

Why is the New Keynesian Phillips Curve apparently flat?*

Keith Kuester[†]
European Central Bank

Gernot J. Müller[‡]
Goethe University Frankfurt

Sarah Stoelting[§]
European University Institute

February 26, 2007

Abstract

Macroeconomic data suggest that the New Keynesian Phillips curve is quite flat - despite microeconomic evidence implying frequent price adjustments. A widely discussed hypothesis is the presence of real rigidities. This paper proposes an alternative explanation: if price markup/markup shocks are persistent and negatively correlated with the labor share, a widely used measure for marginal costs, the pass-through of measured marginal costs into inflation is limited, even if prices are fairly flexible. Using a standard New Keynesian model, we show that the GMM approach to the New Keynesian Phillips curve leads to inconsistent and upward biased estimates if cost-push shocks indeed are persistent. Monte Carlo experiments suggest that the bias is quite sizeable: we find average price durations estimated as high as 11 quarters, when the true value is about 2 quarters. Standard diagnostic tests fail to signal a misspecification of the model.

Keywords: Price Rigidities, New Keynesian Phillips Curve, Cost-push shocks, GMM estimation
JEL-Codes: E30, C15

*The views expressed in the paper are those of the authors and do not necessarily reflect those of the European Central Bank.

[†]Mailing address: European Central Bank, Monetary Policy Strategy Division, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany; E-mail: keith.kuester@ecb.int

[‡]Mailing address: Mertonstrasse 17, 60054 Frankfurt am Main, Germany; E-mail: gernot.mueller@wiwi.uni-frankfurt.de

[§]Mailing address: Villa San Paolo, Via della Piazzuola 43, 50133 Florence, Italy; E-mail: sarah.stolting@eui.eu

1 Introduction

The New Keynesian Phillips curve (NKPC) is a widely used structural model of inflation dynamics. Its key parameter, which governs the pass-through of marginal costs into inflation, is the average time over which prices are kept fixed. This average price duration provides a measure for the degree of price stickiness. If prices are kept unchanged for some time, i.e. if they are sticky, then the pass-through of marginal costs into inflation is limited, or, in other words, the Phillips curve is quite flat. Using Generalized Method of Moments (GMM) estimation techniques on U.S. time series Galí and Gertler (1999), hereafter GG, find average price durations of about 4 quarters. These estimates seem high, or, put differently, the NKPC appears to be quite flat, given recent microeconomic evidence put forward by Bils and Klenow (2004) suggesting average price durations of less than 5 months.¹

Substantial research efforts have been made to reconcile the apparent conflict between micro and macroeconomic evidence on price stickiness. By now it is widely recognized that various kinds of real rigidities allow a reinterpretation of the macroeconomic evidence consistent with much shorter average price durations, see, for instance, Eichenbaum and Fisher (2004) and Altig, Christiano, Eichenbaum, and Lindé (2005). Intuitively, real rigidities induce strategic complementarities in price setting such that price setters change prices by smaller amounts whenever they are able to adjust prices. As a consequence, the pass-through of marginal costs into inflation remains limited, although firms update their price frequently. The appeal of this route to reconcile micro and macroeconomic evidence remains limited, however, to the extent that microevidence not only suggests frequent, but also large price adjustments, a point stressed, for instance, by Maćkowiak and Wiederholt (2006).

The present paper, therefore, suggests a new perspective on the conflicting evidence on price stickiness. We start from the observation that recent estimates of fully specified New Keynesian general equilibrium models on the basis of full information estimation techniques find evidence for autocorrelated cost-push shocks, also labelled ‘price markup shocks’ in the literature, e.g. Smets and Wouters (2006) and Galí and Rabanal (2005). In the present paper, we assess the implications of these shocks for the single equation, GMM-based approach to the NKPC. We find, first, that the orthogonality conditions imposed in the estimation are generally invalid and, second, that because cost-push shocks are correlated with observable marginal costs, GMM estimates are inconsistent and upward biased.

Clearly, while omission of actual features of the economy will induce biases in estimation, it remains an ultimately empirical question, how severe these omissions and how sizeable these biases are. To answer these questions in the context of the NKPC, we conduct Monte Carlo experiments using as

¹Galí, Gertler, and López-Salido (2001, 2003, 2005) provide additional evidence. Depending on the particular specification and sample, the average price duration is found to be between 2.4 and 11.8 quarters for U.S. data. Note that average price durations of about 4 quarters seemed to square well with earlier survey evidence, see Blinder, Canetti, Lebow, and Rudd (1998). Also, a more recent assessment of microeconomic data for the U.S. by Nakamura and Steinsson (2006) suggests price durations between 8 and 11 month once sales are excluded from the sample.

data generating process an empirically plausible model. Specifically, we use the small New Keynesian DSGE model estimated by Galí and Rabanal (2005). This model stays close in flavor to the New Keynesian variants discussed in Woodford (2003) and is by and large representative for a sizeable literature of small and medium scale estimated DSGE models; see e.g. Smets and Wouters (2006). Using the simulated data, we find that GMM estimates are precise and consistent as long as the model is correctly specified. However, once we introduce autocorrelated cost-push shocks into the model, the performance of the GMM estimator deteriorates. The procedure signals a precise estimate even though there is substantial bias. In particular, we find that the bias increases drastically as the persistence of cost-push shocks increases. In fact, for a degree of autocorrelation of cost-push shocks of 0.95, a value reported in Galí and Rabanal (2005), the point estimate for the degree of price stickiness implies a mean duration of prices of about 11 quarters although the true value is just over 2 quarters.

We then turn to diagnostic tests asking whether the econometrician would be able to detect the misspecification of the model. We find that this is unlikely given a realistic sample size of about 150 observations. While the orthogonality conditions imposed in the estimation are, in fact, violated by the model, the value of the J-statistic suggests a 66 percent probability of the null of no violation. Regarding the error terms, there is some evidence for autocorrelation. Yet, it hardly exceeds the extent of autocorrelation in the residuals suggested by tests on U.S. data, say, by Galí et al. (2001). A more sophisticated test, suggested by GG, gives a mixed picture.

As a final test, we resort to Maximum Likelihood (ML) estimation of the misspecified model. Here, just as with GMM, there is again a substantial upward bias in the estimated degree of price rigidity. In the same vein, casual observations of recent full information estimation results of the NKPC also suggest lower degrees of price rigidities, in case autocorrelated cost-push shocks are allowed for, see Galí and Rabanal (2005) and Smets and Wouters (2006).

The paper is structured as follows. Section 2 outlines the model economy used as the data generating process. Section 3 shows that the estimated degree of price rigidity is likely to be upward biased in the presence of autocorrelated cost-push shocks. Section 4 provides a quantitative assessment of this bias using Monte Carlo experiments. In section 5 we provide a brief discussion of alternative approaches, while section 6 concludes.

2 The Model Economy

In this section we outline a variant of the baseline New Keynesian general equilibrium model. Specifically, we focus on the model estimated by Galí and Rabanal (2005) on post World War II U.S. data. Our exposition of the model economy is brief; for more details on the model and the estimation, we refer the interested reader directly to Galí and Rabanal (2005).

The demand side of the model is represented by a consumption Euler equation:

$$b\Delta y_t = E\{\Delta y_{t+1}\} - (1-b)(r_t - E_t\{\pi_{t+1}\}) + (1-\rho_g)(1-b)g_t. \quad (1)$$

Above, Δy_t denotes output growth, r_t the nominal interest rate, π_t the inflation rate. g_t is a demand shock. Parameters ρ_g and b reflect the correlation of this shock and external consumption habits, respectively.

