

The log-linearized equations of the model with backward looking firms are given by

Price index p_t :

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

Index of newly re-set prices p_t^* :

$$p_t^* = (1 - \omega) p_t^f + (1 - \theta) p_t^b \quad (\text{A.2})$$

Forward looking re-set price p_t^f :

$$p_t^f = (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \{ \widehat{mc}_{t,t+k} + p_{t+k} \} \quad (\text{A.3})$$

Marginal cost of forward looking firms that re-set price at t , $\widehat{mc}_{t,t+k}$:

$$\widehat{mc}_{t,t+k} = \widehat{mc}_{t+k} - \frac{\varepsilon\alpha}{1 - \alpha} (p_t^f - p_{t+k}), \quad (\text{A.4})$$

where \widehat{mc}_{t+k} is the percent deviation from steady state of average real marginal cost $MC_t \equiv ((W_t/P_t)/(1 - \alpha)(Y_t/N_t))$.

Backward looking re-set price:

$$p_t^b = p_{t-1}^* + \pi_{t-1}, \quad (\text{A.5})$$

where $\pi_t \equiv p_t - p_{t-1}$.

Rearranging Eqs. (23) and (24) yields

$$\begin{aligned} \pi_t &= \left(\frac{1 - \theta}{\theta} \right) (p_t^* - p_t) \\ &= \left(\frac{1 - \theta}{\theta} \right) [(1 - \omega)(p_t^f - p_t) + \omega(p_t^b - p_t)]. \end{aligned} \quad (\text{A.6})$$

We next obtain expressions for $(p_t^f - p_t)$ and $(p_t^b - p_t)$. Let $\pi_{t,t+k} \equiv p_{t+k} - p_t$. Combining (25) and (26) yields

$$\begin{aligned} p_t^f - p_t &= (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \left\{ \widehat{mc}_{t,t+k} - \frac{\alpha\varepsilon}{1 - \alpha} (p_t^f - p_t) \right. \\ &\quad \left. + \beta\theta \left(1 + \frac{\varepsilon\alpha}{1 - \alpha} \right) \pi_{t,t+k+1} \right\} \\ &= (1 - \beta\theta)\xi \sum_{k=0}^{\infty} (\beta\theta)^k E_t \{ \widehat{mc}_{t,t+k} \} + \sum_{k=1}^{\infty} (\beta\theta)^k E_t \{ \pi_{t,t+k} \}, \end{aligned} \quad (\text{A.7})$$

where as in the text, $\xi \equiv ((1 - \alpha)/1 + \alpha(\varepsilon - 1))$.

Combining (23) with (27) yields

$$p_t^b - p_t = \frac{1}{1 - \theta} \pi_{t-1} - \pi_t. \quad (\text{A.8})$$

Next, insert (29) and (30) into (28) to obtain the following expression for inflation:

$$\pi_t = \left(\frac{1-\theta}{\theta}\right) \left[\omega \left(\frac{1}{1-\theta} \pi_{t-1} - \pi_t \right) + (1-\omega) \left[(1-\beta\theta) \xi \sum_{k=0}^{\infty} (\beta\theta)^k E_t \{ \widehat{mc}_{t+k} \} + \sum_{k=1}^{\infty} (\beta\theta)^k E_t \{ \pi_{t+k} \} \right] \right], \tag{A.9}$$

which, after some algebra, can be rewritten in a more compact form as

$$\pi_t = \tilde{\lambda} \widehat{mc}_t + \gamma_f E_t \{ \pi_{t+1} \} + \gamma_b \pi_{t-1}, \tag{A.10}$$

where

$$\lambda \equiv (1-\theta)(1-\beta\theta)(1-\omega) \xi \phi^{-1},$$

$$\gamma_f \equiv \beta\theta \phi^{-1},$$

$$\gamma_b \equiv \omega \phi^{-1},$$

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

Notice that in the absence of backward looking price setting ($\omega = 0$) Eq. (A.10) becomes the pure forward looking marginal cost-based Phillips with increasing marginal cost, as derived by Sbordone (1999). Under the assumption of a constant marginal cost ($\alpha = 0$) the model becomes the hybrid Phillips curve derived in Galí and Gertler (1999).

References

Balakrishnan, R., López-Salido, D., 2000. Understanding UK inflation: The role of structural change. Mimeo., Bank of England and Bank of Spain.

