

Fuhrer, Jeff

Working Paper

Intrinsic and inherited inflation persistence

Working Papers, No. 05-8

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Fuhrer, Jeff (2005) : Intrinsic and inherited inflation persistence, Working Papers, No. 05-8, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/55559>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

No. 05-8

Intrinsic and Inherited Inflation Persistence

Jeffrey C. Fuhrer

Abstract:

In the now conventional view of the inflation process, the New Keynesian Phillips Curve (NKPC) captures most of the persistence in inflation. The sources of persistence are twofold. First, the “driving process” for inflation—the output gap or, more commonly, real marginal cost—is itself quite persistent, and a casual inspection of the NKPC reveals that inflation must “inherit” this persistence. Second, a modest amount of backward-looking or indexing behavior imparts some “intrinsic” persistence to inflation. This latter source is generally thought to be of less importance than the former, as the degree of autocorrelation in the driving processes is substantial. This paper shows that in practice inflation in the NKPC inherits very little of the persistence of the driving process, and, contrary to conventional wisdom, it is intrinsic persistence that constitutes the dominant source of persistence. The paper explores the reasons for this and links them to two empirical observations. First, it has been difficult to develop a sizable coefficient on the driving process in NKPCs. Second, the shock that enters the NKPC, while often difficult to motivate economically, is large and is critical in distinguishing the sources of inflation persistence. While these observations help to clarify the behavior of inflation in NKPCs, they raise other fundamental questions about how to model inflation.

JEL Classifications: E31, E52

Jeffrey C. Fuhrer is Senior Vice President and Director of Research at the Federal Reserve Bank of Boston. His email address is jeff.fuhrer@bos.frb.org.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/wp/index.htm>.

Thanks to Fabio Ghironi, Matteo Iacoviello, Peter Ireland, Giovanni Olivei, and Paul Willen, and to seminar participants at the Federal Reserve Bank of Boston, the Norges Bank, and Boston College for helpful comments and suggestions.

The views expressed in this paper are solely those of the author and do not necessarily reflect official positions of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: September 23, 2005.

The progression of price-setting models has a long and lively history. Beginning with A. W. Phillips (1958) and continuing with Lucas (1972), Fischer (1977), Gray (1978), Taylor (1980), Calvo (1983), Rotemberg (1983), Gordon (1985), Roberts (1995), Fuhrer and Moore (1995), Galí and Gertler (1999), Erceg, Henderson, and Levin (2000), Galí, Gertler, and Lopez-Salida (2001), Sbordone (2002), Mankiw and Reis (2002), Christiano, Eichenbaum, Evans (2005) and countless others, the specifications have matured to include rational expectations, optimizing foundations, a more persistent driving process (real marginal cost), and a variety of “frictions” that allow the models to mimic the gradual response of inflation to a variety of shocks.

In recent years, much of the development of Phillips curves has centered on two issues: (1) the emergence of real marginal cost (versus an output gap measure) as the preferred driving variable in the specification, on both theoretical and empirical grounds; and (2) the incorporation of frictions into optimizing rational expectations models.¹ The frictions have included indexing (as in Christiano *et al.* 2005) and “rule-of-thumb” or “backward-looking” price setters (as in Galí and Gertler 1999). These frictions have been *ad hoc*, in that they are not micro-founded. Still, the common view is that, after allowing for just a little friction, the baseline model works well. For example, in a fairly recent summary, Galí (2003, Section 3.1) suggests that

The findings...are...quite encouraging for the baseline NKPC: while backward-looking behavior is often statistically significant, it appears to have limited quantitative importance. In other words, while the baseline pure forward-looking model is rejected on statistical grounds, it is still likely to be a reasonable first approximation to the inflation dynamics of both Europe and the U.S.

This paper will provide theoretical analysis and empirical evidence that largely contradicts this emerging consensus on price-setting models. It will show that, regardless of the persistence in the driving process, very little of that persistence is inherited by inflation in the conventional NKPC. This result runs counter to the common intuition that inflation in the NKPC directly inherits the persistence of the driving process, which, in the case of real marginal cost (or proxies thereof), is quite considerable. In fact, inflation does inherit some of the

¹ A third development has been the inclusion of serially correlated shocks to the models. A model with serially correlated shocks is considered in Section II.

persistence of the driving process, but, in the models commonly in use, the amount that it inherits is remarkably small.

So how does this seemingly counterintuitive result arise? There are two reasons: (1) The coefficient on the driving variable in NKPCs is estimated to be very small, on the order of .001 to .05, and (2) In addition to the shock that impels the driving process and thus indirectly influences inflation, there is another shock that disturbs the Phillips curve directly. The paper will show that the variance of that shock is large, generally at least as large as the shock driving real marginal cost or the output gap.