Firms produce differentiated goods which they sell in monopolistically competitive product markets. Their output is linear in employment. From the production side, up to first order, aggregate output, \tilde{y}_t , is therefore linear in aggregate employment, n_t :

$$\tilde{y}_t = n_t. \quad (2)$$

Productivity shocks in the economy are permanent shocks. A tilde on top of a variable indicates that the respective variable has been normalized by productivity before linearization in order to render the linearized model economy stationary. \tilde{y}_t is linked to output growth, Δy_t , and to innovations in permanent technology, ϵ_t^a , by $\Delta y_t = \Delta \tilde{y}_t + \epsilon_t^a$.

Producers of differentiated goods hire bundles of labor in a perfectly competitive market. In each period, which is assumed to be one quarter, each producer of a differentiated good will not be able to reoptimize its price with a certain probability. Instead producers who do not reoptimize mechanistically index to lagged inflation. Linearizing around a zero inflation steady state, the New Keynesian Phillips curve in this model is given by

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t\{\pi_{t+1}\} + \kappa_p (l_t + u_t). \quad (3)$$

Above $\gamma_b = \frac{\eta_p}{1+\beta\eta_p}$, $\gamma_f = \frac{\beta}{1+\beta\eta_p}$ and $\kappa_p = \frac{(1-\beta\theta_p)(1-\theta_p)}{\theta_p(1+\eta_p\beta)}$. θ_p is the probability that a firm cannot reoptimize its price in a given period. Below we will be concerned with the estimation of this parameter which, by the law of large numbers, can also be interpreted as the fraction of firms which keep their prices unchanged in a given period. The average price duration, D , is given by $1/(1-\theta_p)$.

Price indexation is captured by parameter η_p . Parameter β denotes the time discount factor. According to the model, inflation is driven by current and expected marginal costs. These can be decomposed into a measure of the labor share, $l_t = \tilde{w}_t$, where \tilde{w}_t denotes the real wage per efficiency unit, and a shock to the markup/a cost-push shock, u_t . It is the autocorrelation of this shock, ρ_u , which crucially influences estimates of the degree of price stickiness in the economy as we will demonstrate below.

Workers supply their labor monopolistically competitive to intermediaries who bundle these labor services and sell them on to goods producers under perfect competition. Like goods prices, also individual wages are subject to a Calvo nominal rigidity. Those workers who do not update their wage in a given period instead partially index their nominal wage to past inflation. The wage equation in

Table 1: PARAMETER VALUES USED IN SIMULATION

Consumers		
Time-discount factor	β	= 0.99
External habit parameter	b	= 0.42
Inverse of labor supply elasticity	φ	= 0.80
Calvo-stickiness wages	θ_w	= 0.05
Indexation wages	η_w	= 0.42
Price elasticity of demand	ϵ_p	= 6.00
Producers		
Calvo-stickiness prices	θ_p	= 0.53
Indexation prices	η_p	= 0.00
Wage elasticity of labor demand	ϵ_w	= 6.00
Monetary policy		
Response to lagged interest rate	ϕ_r	= 0.69
Response to inflation	ϕ_π	= 1.35
Response to output growth	$\phi_{\Delta y}$	= 0.26
Autocorrelation of shocks		
Persistence of demand shock	ρ_g	= 0.93
Persistence of productivity	ρ_a	= 1.00
Persistence of price-markup shock	ρ_u	= 0.95
Persistence of wage-markup shock	ρ_ν	= 0.91
Standard deviation of innovations		
Innovation to productivity	σ^a	= 0.009
Innovation to demand shock	σ^g	= 0.025
Innovation to price-markup	σ^u	= 0.011
Innovation to wage-markup	σ^w	= 0.012
Innovation to monetary policy	σ^m	= 0.003

Notes: Parameters used to generate data. Parameters are taken from the mean estimates of Galí and Rabanal (2005). Their mean estimate of price indexation is $\eta_p = 0.02$. We set this parameter to zero in order to enhance expositional clarity.

the model of Galí and Rabanal (2005) reads as

$$\begin{aligned} \tilde{w}_t = & \frac{1}{1+\beta}\tilde{w}_{t-1} + \frac{\beta}{1+\beta}E_t\{\tilde{w}_{t+1}\} - \frac{1}{1+\beta}\epsilon_t^a + \frac{\eta_w}{1+\beta}\pi_{t-1} - \frac{1+\beta\eta_w}{1+\beta}\pi_t \\ & + \frac{\beta}{1+\beta}E_t\{\pi_{t+1}\} - \frac{\kappa_w}{1+\beta}(\mu_t^w - \nu_t). \end{aligned} \quad (4)$$

Wages are driven by endogenous variations in the wage markup, μ_t^w , and by exogenous shocks to the markup, ν_t . Parameter $\kappa_w = \frac{(1-\theta_w)(1-\beta\theta_w)}{\theta_w(1+\epsilon_w\varphi)}$ multiplying the markup term measures the strength with which markups influence wages. θ_w is the probability that a worker cannot change its wage, ϵ_w is the elasticity of labor demand with respect to wages and $1/\varphi$ measures the Frisch elasticity of labor supply. The markup links to the rest of the economy via $\mu_t^w = \tilde{w}_t - \left(\frac{1}{1-b}\tilde{y}_t - \frac{b}{1-b}\tilde{y}_{t-1} - g_t + \frac{b}{1-b}\epsilon_t^a + \varphi n_t\right)$. Like price markup shocks, also wage markup shocks are serially correlated, as captured by the autocorrelation parameter, ρ_ν . The economy is closed by a Taylor type rule for monetary policy. The authority sets interest rates in reaction to inflation and output growth. On top of this monetary policy rates are also subject to a monetary policy shock ϵ_t^m :

$$r_t = \phi_r r_{t-1} + (1 - \phi_r)\phi_\pi \pi_t + (1 - \phi_r)\phi_y \Delta y_t + \epsilon_t^m. \quad (5)$$

Table 1 shows the parameterization we choose for the model economy on the basis of the mean estimates reported by Galí and Rabanal (2005). We only deviate with respect to the indexation of prices where their mean estimate is as low as $\eta_p = 0.02$. Following large parts of the GMM literature on the NKPC, we set this parameter to zero.

3 Cost-push shocks and the econometrics of the NKPC

Having outlined a fully specified structural model, we now turn to the econometrics of the NKPC, given by equation (3). This equation takes center stage within the single equation approach based on the Generalized Methods of Moments (GMM). We establish four results. Intuitively, we expect that the estimated pass-through of observed marginal costs (i.e. the labor share), l_t , into inflation, π_t , is lowered by cost-push shocks if they are negatively correlated with the labor share. The pass-through is captured by the estimate of the slope of the NKPC, $\hat{\kappa}_p$. Indeed, this is our first results: whenever cost-push shocks are serially correlated, as much of the recent Bayesian estimation literature finds, GMM estimates of price-duration are inconsistent. In the model at hand we compute, second, that the bias runs in the direction of too much estimated price rigidity. We illustrate third that Q-tests which are typically used to assess the serial correlation of the error terms cannot be used to discern the serial correlation properties of cost-push shocks. Finally, we show that serial correlation of cost-push shocks renders the orthogonality conditions exploited in the GMM estimation invalid.

For the sake of clarity, in the following we restrict ourselves to the estimation of the slope of the Phillips curve, κ_p , and restrict the other parameter in the NKPC, β , to its true value.²

²Appendix A presents the case where also β is estimated. This does not have a strong bearing on the bias in the estimated

Consistency To be slightly more formal note that, absent indexation, NKPC equation (3) can be written as follows

$$y_{t+1} = \kappa_p l_t + \epsilon_{t+1}^{RE} + \tilde{u}_t, \quad (6)$$

where $y_{t+1} := \pi_t - \beta\pi_{t+1}$.