Basu, S., Fernald, J.G., 1997. Returns to scale in US production: Estimates and implications. *Journal of Political Economy* 105, 249–283.

Benabou, R., 1992. Inflation and markups: Theories and evidence from the retail trade sector. *European Economic Review* 36, 566.

Blanchard, O., Muet, P.A., 1992. Competitiveness through disinflation: An assessment of the French macro-strategy. Mimeo., MIT, Cambridge, MA.

Blanchard, O., 1997. The medium run. *Brookings Papers on Economic Activity*, 2, 89–158.

Blanchard, O., Wolfers, J., 2000. The role of shocks and institutions in the rise of European unemployment: The aggregate evidence. *Economic Journal* 110, 1–33.

Benigno, P., López-Salido, D., 2000. Inflation dynamics and optimal monetary policy in Europe. Mimeo., NYU and Bank of Spain.

Calvo, G.A., 1983. Staggered prices in a utility maximizing framework. *Journal of Monetary Economics* 12, 383–398.

- Campbell, J.Y., Shiller, R., 1987. Cointegration and tests of the present value relation. *Journal of Political Economy* 95, 1062–1088.
- Christiano, L., Eichenbaum, M., Evans, C., 1997. Sticky price and limited participation models: A comparison. *European Economic Review* 41, 1201–1249.
- Coenen, G., Wieland, V., 2000. A small estimated Euro-area model with rational expectations and nominal rigidities. ECB Working paper No 30.
- Cooley, T.F., Prescott, E., 1995. Economic growth and business cycles. In: Cooley, T.F. (Ed.), *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, NJ.
- Erceg, C., Henderson, D., Levin, A., 2000. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46, 281–313.
- European Economy, 1999. European Commission, Brussels 1999.
- Fagan, G., Henry, J., Mestre, R., 2001. An area-wide model (AWM) for the Euro area. ECB Working paper series No. 42, January.
- Fuhrer, J., Moore, G., 1995. Inflation persistence. *Quarterly Journal of Economics* 110 (1), 127–159.
- Fuhrer, J., Moore, G., Schuh, S., 1995. Estimating the linear-quadratic inventory: Maximum likelihood versus generalized method of moments. *Journal of Monetary Economics* 35, 115–157.
- Galí, J., Gertler, M., 1999. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44, 195–222.
- Goodfriend, M., King, R., 1997. The new neoclassical synthesis. In: Bernanke, B., Rotemberg, J. (Eds.), *NBER MacroAnnual*. MIT Press, Cambridge, MA.
- Gordon, R.J., 1998. Foundations of the Goldilocks economy. *Brookings Papers on Economic Activity* 2, 297–346.
- Mankiw, N.G., 2000. The inexorable and mysterious tradeoff between inflation and unemployment. NBER Working paper # 7884, September.
- Roberts, J., 1997. Is inflation sticky. *Journal of Monetary Economics* 39 (2), 173–196.
- Rotemberg, J., Woodford, M., 1995. Dynamic general equilibrium models with imperfectly competitive product markets. In: Cooley, T.F. (Ed.), *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, NJ, pp. 243–293.
- Rotemberg, J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Bernanke, B., Rotemberg, J. (Eds.), *NBER Macroeconomic Annual 1997*. MIT Press, Cambridge, MA, pp. 297–346.
- Rudebusch, G.D., Svensson, L.E.O., 1999. Policy rules for inflation targeting. In: Taylor, J. (Ed.), *Monetary Policy Rules*. The University of Chicago Press, Chicago, pp. 203–246.
- Sbordone, A.M., 1999. Prices and unit labor costs: A new test of price stickiness. Mimeo., Rutgers University.
- Shimer, R., 1998. Why is the U.S. unemployment rate so much lower. In: Bernanke, B., Rotemberg, J. (Eds.), *NBER Macro Annual*. MIT Press, Cambridge, MA.
- Staiger, D., Stock, J., 1997. Instrumental variables regression with weak instruments. *Econometrica* 65 (3), 557–586.
- Stock, J., 1998. Comment on Gordon. *Brookings Papers on Economic Activity*, 2, 347–360.
- Stock, J., Watson, M., 1999. Forecasting inflation. *Journal of Monetary Economics* 44, 293–335.
- Taylor, J.B., 1980. Aggregate dynamics and staggered contracts. *Journal of Political Economy* 88, 1–23.
- Woodford, M., 1996. Control of the public debt: A requirement for price stability? NBER Working paper No. 5684, July.