As the paper demonstrates below, those two facts together imply a very attenuated inheritance of the driving variable's persistence into the inflation process. A simple intuition for this result is as follows. Consider the purely forward-looking version of the NKPC displayed below.

$$\begin{aligned}\pi_t &= \beta E_t \pi_{t+1} + \gamma y_t + e_t \\ y_t &= \rho y_{t-1} + u_t\end{aligned}\tag{1.0}$$

If one iterates the top equation forward, one sees that inflation is simply a discounted sum of future y 's plus the error term e_t , which is assumed to be *iid* for the moment.² If there were no other shock in the model—if e_t were identically zero—then inflation's dynamic properties would be solely determined by those of the driving process y .

However, in the presence of a second shock, the intuition about inflation persistence changes. In that case, one can think of the simple forward-looking model as the sum of an AR(1) process y and an uncorrelated shock e . The persistence of the AR(1) process—summarized by its autocorrelation function—decays geometrically at rate ρ . The persistence of the shock process is rather uninteresting: Its autocorrelation function equals one at lag zero, and zero at all other lags. Which of these two processes dominates inflation's autocorrelation properties depends on two parameters, γ and the variance of e (relative to the variance of u). The larger is γ , the more the mix looks like the AR(1) process, and the less it looks like white noise. The larger is the

² We consider the ramifications of a serially correlated shock below.

variance of e relative to that of u , the more the process looks like white noise and the less it looks like an AR(1) process.

If γ is relatively small, and the variance of e relatively large, then the frictions added to the NKPC—the sources of “intrinsic persistence” in the model—will no longer be quantitatively unimportant but statistically significant additions. They will be of first-order importance to the model. But it also follows that the optimizing foundations, through which the forward-looking model with marginal cost as the driving process is motivated, become correspondingly less important for explaining inflation behavior. Thus, it becomes critical to understand *what* the inflation shock is and why the estimated coefficient on the driving process is so small.³ This paper will provide only partial answers to these questions.

The paper demonstrates analytically the propositions about inherited persistence for the forward-looking model in Section I. It analyzes the case of the hybrid model in Section II. Section III considers some extensions, including a model with explicit monetary policy. It also considers the implications of possible recent changes in the persistence of inflation. Section IV examines reduced-form properties in the data that will lead to structural models that embody a small γ and a relatively large variance of the inflation shock. Section V concludes.

I. The Purely Forward-Looking Model

Consider the canonical hybrid New-Keynesian Phillips curve (HNKPC), which may be expressed as⁴

³ The same points are demonstrated for the Mankiw-Reis model of price-setting in Fuhrer (2002). There, the presence of large “mark-up shocks,” which are the equivalent of inflation shocks in the HNKPC, similarly imply that inflation inherits very little of the driving variable’s persistence.

⁴ A related specification allows lagged inflation to Granger-cause the driving process y .

$$\begin{aligned}\pi_t &= (\beta - \mu) E_t \pi_{t+1} + \mu \pi_{t-1} + \gamma y_t + e_t \\ y_t &= \rho y_{t-1} + \delta \pi_{t-1} + u_t\end{aligned}$$

Because this modification adds no intrinsic persistence to inflation, its implications for the autocorrelation properties of inflation are virtually identical to those of the model in equation (1.1), for any plausible value of δ . In addition, for the data employed in this paper, estimates of δ tend to be nearly zero and insignificantly different from zero.

$$\begin{aligned}
\pi_t &= (\beta - \mu) E_t \pi_{t+1} + \mu \pi_{t-1} + \gamma y_t + e_t \\
y_t &= \rho y_{t-1} + u_t \\
Var(e_t, u_t) &= \Sigma
\end{aligned} \tag{1.1}$$

where π denotes inflation, y is a driving variable, typically a proxy for real marginal cost or the output gap, μ and $(\beta - \mu)$ are the weights on past and expected inflation, and γ is the coefficient on the driving process.⁵ The baseline case will assume that e , the “inflation shock,” is a white noise *iid* shock, although that assumption is relaxed below. The second equation specifies the simplest persistent process for the driving variable y , a first-order autoregression with autoregressive parameter ρ , which is set to 0.9 in all of the exercises below.⁶ The covariance matrix of the error processes is denoted by Σ and will be assumed diagonal throughout. However, the relative sizes of the shock variances will be allowed to vary and will be shown to have important effects on inflation persistence.