Here ϵ_{t+1}^{RE} is the rational expectations error that ensures $E_t(\pi_{t+1}) \equiv \pi_{t+1} + \frac{1}{\beta}\epsilon_{t+1}^{RE}$ and $\tilde{u}_t = \kappa_p u_t$. Equation (6) is linear. Estimating it by GMM is thus equivalent to applying two stage least squares estimation (2SLS henceforth).³ Let the instrument vector be $z_{t-1} = [\pi_{t-1}, l_{t-1}]'$. As Appendix A shows, in probability the estimator converges to

$$\hat{\kappa}_{p,T}^{2SLS} \xrightarrow{p} \kappa_p + q^{-1} \bar{\delta}' \kappa_p \rho_u E \left\{ \left[\pi_{t-1}, l_{t-1} \right]' u_{t-1} \right\}. \quad (7)$$

Here $\bar{\delta}$ is the probability limit of $\hat{\delta}_T$, the estimator in the first-stage regression, and q is the second cross-moment of regressor and instrumented regressor. Most notably, to the extent that the instruments correlate with the cost-push shock contemporaneously, as is the case in our data-generating model, i.e. $E \left\{ \left[\pi_{t-1}, l_{t-1} \right]' u_{t-1} \right\} \neq 0$, the estimates will be inconsistent whenever the cost-push shock is persistent ($\rho_u \neq 0$).




Figure 1: THEORETICAL PROBABILITY LIMITS OF 2SLS ESTIMATORS FOR VARYING DEGREES OF SERIAL CORRELATION, ρ_u , IN THE COST-PUSH SHOCK. Dashed lines show the true values. From left to right: estimator for the slope of the Phillips curve, κ_p , and for the fraction of firms which do not reoptimize their price, θ_p .

slope of the NKPC. Also in this case the implied price durations show an upward bias. The strength of this bias, however, is mitigated by an upward bias also in the estimate of β .

³Rudd and Whelan (2005) also exploit this property to analyze a possible bias in GMM estimates of the NKPC. Their focus, however, is different from ours as they focus on the possibility that the true process for inflation is a backward-looking Phillips curve.

Assessing the direction of the bias Under the calibration of the model, computing $q, \bar{\delta}$ and the covariance between instruments and the cost-push term, we find that whenever cost-push shocks are persistent, estimates of the slope of the NKPC, κ_p , are inconsistent and converge to too low values. The left panel of Figure 1 plots the probability limit of $\hat{\kappa}_p$ as a function of ρ_u . The downward bias resulting from an increase in ρ_u is apparent. It translates into too large estimates of the mean price duration even in the limit, shown in the right panel of Figure 1.

The economics behind this finding is the following. Consider for the sake of the argument a higher than average realization of the labor share. Following the NKPC, ceteris paribus a high value of the labor share in period t has an inflationary impact. And, all else equal again, a negative realization of the cost-push shock induces a rise in the labor share. To the extent that cost-push shocks are frequent and a sizeable source of fluctuations in the labor share, a higher than average labor share will frequently come along with a low realization of the cost-push/markup shock. This will be persistently the case if cost-push shocks are serially correlated. On these occasions a high labor share is not as inflationary as in the absence of these correlated cost-push shocks. In order to reconcile the frequent absence of an inflationary impact from persistently high labor shares with the NKPC, the pass-through parameter κ_p will be therefore be underestimated.

Implications of cost-push shocks for the Q-Test In the following we argue that the Q-Test for serial correlation of the error term in any estimated NKPC will rightfully point to serial correlation in the estimation residuals whenever cost-push shocks are present ($\sigma_u > 0$), i.e. even if $\rho_u = 0$. The Q-Test statistic is given by

$$Q(p) = T(T+2) \sum_{j=1}^p \frac{\hat{r}_{j,T}^2}{T-j}, \quad (8)$$

where $\hat{r}_{j,T}$ is the j^{th} correlation coefficient of the residuals \hat{e}_t with

$$\hat{r}_{j,T} = \frac{\sum_{t=j+1}^T \hat{e}_t \hat{e}_{t-j}}{\sum_{t=1}^T (\hat{e}_t)^2}.$$

Under the null of no serial correlation in the residual e_t , invoking a suitable central limit theorem for martingale difference sequences, the Q-statistic asymptotically follows a χ^2 distribution with p degrees of freedom:

$$Q(p) \xrightarrow{d} \chi^2(p).$$

For the sake of clarity, let us focus on the case $p = 1$. The asymptotics of the Q-Test relies on

$$\frac{1}{\sqrt{T}} \sum_{t=2}^T \hat{e}_t \hat{e}_{t-1} \xrightarrow{d} N(0, \sigma_e^2).$$

Evaluated at the true parameter values, the combined residual in the New Keynesian Phillips curve is given by $e_t := \epsilon_t^{\text{RE}} + \tilde{u}_{t-1}$. The point is that $\epsilon_{t-1}^{\text{RE}}$ in the model is not orthogonal to \tilde{u}_{t-1} . e_t is

therefore serially correlated whenever cost-push shocks are present. More concretely, to the extent that cost-push shocks are inflationary, they tend to induce positive inflation surprises such that in the New Keynesian model,

$$E \{ \tilde{u}_{t-1} \epsilon_{t-1}^{RE} \} < 0.$$

In other words, even when cost-push shocks are white noise, the residual e_t will not be a martingale difference sequence. The Q statistic in this case will tend to be non-central chi-square with a strictly positive-non-centrality parameter. It will therefore consistently signal autocorrelation in the empirical residual of the NKPC even in the absence of autocorrelation in the cost-push shock itself.

Evidence obtained on the basis of the Q-Test signalling serial correlation in the NKPC residual is therefore as much a test of the mere presence of shocks to the markup/cost-push shocks as a test of serial correlation these shocks. The Q-Test cannot be used to discern whether cost-push shocks/shocks to the markup are serially correlated.

Implications of serial correlation for the J-Test Given that the GMM estimation is based on orthogonality conditions one may ask whether the autocorrelation of cost-push shocks also has a bearing on these. Assuming cost-push shocks are unobserved, the GMM approach to the NKPC exploits the following set of moment conditions:

$$\begin{aligned} E \{ [y_{t+1} - \kappa_p l_t] \mathbf{z}_{t-1} \} &= E \{ [\epsilon_{t+1}^{RE} + \tilde{u}_t] \mathbf{z}_{t-1} \} \\ &= \rho_u E \{ \tilde{u}_{t-1} \mathbf{z}_{t-1} \} \neq 0, \quad \text{if } \rho_u > 0. \end{aligned} \quad (9)$$

In words, the orthogonality conditions are unlikely to hold in the baseline New Keynesian model with autocorrelated cost-push shocks. In principle, the J-test used in GMM estimation should signal the violation of these conditions; in practical applications, however, lack of power leads to a frequent failure to reject the null as we shall demonstrate in the next section.

4 Monte Carlo Experiments

In order to perform a quantitative assessment of the bias due to autocorrelated cost-push shocks we perform a Monte Carlo experiment. In a first step, we consider the standard GMM approach to the NKPC. In a second step we also look at the effects of an incorrectly specified full information approach. Both experiments suggest that the bias can be substantial. Standard diagnostic tests, however, are not necessarily signalling a misspecification of the model.

4.1 GMM estimation

We simulate the model outlined in section 2 and generate 1000 random time series using 100 observations to initiate the model. Assuming that the value of β is known to be 0.99, the value used

in the simulation of the model, we estimate the parameter θ_p on the basis of the following moment condition,⁴

$$E \{ [\theta_p \pi_t - (1 - \theta_p)(1 - \beta \theta_p) l_t - \theta_p \beta \pi_{t+1}] z_{t-1} \} = 0. \quad (10)$$

Now the instrument vector z_{t-1} contains four lags of inflation, the labor income share and output growth, i.e. variables that are close to the instruments used by GG and that were at the same time available from our simulation of the model.⁵ Our optimal weighting matrix uses the Newey-West correction for the likely serial correlation of the orthogonality conditions. We repeat the experiment for different degrees of persistence of cost-push shocks. We then compare the distribution of point estimates, $\hat{\theta}_p$, to the true value, which is $\theta_p = 0.53$ throughout.

