Is the New Keynesian Phillips Curve Flat?*

Keith Kuester
Federal Reserve Bank of Philadelphia

Gernot J. Müller
University of Bonn

Sarah Stölting
European University Institute, Florence

January 14, 2009

Abstract

This paper provides Monte Carlo evidence that GMM estimates of the New Keynesian Phillips curve are biased towards finding too much price rigidity if cost-push shocks are auto-correlated. This result may reconcile GMM estimates with the microevidence on price rigidities.

JEL Classification System: E31,E32,C22

Keywords: Price Rigidities, New Keynesian Phillips Curve, Cost-push shocks, GMM estimation

---

* Keith Kuester, Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106, Tel./Fax: +1 215 574-3415/4303, e-mail: keith.kuester@phil.frb.org.
Gernot Müller (corresponding author), University of Bonn, Lennéstr. 37, 53113 Bonn, Germany, Tel./Fax: +49 228 73-1979/7953, e-mail: gernot.mueller@uni-bonn.de.
Sarah Stölting, European University Institute, Via della Piazzuola 43, 50133 Florence, Italy, e-mail: Sarah.Stolting@EUI.eu.
We thank Peter McAdam, participants at the CEF 2007 conference in Montreal, the EEA/ESEM Meeting 2007 in Budapest and the German Economic Association Meeting 2007 in Munich, as well as two anonymous referees for comments. The views expressed in this paper are those of the authors. They do not necessarily coincide with those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

The key parameter of the New Keynesian Phillips curve (NKPC) provides a measure for the extent of price rigidity. Using the Generalized Method of Moments (GMM) on U.S. time series Galí and Gertler (1999), hereafter GG, report estimates for their baseline specification implying that prices are fixed between five and six *quarters* on average. This is high or, put differently, the Phillips curve appears to be too flat given recent microeconomic evidence which suggests average price durations of about five *months*, see Bils and Klenow (2004).[1] Real rigidities reconcile this evidence on the presumption that the NKPC is actually as flat as the GMM evidence suggests, see e.g. Eichenbaum and Fisher (2007).

We present an alternative resolution to the macro-micro divide which suggests that the NKPC may not be flat in the first place. Full information estimates of New Keynesian models provide evidence of autocorrelated cost-push shocks, see e.g. Galí and Rabanal (2005) and Smets and Wouters (2007). This autocorrelation renders the GMM orthogonality conditions invalid. We conduct Monte Carlo experiments using the model estimated by Galí and Rabanal (2005), henceforth GR, as data-generating process. The GMM estimates of the NKPC suggest too much price stickiness: for the autocorrelation of cost-push shocks estimated by GR the GMM estimates imply price durations of up to 12 quarters although the true duration is just 2 quarters. Diagnostic tests fail to detect this violation.

2 GMM estimation of the NKPC with autocorrelated cost-push shocks

The standard New Keynesian Phillips curve is given by

$$\pi_t = \beta E_t \left\{ \pi_{t+1} \right\} + \kappa_p \left(l_t + u_t \right), \tag{1}$$

[1] Note that GG's estimates range up to price durations of eleven quarters—as do those of other macroeconometric studies, see our working paper version for further references. Regarding the microeconomic evidence there has been some debate recently, which mainly rests on how to treat sales. Sales are frequent and sales prices tend to revert to their previous level. For this reason some authors advocate to exclude sales when measuring price adjustments/price stickiness. Doing so, Nakamura and Steinsson (2008) find average price durations of 8 to 11 months (7 to 9 months when they take product replacement into account). While these finding imply more price stickiness relative to the findings of Bils and Klenow, it still falls short of what macroeconomic studies suggest.

where time discount factor $\beta \in (0,1)$ and $\kappa_p > 0$ is the slope of the NKPC. π_t denotes inflation and l_t the labor share. u_t denotes an exogenous cost-push shock where $u_t = \rho_u u_{t-1} + \varepsilon_t^u$. ε_t^u is a zero mean innovation and $\rho_u \in [0,1)$ is the autocorrelation of the cost-push shock.[2] In the Calvo-staggered formulation, $\kappa_p = \frac{(1-\beta\theta_p)(1-\theta_p)}{\theta_p}$, where θ_p measures the probability that a firm cannot reoptimize its price in a given period. The higher θ_p the lower the slope κ_p, i.e. the flatter the Phillips curve. The average price duration is given by $D = 1/(1-\theta_p)$. We briefly highlight implications of the autocorrelation of cost-push shocks for the estimate of price stickiness. We restrict the other parameter in the NKPC, β, to its true value.