The analytical autocorrelation function for inflation

The solution to the HNKPC model may be expressed as a vector first-order state-space system:

$$x_t = \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = A \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + S_0^{-1} \begin{bmatrix} e_t \\ u_t \end{bmatrix} \tag{1.2}$$

where A is the matrix of reduced-form solution coefficients (see Anderson and Moore, 1985), and S_0 is defined in Fuhrer and Moore (1995). For this simple model with $\mu=0$, A and S_0 are

$$A = \begin{bmatrix} 0 & \frac{\gamma\rho}{1 - \rho\beta} \\ 0 & \rho \end{bmatrix}; S_0^{-1} = \begin{bmatrix} 1 & \frac{\gamma}{1 - \rho\beta} \\ 0 & 1 \end{bmatrix} \tag{1.3}$$

⁵ While I do not make this explicit here, one can map the coefficient γ into the underlying frequency of price adjustments, as in Woodford (2003), chapter 3, eq. 2.13, or Galí and Gertler (1999), equation 16. As the fraction of prices that remain fixed each period increases, the coefficient on marginal cost declines. Thus a rise in γ implicitly corresponds to an increase in the frequency of price adjustment or, equivalently, to an increase in price flexibility.

⁶ This value corresponds to estimates obtained later in the paper. The qualitative results in the paper are unchanged by a value for ρ up to 0.95.

Note that the structure of A implies that the lagged inflation rate does not enter the solution for current inflation. The structure of S_0 implies that the relative effects of the two shocks on inflation will depend critically on γ .

Denote the k -period ahead variance of x by V_k , where $V_k = AV_{k-1}A'$, with V_0 initialized to $S_0^{-1}\Sigma S_0^{-1'}$. The unconditional variance of x , denoted V , is the convergent sum of the V_k . Then the correlation of the vector x_t with x_{t-k} is simply $\Gamma_k = A^{k-1}\Gamma_1$, with $\Gamma_1 = V$. Hence, the matrices that determine the autocorrelation properties of x are the transition matrix A and the unconditional variance matrix V .

Using these two matrices and the definition of the unconditional variance, we can show that the unconditional variance of inflation for the NKPC model is a linear combination of the variances of e and u :

$$V = Var \begin{bmatrix} \pi \\ y \end{bmatrix} = \begin{bmatrix} \frac{\left[\gamma + \frac{\beta\gamma\rho}{1-\rho\beta} \right]^2}{1-\rho^2} \sigma_u^2 + \sigma_e^2 & \frac{\gamma\sigma_u^2}{(1-\rho\beta)(1-\rho^2)} \\ \frac{\gamma\sigma_u^2}{(1-\rho\beta)(1-\rho^2)} & \frac{\sigma_u^2}{1-\rho^2} \end{bmatrix} \quad (1.4)$$

The first term in the unconditional variance of inflation is the unconditional variance of y scaled by $\left[\gamma + \frac{\beta\gamma\rho}{1-\rho\beta} \right]^2$. The weight on the variance of y is strictly increasing in γ and ρ , so the larger are γ and ρ , the larger is the relative influence of σ_u^2 and the smaller is the relative influence of σ_e^2 in the variance of π .

The autocorrelations for y_t take the expected form for an AR(1) process. The autocorrelations for inflation at horizon i are denoted by Γ_i ; they are⁷

⁷ In an unpublished comment, Galí (2005) derives the solution to these models in the case of $e=0$.

$$\Gamma_i = \frac{\rho^i \gamma^2}{a \sigma_e^2 + \gamma^2} \quad (1.5)$$

$$a = (1 - \rho^2)(1 - \rho\beta)^2$$

Clearly, the autocorrelations decay at the rate ρ . The term $\frac{\gamma^2}{a\sigma_e^2 + \gamma^2}$ sets the initial “level” of the autocorrelation function, with the rate of decay from the initial level dependent only on ρ . Thus, the difference between a persistent and a non-persistent inflation rate in this model will hinge on how large is the first autocorrelation: Do the autocorrelations jump down toward zero immediately, or do they decay from near one? As suggested above, the answer to this question must depend upon the extent to which y feeds into π (that is, how large is γ), and the relative size of the variances of the shock hitting the inflation equation and the shock hitting the driving process. With σ_u^2 normalized to 1, it is straightforward to show that equation (1.5) implies

$$\frac{\partial \Gamma_1}{\partial \gamma} > 0; \frac{\partial \Gamma_1}{\partial \sigma_e^2} < 0 \quad (1.6)$$

That is, the smaller is the influence of y on π , the smaller is the initial autocorrelation. The larger is the variance of e relative to u , the smaller is the initial autocorrelation.