Table 2: RESULTS OF GMM ESTIMATION

	Point Estimates		Diagnostics		
	$\hat{\theta}_p$	D	F-Statistic	J-Statistic	Q(4)
<i>152 Observations:</i>					
$\rho_u = 0$	0.51 (0.06)	2.08 (0.26)	5.06 (0.00)	9.26 (0.61)	10.12 (0.21)
$\rho_u = 0.95$	0.91 (0.04)	16.28 (14.09)	123.36 (0.00)	8.46 (0.66)	14.54 (0.03)
<i>2,000 Observations:</i>					
$\rho_u = 0$	0.53 (0.01)	2.12 (0.06)	59.90 (0.00)	10.56 (0.52)	101.45 (0.00)
$\rho_u = 0.95$	0.95 (0.02)	26.65 (27.57)	2150.11 (0.00)	28.14 (0.06)	159.63 (0.00)

Notes: values (in parentheses) are averages (and standard deviations) over 1000 draws; except for the diagnostics where average p-values are given in parentheses.

The results of our experiments are displayed in Table 2. In the upper panel, we consider a sample size which is typical for macroeconomic studies: 152 usable observations as in Galí and Gertler (1999). The first row gives the estimated degree of price stickiness and the diagnostics if price markup shocks are uncorrelated in the true model economy, while the second row of the first panel shows results obtained on the basis of a higher degree of autocorrelation: $\rho_u = 0.95$, the value reported by Galí and Rabanal (2005). In the first case, the mean estimate of $\hat{\theta}_p = 0.51$ is close to the true value of

⁴Using the alternative moment condition, i.e. dividing equation (10) by θ_p , we find that the estimator for θ_p converges relatively often to 1. We therefore report results obtained on the basis 10.

⁵GG use in addition four lags of a short-long interest rate spread, wage inflation and commodity price inflation.

$\theta_p = 0.53$. The implied mean duration of prices, computed as the average over the mean duration implied by each $\hat{\theta}_p$, is displayed in the second column labeled ‘D’. It is about 2 quarters. The small downward bias disappears if a larger sample is considered. In the lower panel of Table 2 we repeat the experiment but use 2000 observations in each draw. As a result, we obtain a point estimate of $\hat{\theta}_p = 0.53$ suggesting that the degree of price rigidity can be consistently estimated in the absence of autocorrelation of markup shocks.⁶

In contrast, if cost push shocks are autocorrelated, we obtain quite different estimates for the degree of price stickiness. In the small sample, we find $\hat{\theta}_p = 0.91$, which, in turn, implies average price durations of about 11 quarters.⁷ We are thus confronted with a substantial upward bias and - by any measure - an enormous degree of price rigidity. It is noteworthy, though, that such a high degree of price rigidity is also found by Galí and Gertler (1999) for some specifications of the NKPC model estimated on U.S. time series. In our setup, the estimator is also inconsistent in this case. If 2000 observations are used instead of 152, we still find mean estimates of $\hat{\theta}_p = 0.95$.




Figure 2: DISTRIBUTION OF $\hat{\theta}$, GIVEN 1000 DRAWS (152 OBSERVATIONS EACH) AND $\theta = 0.53$

To summarize, Figure 2 gives a graphical representation of the results. The left panel shows the distribution of $\hat{\theta}_p$ obtained if cost-push shocks are not persistent. These are relatively well centered around the true value of $\theta_p = 0.53$. The right panel makes clear the bias in the estimates in the presence of serial correlation of the markup shock ($\rho_u = 0.95$).

Given that the empirical model is misspecified under autocorrelated markup shocks, this result and in particular the direction of the bias may not come as a surprise given the arguments put forward in

⁶Lindé (2005) investigates the properties of GMM in estimating the NKPC focusing on the extent of forward-looking behavior in price setting. He finds that even an uncorrelated cost-push shock induces a some bias. In light of our results and the arguments put forward in Galí et al. (2005) this is likely to be the result of his choice of instruments.

⁷Note in the second column of Table 2 we report $D = 16.28$. Recall that we compute the average over the mean duration implied by each estimate for θ_p . This is different from the mean duration implied by the average estimate of θ_p , because of Jensen’s inequality, see Baharad and Eden (2004) for further exploration of this issue.

section 3. The key question is whether standard diagnostic tests are able to detect the violation of the moment condition imposed in the GMM estimation in case of autocorrelated markup shocks. Before turning to this question, we note that in all specifications, we reject the null of no significance for the instruments in a first state regression of the labor share on the set of instruments. The results of the F-test are displayed in the third column of Table 2.

In the fourth column we report the J-Statistic. For the small sample, the J-test clearly fails to detect the violations of the orthogonality conditions which were imposed in the GMM estimation: the p-value in favor of the null of orthogonality is above 60 percent. Yet this test is known to have low power. Turning to the the hypothetical sample size of 2,000 observations shows that it works in principle. But still the null of orthogonality can only be rejected at the 10 percent level in case of $\rho_u = 0.95$.

We then turn to the Box-Ljung test for autocorrelation of the residuals. There is no evidence for autocorrelation in the small sample if cost-push shocks are not persistent ($\rho_u = 0$). In contrast, in case markup shocks are autocorrelated, the p-value of 3 percent suggests a rejection of the null of no autocorrelation. It is quite likely, however, that such a result would not necessarily lead an econometrician to reject the null and hence the model. For instance, Galí et al. (2001) report a $Q(4) = 10.2$ with a p-value of 4 percent, when estimating the NKPC on U.S. data. Turning to the large sample, we observe that the null is also rejected due to the mere presence of cost-push shocks even if they are not autocorrelated, as argued in section 3 above.⁸

As a final test of the model, we assess the performance of the NKPC by considering model based forecasts for actual inflation, following GG and Galí et al. (2001). Solving the NKPC forward gives

$$\pi_t = \kappa_p \sum_{k=0}^{\infty} \beta^k E_t \{l_{t+k} + u_{t+k}\} = \kappa_p \sum_{k=0}^{\infty} \beta^k E_t \{l_{t+k}\} + \frac{\kappa_p}{1 - \beta\rho_u} u_t. \quad (11)$$

Forecasts for the future labor share are obtained by using a VAR model. Specifically, we consider a 3-variable VAR(4) model in v_t , which includes the labor share, output growth and inflation. Further, let A denote the companion matrix of the VAR(4) model and $x_t = [v_t' \ v_{t-1}' \ v_{t-2}' \ v_{t-3}']'$. Let e_1 be a vector with zeros except for a 1 in the position of the labor share. Then $l_{t+k} = e_1' A^k x_t$ is the VAR based forecast for the labor share which is taken to proxy for the rational expectations of the labor share in the NKPC. Substituting this for rational expectations in (11), we have for the prediction of current inflation

$$\pi_t = \kappa_p e_1' (I - \beta A)^{-1} x_t + \frac{\kappa_p}{1 - \beta\rho_u} u_t.$$

Assuming that the markup shock, u_t , is not observed, the following expression gives the model based forecast of inflation, referred to as fundamental inflation by GG:

$$\pi_t^* = \kappa_p e_1' (I - \beta A)^{-1} x_t. \quad (12)$$

⁸If one assumes that there are no cost-push shocks at all ($\sigma_u=0$), a case not reported in Table 2, one cannot reject the null of no autocorrelation in the residuals for large samples. Specifically, we find an average p-value of 0.43