The moment condition we use in GMM estimation is

$$E\left\{\left[\theta_p \pi_t - (1-\theta_p)(1-\beta\theta_p) l_t - \theta_p \beta \pi_{t+1}\right] \boldsymbol{z}_{t-1}\right\} = 0, \tag{2}$$

where $\boldsymbol{z}_{t-1}$ is a vector of instruments.

Rewriting (1) gives

$$y_{t+1} = \kappa_p l_t + \epsilon_{t+1}^{RE} + \tilde{u}_t, \tag{3}$$

where $y_{t+1} := \pi_t - \beta\pi_{t+1}$. Here ϵ_{t+1}^{RE} is the rational expectations error that ensures $\beta E_t(\pi_{t+1}) \equiv \beta\pi_{t+1} + \epsilon_{t+1}^{RE}$ and $\tilde{u}_t = \kappa_p u_t$.

Moment condition (2) is violated whenever cost-push shocks are serially correlated:

$$\begin{aligned} E\left\{\left[y_{t+1} - \kappa_p l_t\right] \boldsymbol{z}_{t-1}\right\} &= E\left\{\left[\epsilon_{t+1}^{RE} + \tilde{u}_t\right] \boldsymbol{z}_{t-1}\right\} \\ &= \rho_u E\left\{\tilde{u}_{t-1} \boldsymbol{z}_{t-1}\right\} \neq 0, \quad \text{if} \quad \rho_u > 0. \end{aligned} \tag{4}$$

This renders the GMM estimate of price stickiness, θ_p, or alternatively an estimate of the slope of the NKPC, κ_p, inconsistent.

3 Monte Carlo experiments to assess the size of the bias

In order to assess quantitatively the bias induced by the autocorrelation of cost-push shocks, we perform Monte Carlo experiments using the model of Galí and Rabanal (2005) as data generating process. The model is representative of the recent small to medium-scale generation of New Keynesian models. The model is a closed economy, in which labor is the only factor of production.

[2] See Woodford (2003, p. 448 ff.) for a detailed discussion how autocorrelated cost-push shocks result from exogenous variations in market power, variable tax distortions or other inefficient supply shocks.

Monopolistic competition in the product and labor markets gives firms and workers price-setting power. Yet, both prices and wages are subject to nominal rigidity, the dynamics of inflation—abstracting from indexation—being described by the NKPC (1). Consumption is subject to external habit persistence. Five shocks drive the economy: a productivity shock, a demand shock, a cost-push shock, a wage-markup shock and an innovation to the monetary (Taylor type) policy rule, that closes the model. The monetary policy shock is iid. All other shocks are highly serially correlated. Galí and Rabanal estimate this model by Bayesian techniques on post World War II quarterly U.S. data, and find that it provides a good account of the U.S. time series. We use the model to generate 1000 random time series assuming parameter values corresponding to the mean estimates reported by Galí and Rabanal (apart from a tiny backward-looking component in the Phillips curve, which we set to zero).

Regarding the sample size we consider i) a realistic sample size of 152 observations and ii) one of 2000 observations. Throughout we assume that the value of β is known to be 0.99, the value used in the simulation of the model and focus on the estimates of the degree of price rigidity, $\hat{\theta}_p$.[3] The value used in the simulation of the model is $\theta_p = 0.53$. Our estimation of θ_p in the NKPC mimics the seminal work of GG using moment condition (2) and four lags of inflation, the labor share, and output growth as instruments.[4]

Table 1 (Results of GMM estimation) about here.