Note also that in the case in which the stochastic dimension of $[\pi, y]$ is 1—that is, for simplicity $\sigma_e^2 = 0$, so that the only shock in the system is u —then the autocorrelations of inflation take the simpler form

$$\Gamma_i = \rho^i$$

Not surprisingly, in this special case, inflation follows exactly the same AR(1) process as y .⁸

As it turns out, extreme values of σ_e^2 and/or γ are not required to imply a very small first-period autocorrelation for inflation, even when the autocorrelation of y is considerable. Table 1 displays the value of the first autocorrelation of inflation for various values of σ_e^2 and γ . Because the autocorrelations following the first will die out geometrically at rate ρ , this first

⁸ The autocovariances of π will still depend upon γ , but the autocorrelations for π are identical to those for y .

autocorrelation is a sufficient statistic for the entire function, once one knows ρ . For reference, the last row of the table displays the value of the first autocorrelation when σ_e^2 and γ take on values that are implied by common estimates in the empirical literature on NKPCs.⁹

Table 1

Value of Γ_1 for selected values of σ_e^2 and γ

σ_e^2	Γ			
	.02	.05	.1	.2
0	0.90	0.90	0.90	0.90
1	0.12	0.44	0.71	0.84
2	0.06	0.29	0.59	0.79
3	0.04	0.22	0.50	0.75
5	0.03	0.14	0.39	0.68
$\sigma_e^2=3, \gamma=.03$	0.092			

As the table suggests, depending on the parameter configuration, the first autocorrelation of inflation can range from essentially zero for high values of σ_e^2 and low values of γ to ρ for the opposite. The relative size of the shock variances is not often reported, but given the evidence presented below, it will be unusual to find $\sigma_e^2 \ll 1$. Thus, this table

⁹ Note that the empirical analysis in this paper employs annualized inflation rates, so the appropriate adjustment must be made to scale both γ and σ_e^2 . Raising γ by a factor of four (converting the highest estimates in Galí and Gertler to an annualized basis) shifts the autocorrelations up, but they still lie outside the 90 percent confidence intervals. More broadly, one must work very hard to obtain sizable and significant estimates of γ . Estimates on quarterly inflation rates range from 0.001 or below in Rudd and Whelan (2001) to 0.037 in Galí and Gertler (1999). The estimates presented below generally lie well below 0.01, with only one estimate on *annualized growth rates* exceeding 0.03, and none significantly different from zero. Thus a γ of 0.03 is a quite generous annualized coefficient, given the number of near-zero estimates in the literature, and given the difficulty in developing significant estimates. The maximum likelihood estimates presented below, using annualized inflation rates and employing either the output gap or real marginal cost as the driving variable, develop estimates of γ of 0.011 and 0.001, respectively.

implies that most often, the first autocorrelation for inflation implied by the NKPC will fall at or below 0.1.

Therefore, the purely forward-looking version of the NKPC imparts only a trivial amount of persistence to inflation. Of course, when one includes a lag of inflation, as in the hybrid model, the interaction between lagged inflation and the forward-looking component of the model must be taken into account. The hybrid model is the subject of the next section.¹⁰

II. The Hybrid Model

Now, consider the hybrid NKPC (HNKPC), which sets $\mu > 0$ in equation 1.1. The algebra becomes somewhat more complex (see the details in Appendix A), but much of the intuition from the simple NKPC remains. Larger values of σ_e^2 and γ will imply lesser and greater inheritance, respectively, of the persistence in the driving process. Now, however, the degree of “backward-looking” or indexing behavior—the size of μ —becomes critical in determining the persistence of inflation implied by the model.

The key matrices A and S_0 for the state-space representation (equation (1.2)) of the hybrid model are

$$A = \begin{bmatrix} \lambda_s & \frac{\gamma\rho}{(\beta-\mu)(\lambda_b-\rho)} \\ 0 & \rho \end{bmatrix}; S_0^{-1} = \begin{bmatrix} \frac{\lambda_s}{\mu} & \frac{-\gamma\lambda_b\lambda_s}{\mu(\rho-\lambda_b)} \\ 0 & 1 \end{bmatrix} \quad (2.1)$$

where λ_b and λ_s are the unstable and stable roots, respectively, of the system.¹¹ We can write the solution to the model as ¹²

¹⁰ Ireland (2004) also emphasizes the centrality of this shock, in his model a “cost-push” shock, in achieving data consistency.

¹¹ The third root is always ρ . Appendix A shows that $\lambda_b\lambda_s = \frac{\mu}{\beta-\mu}$.