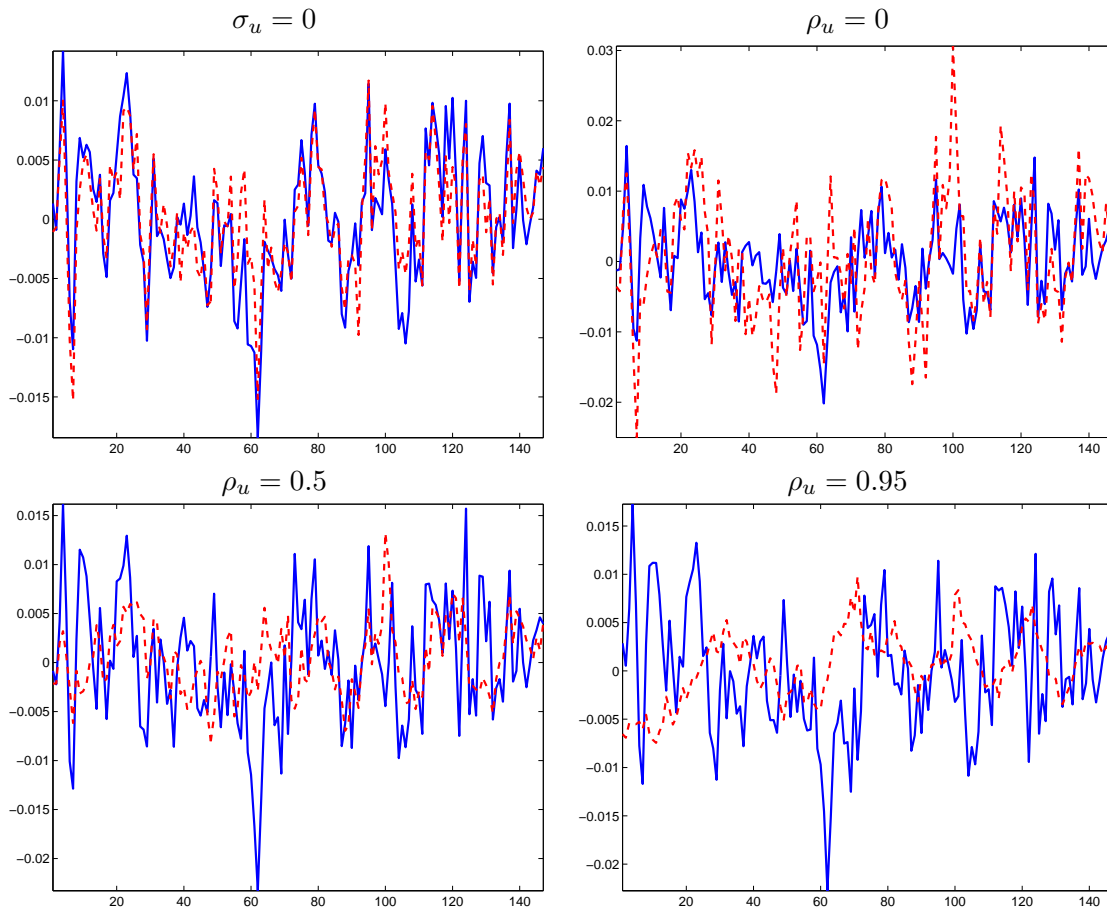


Figure 3: ACTUAL (SOLID LINE) VS. FUNDAMENTAL (DASHED LINE) INFLATION

Following GG and Galí et al. (2001) we assess the model performance by comparing actual and fundamental inflation. We consider four cases: in each case we draw one set of time series from the model. In the first case, we switch off cost-push shocks altogether. Then we consider, in turn, the case of uncorrelated markup shocks, a degree of autocorrelation of 0.5 and, finally, a degree of autocorrelation of 0.95. Figure 3 displays the results.

In the upper left panel (case 1), fundamental inflation tracks actual inflation quite closely. The estimate obtained in this case is $\hat{\theta}_p = 0.53$. In the upper right panel (case 2), with uncorrelated cost push shocks fundamental inflation somewhat overpredicts the movements in actual inflation ($\hat{\theta}_p = 0.49$) as a result of not taking into account cost push shocks. In the lower panels, cases 3 and 4 are displayed: with an intermediate degree of autocorrelation in the cost push shock, the model already has some problems to predict inflation correctly. There is, however, already substantial bias in the estimated degree of price rigidity as $\hat{\theta}_p = 0.69$. Finally, the model performance deteriorates further in case of highly correlated cost push shocks, i.e. the case in which the bias is largest ($\hat{\theta}_p = 0.88$).

Assuming autocorrelated cost-push shocks in the actual data generating process, it is not clear whether




Figure 4: ESTIMATES FOR θ_p , FOR INCREASING VALUES OF ρ_u (AVERAGE OF 1000 DRAWS, 152 OBSERVATIONS EACH)

the evidence depicted in Figure 3, would be sufficient to reject the NKPC model estimated using GMM. Casual inspection suggests that the model does a better job if confronted with actual data in GG and Galí et al. (2001). It should also be noted, however, that the assessment in Figure 3 is based on single draws and one cannot rule out that the model would do better or even worse for another realization of time series.

In sum, the presence of an autocorrelated cost push shock induces a substantial upward bias in the GMM estimate of the degree of price rigidity. At the same time it is unclear whether such a bias would be detectable in actual data, given the ambivalence of the signals obtained from standard diagnostic tests. Finally, to assess the role of persistence for the GMM estimates of θ_p more formally, Figure 4 plots the mean estimate of θ_p against the degree of persistence in the markup shock assumed in the simulation of the data in the model: the bias increases systematically with the degree of autocorrelation of the cost-push shock.

4.2 Maximum Likelihood Estimation

This section evaluates to which extend the omission of an autocorrelated cost-push term may induce a bias towards too much price stickiness not only in the GMM but also in the maximum likelihood estimates of the model. At the same time, we investigate whether this bias would be identifiable in at least a subset of the estimation statistics. As above, we generate 1000 random time series using again the model outlined in section 2, where $\rho_u = 0.95$. We focus on a realistic sample size of 152 observations and use a further 100 observations are used to initiate the model.