The upper panel of Table 1 shows results for the small sample size. The first row gives the estimated degree of price stickiness and the diagnostics if cost-push shocks are uncorrelated, 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 GR.[5] In the first case, the median estimate of $\hat{\theta}_p = 0.53$ corresponds to its true value. The median duration of prices, D, is about 2 quarters.

In contrast, if cost push shocks are autocorrelated we find a median estimate of $\hat{\theta}_p = 0.92$ which

[3] Setting a parameter to a fixed value typically helps in identifying others, and particularly so if this value is the true one. In particular, setting β to its true value and focusing exclusively on the slope of the NKPC alleviates the weak instrument problems in GMM estimation of the NKPC stressed by Mavroeidis (2005).

[4] The optimal weighting matrix uses the Newey-West correction for the likely serial correlation of the orthogonality conditions.

[5] We keep the volatility of cost-push shocks, u_t, at the value estimated by GR when we vary the autocorrelation.

implies average price durations of more than 12 quarters. The estimator is also inconsistent in this case, see the large sample experiment reported in the lower panel of Table 1.[6]

Standard diagnostic tests fail to detect that the empirical model used in the GMM estimation is misspecified under autocorrelated cost-push shocks. For the small sample, the J-Test (see the fourth column of Table 1) fails to detect the violation of the orthogonality conditions: the null is rejected only in 2 percent of the draws.[7] Turning to the hypothetical sample size of 2000 observations, the power of the J-Test increases and the null is correctly rejected for about 80 percent of the draws in case of serially correlated cost-push shocks.

The Ljung-Box Q-test for autocorrelation of the residuals is frequently used in empirical work but cannot discern the mere presence of cost-push shocks from serially correlated shocks. The combined residual in the NKPC is given by $e_t := \epsilon_t^{\mathrm{RE}} + \tilde{u}_{t-1}$. But $\epsilon_{t-1}^{\mathrm{RE}}$ in the model is not orthogonal to $\tilde{u}_{t-1}$. e_t is therefore serially correlated whenever cost-push shocks are present. We generally find high rejection frequencies for all specifications (see the fifth column).

For the small sample size, Figure 1 plots the median estimate of θ_p against the various degrees of autocorrelation in the cost-push shock in the data-generating process.

Figure 1 about here.

Summarizing, under autocorrelated cost-push shocks GMM estimates imply too much price rigidity and diagnostic tests do not detect a model misspecification.[8]

4 Conclusion

Is the New Keynesian Phillips curve actually flat? In this paper we suggest a new interpretation of the macroeconometric evidence, consistent with microeconometric studies implying more frequent and sizeable price adjustments. If autocorrelated cost-push shocks are a pervasive feature of the data, as suggested by recent full information estimation of New Keynesian general equilibrium

[6] This inconsistency is not due to a weak instrument problem as the corresponding F-Tests in the third column of Table 1 show.

[7] This limited power squares well with results by Mavroeidis (2005).

[8] In the working paper version of this paper we show that the bias in the estimated degree of price rigidity is not limited to GMM estimation but also arises if minimum distance estimation or classical ML techniques are used.

models, GMM estimates are likely to be biased towards finding too much price rigidity. The bias can be substantial: in our Monte Carlo experiment we find GMM estimates implying average price durations of about 12 quarters, while, in fact, the true value in the simulation is about 2 quarters. Interestingly, standard tests fail to detect this misspecification.

GMM estimates are biased upwards for the following economic reason. In general equilibrium cost-push shocks lower the labor share, l_t, which is used as empirical proxy for marginal costs. To the extent that an observed fall in this proxy is triggered by a positive cost-push shock, a low realization of this proxy for marginal costs does not translate into lower inflation since the dis-inflationary effect of a lower labor share will be offset by the unobserved higher cost-push shock. As a result, the estimated pass-through of marginal costs (as measured by the labor share) to inflation appears to be small, i.e. the NKPC looks flatter than it actually is.

References

Bils, M., and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5), 947–985.