152 Observations, $\rho_u = 0$

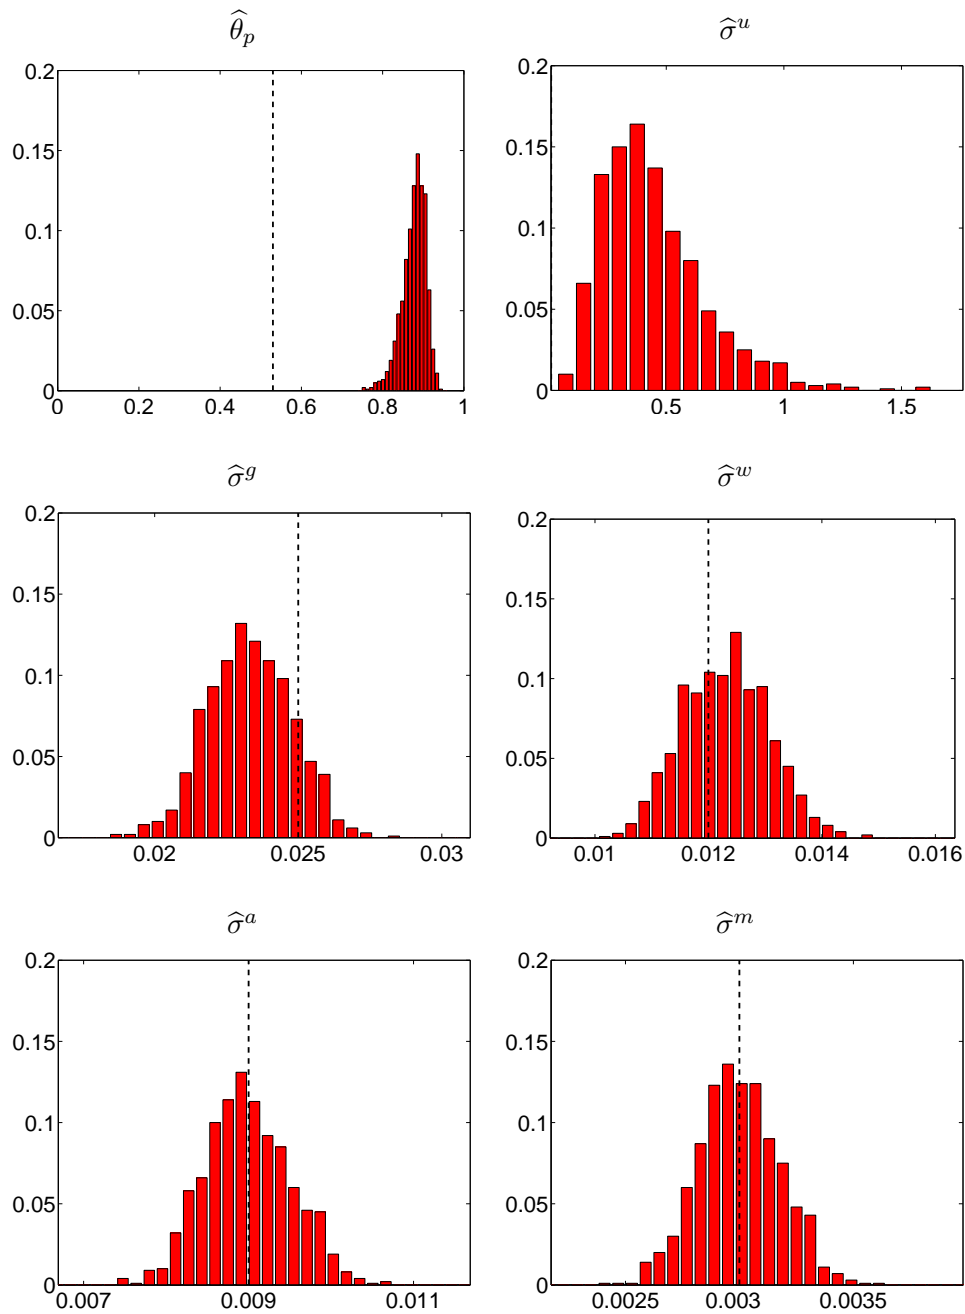


Figure 5: DISTRIBUTION OF ESTIMATED PARAMETERS GIVEN 1000 DRAWS AND A SAMPLE LENGTH OF 152 QUARTERS IN THE MIS- SPECIFIED MODEL.

Notes: The red columns pertain to histograms of the estimated parameter values. In each of the plots, a vertical black dashed line marks the true parameter value which had been used to generate the data. From top left to bottom right, the parameters are: price stickiness θ_p , standard deviations of the price markup shock, of the shocks to government spending, to the wage markup and to productivity. The final panel shows the estimated standard deviation of the monetary policy shock. For the estimate of the standard deviation of innovations to the cost-push shock, $\hat{\sigma}_u$, the dashed line signalling the true value is adjacent to the left vertical axis.

In the first experiment, we assume that estimation is conducted with the wrong model, i.e. the model without allowing for serial correlation in the price markup shock (false perception). In the estimation we give the econometrician more than a fair chance by assuming that she has full knowledge about all the parameters of the model apart from the degree of price stickiness, θ_p , and the stochastic structure, i.e. the standard deviations of the shocks. Apart from the serial correlation of the price markup shock, the corresponding serial correlation parameters of all the shocks are assumed to be known and are set to their true values. For each draw, θ_p and the standard deviations of the innovations to the shock processes are estimated.

Figure 5 plots the results by showing the histograms of the distribution of the estimators over the different draws when estimation is conducted with the wrong model which does not feature serial correlation in the price markup shock. As is evident from the top-left panel, a considerable bias in the estimation of the Phillips curve results from omitting serially correlated cost-push shocks from the estimation. The first row of Table 3 summarizes the above in numbers.

Table 3: PARAMETER ESTIMATES WITH MAXIMUM LIKELIHOOD

	$\hat{\theta}_p$	$\hat{\sigma}^a \cdot 100$	$\hat{\sigma}^g \cdot 100$	$\hat{\sigma}^u \cdot 100$	$\hat{\sigma}^w \cdot 100$	$\hat{\sigma}^m \cdot 100$	$\hat{\rho}_u$
False perception	0.88 (0.03)	0.90 (0.05)	2.30 (0.42)	43.92 (23.27)	1.21 (0.23)	0.13 (0.27)	0.00 –
Correct perception	0.53 (0.03)	0.90 (0.05)	2.49 (0.14)	1.10 (0.08)	1.20 (0.08)	0.30 (0.05)	0.95 (0.02)
True values	0.53	0.90	2.50	1.10	1.20	0.30	0.95

Notes: Mean estimates using 1000 runs with a sample length of 152 observations. Standard deviations in parenthesis.

The price stickiness parameter is estimated to be equal to $\hat{\theta}_p = 0.88$ on average with quite a narrow empirical standard deviation. This parameter value would imply that prices on average remain non-optimized for a little bit more than two *years* instead of the two *quarters* in the true data generating process. Not surprisingly perhaps, a noticeable bias also emerges for the standard deviation of the innovation to the price markup shock (first row, second column in Figure 5).

In a second experiment, we consider the case in which the econometrician also allows for serial correlation in the cost-push shock and estimates ρ_u (correct perception). In this case the bias vanishes altogether, as the results shown in the second column of table 3 make clear.

We next turn to the question whether this bias would be recognizable to an economist/econometrician who is conducting estimation and inference on just one set of data. Towards this end Table 4 shows mean root-mean squared forecast errors (in sample) for both estimated variants of the model. Evidently, the correctly specified model dominates the misspecified model in both the inflation forecast

Table 4: RMSE WITH MAXIMUM LIKELIHOOD

	r_t	π_t	$\tilde{w}_t - \tilde{y}_t$	n_t	Δy_t
False perception	0.27	0.66	1.35	1.21	1.22
Correct perception	0.22	0.51	1.29	1.17	1.18
Relative RMSE (in %)	22.35	29.35	5.00	3.41	3.22
% of time false model has lower RMSE	0.00	0.00	3.40	4.90	5.90

Notes: Root mean squared errors (in sample, multiplied by 100) using 1000 runs with a sample length of 152 observations. The true correlation of price markup shock is $\rho_u = 0.95$. The perceived correlation of price markup shock is $\rho_u = 0.00$.

and the interest rate forecast dimension. Yet, while the RMSEs speak a clear language and so would likelihood ratio tests, most economists who firmly believe that their estimated model should feature no serial correlation in the cost-push shock will, we suppose, be quite hesitant to opt for the larger model despite the econometric evidence.⁹

Table 5: MEAN ABSOLUTE DIFFERENCE IN IMPLIED STANDARD DEVIATIONS

	r_t	π_t	$\tilde{w}_t - \tilde{y}_t$	n_t	Δy_t
False perception	0.09	0.05	0.60	0.79	0.07
Correct perception	0.11	0.06	0.72	0.78	0.06

Notes: Mean difference of standard deviation implied by the estimates of the model and the standard deviation of the data (both multiplied by 100). 1000 runs with a sample length of 152 observations were drawn. True correlation of price markup shock is $\rho_u = 0.95$. Perceived correlation of price markup shock: $\rho_u = 0.00$.

As a final diagnostic, Table 5 shows the mean absolute difference between the standard deviations implied by the estimation of the respective model on the respective time series and the standard deviation of the data series on which the model was estimated. Here, the misspecification is less visible, if at all.

We conclude this section noting that, as with GMM, autocorrelated cost-push shocks if not accounted for in the ML estimated induce a substantial upward bias in the estimated degree of price rigidity. In other words, the problem is not specific to GMM, where the orthogonality conditions are violated in the presence of autocorrelated cost-push shocks. Our results thus suggest that there is generic economic mechanism which induces the bias: cost-push shocks by lowering the labor share reduce

⁹On the economic side, current generation DSGE models are often criticized for lacking internal propagation, see, for instance, Cogley and Nason (1995). Serially correlated shocks can partially make up of the resulting lack of persistence. This may, however, not appeal to quite a few economists since economic theory provides little guidance with regard to the autocorrelation structure of the shocks. While the tests speak a clear language, an economist may still ask parsimony.

the observable pass-through of the latter into inflation.

5 Other approaches

There is a vast literature on the NKPC and different routes have been taken to confront it with the data. To compare results from different studies it is important to note that only in the baseline NKPC model there is a direct mapping from the coefficient κ_p , which measures the pass-through of marginal costs into inflation, to the probability that prices are kept fixed, measured by the parameter θ_p . Recall from equation (3) that absent indexation, we have

$$\kappa_p = \frac{(1 - \beta\theta_p)(1 - \theta_p)}{\theta_p}. \quad (13)$$

Introducing real rigidities or indexation in the New Keynesian model, implies that a given value of the slope, κ_p , corresponds to a lower degree of stickiness, θ_p , relative to the baseline model. It should be stressed that such rigidities do not alter the pass-through of marginal costs into inflation, but simply provide a modified structural interpretation of the same reduced form slope coefficient κ_p . Autocorrelated cost-push shocks, in contrast, directly affect the estimate of κ_p . To make results comparable we therefore focus on this parameter, or on the degree of price rigidity which κ_p would be implied in the baseline specification of the NKPC, i.e. by equation (13).

A widely discussed study by Sbordone (2002) estimates the inverse of κ_p by matching the actual price/unit-labor-costs ratio with the one implied by the model. She finds a value of 18.3 which in the baseline model would imply an estimate of θ_p of about 0.8. Is this approach also likely to lead to biased estimates in the presence of autocorrelated markup shocks? To address this questions we estimated θ_p by minimizing the distance the between actual and fundamental inflation . Results from these experiments, which are available on request, suggest that this method also leads to substantially biased estimates in case cost-push shocks are autocorrelated.

An alternative way to estimate the NKPC is to match the impulse responses to a monetary policy shock in a full-fledged New Keynesian general equilibrium model with with those obtained from a structural VAR model. In contrast to full information ML estimation or Bayesian approaches only a certain aspect of the data is used in the structural estimation by matching impulse response functions to a monetary policy shock. Therefore this limited information approach allows one to remain agnostic about the full stochastic structure of the estimated model and to focus on the shock process of interest; see Meier and Müller (2006) for further discussion of this point.

At the same time, however, for this approach to make sense it is necessary that the identification assumptions imposed in the structural VAR model are consistent with the theoretical model. Generally, monetary policy shocks are identified in VAR models by assuming that prices and output do not respond contemporaneously to innovations in the interest rate (recursiveness assumption). This

assumption is violated in our baseline New Keynesian model and, in fact, by a variety of New Keynesian models, e.g. Smets and Wouters (2006). As a remedy several authors have assumed that prices and quantities are predetermined for one or several periods in addition to Calvo-type rigidities, see, for instance, Rotemberg and Woodford (1997).

We also performed a Monte Carlo experiment using a variant of the the above model modified to satisfy the recursiveness assumption. It turns out that in this case estimates of θ_p are not biased even if cost-push shocks are autocorrelated. Hence, these results highlight the advantage of the limited information approach. Nevertheless, its validity critically depends on the plausibility of the recursiveness assumption.¹⁰ Irrespective of one's stand regarding the identifying assumptions, it is interesting to observe that results obtained within this literature are quite heterogenous. To consider two of the most widely cited studies: Christiano, Eichenbaum, and Evans (2005) find $\hat{\theta}_p = 0.6$, where the model is explicitly designed to imply low price rigidities (by muting movements in marginal costs). Rotemberg and Woodford (1997), in contrast, find substantial price rigidities, once one accounts for the effects of strategic complementarities arising from real rigidities in their model.

6 Conclusion

Why is the New Keynesian Phillips curve apparently flat? In this paper we suggest a new interpretation of the macroeconometric evidence. Our explanation is consistent with microeconomic studies which imply that prices are adjusted frequently and by sizeable amounts. We start from the observation that full information estimation of medium-scale New Keynesian general equilibrium models provides evidence in favor of highly autocorrelated cost-push/markup shocks, see, e.g. Galí and Rabanal (2005) and Smets and Wouters (2006).

We show that in this case the orthogonality conditions exploited in the GMM approach to the New Keynesian Phillips Curve pioneered by Galí and Gertler (1999) are likely to be invalid. The estimator can be shown to be inconsistent and upward biased. We assess the quantitative implications of these complications through Monte Carlo experiments. As data generating process we use a variant of the New Keynesian model estimated and shown to provide a satisfactory account of U.S. time series by Galí and Rabanal (2005).

This data generating process assumes that prices have a mean duration of about 2 quarters and that the degree of autocorrelation of cost-push shocks is 0.95. Using a realistic sample size of 152 observations, we find that the average estimate for the mean duration of prices is about 11 quarters. Interestingly, standard tests fail to detect the misspecification of the model.

¹⁰Kehoe (2006) stresses that most monetary models do not satisfy the recursiveness assumption. Without the assumption of predetermined prices and output, the responses of the baseline New Keynesian cannot be successfully matched with those obtained from a VAR estimated on data generated by the model.

Moreover, the bias is not limited to GMM estimation. We also estimate the degree of price rigidity using ML techniques on the simulated time series and find a similar degree of price rigidity if we do not allow for autocorrelation in the cost-push shocks in the estimated model.

An intuitive interpretation of our results is as follows. Positive cost-push/markup shocks induce a negative response of the labor share, the measure for marginal costs used in the literature. To the extent that cost-push shocks play a role in driving the business cycle, a high labor share will be associated with a lower than average realization of the cost push-shock and this will be persistently the case if cost-push shocks are autocorrelated. A standard omitted variable logic suggests that the (partial) inflationary effect of the labor share (wages) on inflation is underestimated whenever the serial correlation in the markup shock is not accounted for.

We illustrate both theoretically and by means of simulations that this bias exists. Serially cost-push/markup shocks render the instrumental variables or GMM estimates of the NKPC inconsistent. Put differently, if cost-push shocks are an important source of business cycle fluctuations, but not considered in the estimation, as in the standard GMM approach, the New Keynesian Phillips Curve will appear to be implausibly flat, while, in fact, it is quite steep.

A Properties of the GMM estimator

This section derives the probability limit of the estimator in equation (7). To assess this more formally we follow Hamilton (1994, p. 238f.). This appendix considers the convergence properties when both the discount factor, β , and the slope of the Phillips curve, κ_p , are estimated. The formulae in Section 3, which restrict estimation to κ_p can easily be obtained following the lines of the exposition below. First, we note that absent indexation equation (3) can be written as follows

$$y_t = \boldsymbol{\alpha}' \mathbf{x}_{t+1} + \epsilon_{t+1}^{RE} + \tilde{u}_t,$$

where $\mathbf{x}_{t+1} = [\pi_{t+1}, l_t]'$, $\boldsymbol{\alpha} = [\beta, \kappa_p]'$ and ϵ_{t+1}^{RE} is the rational expectations error that ensures $E_t(\pi_{t+1}) \equiv \pi_{t+1} + \epsilon_{t+1}^{RE}$ and $\tilde{u}_t = \kappa_p u_t$. As above equation is linear, estimating it by GMM is equivalent to applying two stage least squares estimation (2SLS henceforth). Let the instrument vector be $\mathbf{z}_{t-1} = [\pi_{t-1}, l_{t-1}]'$.