Eichenbaum, M., and J. D. Fisher (2007): "Estimating the Frequency of Price Re-optimization in Calvo-style Models," *Journal of Monetary Economics*, 54, 2032–2047.

Galí, J., and M. Gertler (1999): "Inflation Dynamics: A Structural Economometric Analysis," *Journal of Monetary Economics*, 44, 195–222.

Galí, J., and P. Rabanal (2005): "Technology Shocks and Aggregate Fluctuations: How well does the Real Business Cycle Model fit postwar U.S. data?," in *NBER Macroeconomics Annual 2004, Volume 19*, ed. by M. Gertler, and K. Rogoff, pp. 225–288. MIT Press, Cambridge MA.

Kuester, K., G. J. Müller, and S. Stölting (2007): "Is the New Keynesian Phillips Curve flat?," ECB Working Paper 809.

Mavroeidis, S. (2005): "Identification Issues in Forward-looking Models Estimated by GMM, with an Application to the Phillips Curve," *Journal of Money, Credit, and Banking*, 37(3), 421–448.

Nakamura, E., and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles," *American Economic Review*, 97, 586–606.

Woodford, M. (2003): *Interest & Prices*. Princeton University Press, Princeton, New Jersey.

Table 1: Results of GMM estimation

	Point Estimates		Diagnostics		
	$\hat{\theta}_p$	D	F-Statistic	J-Statistic	Q(4)
152 Observations:					
$\rho_u = 0$	0.53	2.14	3.86	7.54	22.03
	(0.36,0.69)	(1.57,3.22)	[0.96]	[0.00]	[0.94]
$\rho_u = 0.95$	0.92	12.55	109.78	7.99	14.08
	(0.81,0.98)	(5.22,52.19)	[1.00]	[0.02]	[0.88]
2,000 Observations:					
$\rho_u = 0$	0.53	2.13	44.82	9.72	269.32
	(0.49,0.57)	(1.98,2.32)	[1.00]	[0.02]	[1.00]
$\rho_u = 0.95$	0.95	21.59	2131.66	27.35	159.96
	(0.90,0.99)	(9.59,70.53)	[1.00]	[0.81]	[1.00]

Notes: Median values over 1000 draws. Values in parenthesis are 2.5% and 97.5% quantiles; for the diagnostics rejection frequencies for the null at a 5% level are given in square brackets. The "F-statistic" refers to an F-Test of the joint significance of the instruments when regressing the labor share on the instruments. "J-statistic" refers to a J-Test for violation of the orthogonality conditions used in the GMM estimation. "Q(4)" refers to a Ljung-Box test for the presence of serial correlation of the error terms in the GMM estimation—the null being that these are serially uncorrelated.

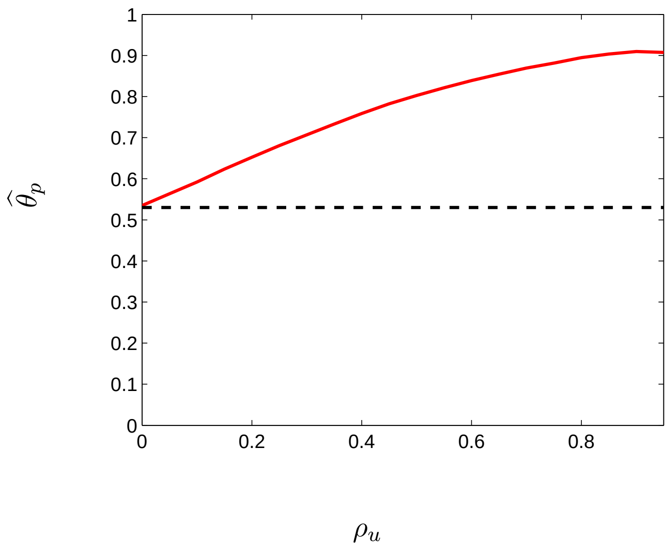

Figure 1: Median estimate for θ_p using GMM, for increasing values of ρ_u. *Notes:* Based on 1000 time series of length 152 observations. The dashed line displays the true value of θ_p.