The 2SLS estimator of $\boldsymbol{\alpha}$ when the sample size is T is

$$\hat{\boldsymbol{\alpha}}_{2SLS,T} = \left[\sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \mathbf{x}'_{t+1} \right]^{-1} \left[\sum_{t=1}^T \hat{\mathbf{x}}_{t+1} y'_t \right].$$

Above $\hat{\mathbf{x}}_{t+1} = \hat{\boldsymbol{\delta}}_T' \mathbf{z}_{t-1}$, the OLS fitted value when regressing \mathbf{x}_{t+1} on the instruments with

$$\hat{\boldsymbol{\delta}}_T = \left[\sum_{t=1}^T \mathbf{z}_{t-1} \mathbf{z}'_{t-1} \right]^{-1} \left[\sum_{t=1}^T \mathbf{z}_{t-1} \mathbf{x}'_{t+1} \right].$$

Using the Phillips curve to substitute for y_t in the 2SLS estimator, we obtain

$$\hat{\boldsymbol{\alpha}}_{2SLS,T} = \left[\sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \mathbf{x}'_{t+1} \right]^{-1} \left[\sum_{t=1}^T \hat{\mathbf{x}}_{t+1} (\boldsymbol{\alpha}' \mathbf{x}_{t+1} + \epsilon_{t+1}^{RE} + \tilde{u}_t)' \right],$$

So

$$\hat{\boldsymbol{\alpha}}_{2SLS,T} = \boldsymbol{\alpha} + \left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \mathbf{x}'_{t+1} \right]^{-1} \left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} (\epsilon_{t+1}^{RE} + \tilde{u}_t) \right].$$

Assuming that \mathbf{x}_{t+1} and \mathbf{z}_{t-1} are jointly covariance-stationary and ergodic for second moments, we have that

$$\hat{\boldsymbol{\delta}}_T \xrightarrow{p} [E(\mathbf{z}_{t-1} \mathbf{z}'_{t-1})]^{-1} [E(\mathbf{z}_{t-1} \mathbf{x}'_{t+1})] =: \bar{\boldsymbol{\delta}}.$$

Under the same assumptions

$$\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \mathbf{x}'_{t+1} \xrightarrow{p} Q,$$

where $Q = \bar{\boldsymbol{\delta}}' [E(\mathbf{z}_{t-1} \mathbf{x}'_{t+1})]$.

In addition,

$$\left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} (\epsilon_{t+1}^{RE} + \tilde{u}_t) \right] = \left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \epsilon_{t+1}^{RE} \right] + \left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \tilde{u}_t \right],$$

the first term of which will typically converge to zero in probability by means of a suitable law of large numbers. The final term, however, is given by

$$\left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} \tilde{u}_t \right] = \hat{\boldsymbol{\delta}}_T' \left[\frac{1}{T} \sum_{t=1}^T \mathbf{z}_{t-1} \tilde{u}_t \right].$$

Taking the probability limit we thus have that

$$\left[\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{x}}_{t+1} (\epsilon_{t+1}^{RE} + \tilde{u}_t) \right] \xrightarrow{p} \bar{\boldsymbol{\delta}}' [E(\mathbf{z}_{t-1} \tilde{u}_t)].$$

Note that in our example

$$E(\mathbf{z}_{t-1} \tilde{u}_t) = E \left(\left[\begin{array}{c} \pi_{t-1} \\ l_{t-1} \end{array} \right]' \tilde{u}_t \right) = \rho_u \kappa_p E \left(\left[\begin{array}{c} \pi_{t-1} \\ l_{t-1} \end{array} \right]' u_{t-1} \right) \neq 0 \text{ unless } \rho_u = 0.$$

Summarizing,

$$\hat{\boldsymbol{\alpha}}_{2SLS,T} \xrightarrow{p} \boldsymbol{\alpha} + Q^{-1} \bar{\boldsymbol{\delta}}' \kappa_p \rho_u E \left(\left[\begin{array}{c} \pi_{t-1} \\ l_{t-1} \end{array} \right]' u_{t-1} \right),$$

illustrating the inconsistency of the GMM approach with serially correlated price-markup shocks.

In figure 6, we illustrate the inconsistency of the estimates for β and κ_p in the calibrated model outlined in section 2 - for various values of ρ_u , confirming the results discussed in the main text. Note that as β is left unrestricted in the estimation, the bias in the slope of the Phillips curve is hardly affected. Since estimates of β are upward biased, the implied price stickiness shows a smaller bias than in the case presented in Section 3.




Figure 6: β AND κ_p ARE ESTIMATED. THEORETICAL PROBABILITY LIMITS OF 2SLS ESTIMATORS FOR VARYING DEGREES OF SERIAL CORRELATION, ρ_u , IN THE COST-PUSH SHOCK. Dashed lines show the true values. From left to right: estimator for the discount-factor, β , for the slope of the Phillips curve, κ_p , and for the fraction of firms which do not reoptimize their price, θ_p .

References

- D. Altig, L. J. Christiano, M. Eichenbaum, and J. Lindé. Firm specific capital, nominal rigidities and the business cycle. NBER Working Paper 11034, 2005.
- E. Baharad and B. Eden. Price rigidity and price dispersion: Evidence from micro data. *Review of Economic Dynamics*, 7:613–641, 2004.
- M. Bils and P. J. Klenow. Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5):947–985, 2004.
- A. Blinder, E. Canetti, E. Lebow, and J. Rudd. *Asking about Prices*. Russell Sage Foundation, New York, 1998.
- L. J. Christiano, M. Eichenbaum, and C. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, 2005.
- T. Cogley and J. M. Nason. Output dynamics in real-business-cycle models. *American Economic Review*, 85(3):492–511, June 1995.
- M. Eichenbaum and J. D. Fisher. Evaluating the calvo model of sticky prices. NBER Working Paper 10617, 2004.
- J. Galí and M. Gertler. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44:195–222, 1999.
- J. Galí and P. Rabanal. Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar U.S. data? In M. Gertler and K. Rogoff, editors, *NBER Macroeconomics Annual 2004, Volume 19*, pages 225–288. MIT Press, Cambridge MA, 2005.
- J. Galí, M. Gertler, and J. D. López-Salido. European inflation dynamics. *European Economic Review*, 45:1237–1270, 2001.
- J. Galí, M. Gertler, and J. D. López-Salido. Erratum to 'European inflation dynamics'. *European Economic Review*, 47:759–760, 2003.
- J. Galí, M. Gertler, and J. D. López-Salido. Robustness of the estimates of the hybrid new keynesian phillips curve. *Journal of Monetary Economics*, 52:1107–1118, 2005.
- J. D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, New Jersey, 1994.
- P. J. Kehoe. How to advance theory with structural vars: Use the sims-cogley-nason approach. NBER Working Paper 12575, 2006.
- J. Lindé. Estimating new-keynesian phillips curves: A full information maximum likelihood ap-

- proach. *Journal of Monetary Economics*, 52:1135–1149, 2005.
- B. Maćkowiak and M. Wiederholt. Optimal sticky prices under rational inattention. mimeo, 2006.
- A. Meier and G. J. Müller. Fleshing out the monetary transmission mechanism: Output composition and the role of financial frictions. *Journal of Money, Credit and Banking*, 38:2099–2134, 2006.
- E. Nakamura and J. Steinsson. Five facts about prices: A reevaluation of menu cost models. mimeo, 2006.
- J. R. Rotemberg and M. Woodford. Optimization-based econometric framework for the evaluation of monetary policy. In B. S. Bernanke and J. R. Rotemberg, editors, *NBER Macroeconomics Annual 2004, Volume 19*, pages 297–346. MIT Press, Cambridge MA, 1997.
- A. M. Sbordone. Prices and unit labor costs: A new test of price stickiness. *Journal of Monetary Economics*, 49:265–292, 2002.
- F. Smets and R. Wouters. Shocks and frictions in US business cycles. forthcoming in *American Economic Review*, 2006.
- M. Woodford. *Interest & Prices*. Princeton University Press, Princeton, New Jersey, 2003.