¹² This representation is a version of the familiar solution to the second-order difference equation,

$$\pi_t = \lambda_s \pi_{t-1} + \gamma f(\lambda_s, \lambda_b) \sum_{i=0}^{\infty} \lambda_b^i E_t y_{t+i} + \varepsilon_t$$

in which the forward-sum term in the equation is solved for the sum of the t -period expectations of y_t .

$$\pi_t = \frac{\mu\pi_{t-1} - \frac{\gamma\lambda_b}{\rho - \lambda_b} y_t + e_t}{(\beta - \mu)\lambda_b} = \lambda_s\pi_{t-1} + \frac{\gamma\lambda_s}{(\beta - \mu)\rho\lambda_s - \mu} y_t + \frac{\lambda_s}{\mu} e_t \quad (2.2)$$

This representation shows that, as is common for simple second-order difference equations of this type, the coefficient on lagged inflation in the HNKPC solution is the stable root of the system. The stable root, in turn, is a function of the parameters μ and β ; the dependence of the stable root and the first autocorrelation on μ is examined below.

The unconditional variance of inflation, denoted here by V_π , is again a weighted average of the underlying shock variances. The weights are given by

$$V_\pi = w_e\sigma_e^2 + w_u\sigma_u^2 = \frac{\lambda_s^2}{\mu^2(1 - \lambda_s^2)}\sigma_e^2 + \frac{\lambda_s\gamma^2[2\mu^2\rho(1 - \rho\lambda_s) + \lambda_s]}{(1 - \rho^2)(1 - \lambda_s^2)[\lambda_s\rho(\beta - \mu) - \mu]^2}\sigma_u^2 \quad (2.3)$$


The stable root plays a key role in determining the contributions of the two conditional variances to the unconditional variance of inflation. Figure 1 displays the variation in the ratio w_e/w_u as μ varies. The effect of μ on the contribution of σ_e^2 to V_π is not monotonic. Increasing μ from 0 to 0.4 slightly depresses the contribution of σ_e^2 . But as μ increases from 0.4 to 0.9, the relative weight on σ_e^2 rises by a factor of six. Thus relatively modest differences in μ imply significant differences in the contributions of the two variances. As the lower panel of the figure indicates, the larger is γ , the smaller is the relative contribution of σ_e^2 to V_π , but, in any case, the effect is relatively small.

The expression for the autocorrelations in the HNKPC is somewhat more complex than in the simple NKPC. Nonetheless, the autocorrelation function can be shown to decay approximately geometrically after the first few autocorrelations.¹³ As a result, again a critical question is how large is the first autocorrelation. It can be expressed as

$$\Gamma_1 = \frac{a}{b\sigma_e^2 - c\rho\mu} d + \lambda_s \quad (2.4)$$

¹³ The rate of decay is slower than ρ for the first few autocorrelations, and then converges to ρ as k gets large.

where $[a, b, c, d]$ are functions of the stable root λ_s (in turn a function of μ and β) and the underlying parameters $[\mu, \beta, \gamma, \sigma_e^2]$. As is the case for the purely forward-looking model above, it can be shown that Γ_1 is decreasing in σ_e^2 . As will be shown below, the additive term in λ_s



dominates Γ_1 , and both λ_s and Γ_1 rise almost one-for-one with μ .

Table 2 shows the value of Γ_1 for an array of values for σ_e^2 and μ . The table illustrates that, for values of these parameters in the range commonly estimated, one obtains a relatively small first autocorrelation—0.6 or below. The bottom panel of the table shows the first eight autocorrelations of inflation when γ , σ_e^2 , and μ are set to values consistent with parameter estimates in the literature. The autocorrelations die out quickly, and in the following section we will see that they die out significantly more quickly than those exhibited in the data.

Note in the first row of Table 2 that as σ_e^2 goes to zero, the autocorrelation of inflation is bounded below by ρ and rises quickly to one as μ increases.¹⁴ The message of this table is that if

¹⁴ Of course with $\sigma_e^2 = 0$ and $\mu = 0$, the first autocorrelation is ρ .

one wishes to be roughly data-consistent and to assume a relatively small fraction of backward-looking or rule-of-thumb price-setters—say 0.3—one must impose a relative variance of σ_e^2 near zero. As will be shown below, such an estimate appears to be strongly at odds with the data.

Table 2

Value of Γ_1 for selected values of σ_e^2 and μ , hybrid model


$\gamma=0.03, \beta=0.98, \rho=0.9$

		M				
σ_e^2		.1	.3	.5	.7	.9
0		0.92	0.96	0.99	1.00	1.00
1		0.32	0.60	0.89	0.96	0.98
2		0.23	0.53	0.86	0.96	0.98
3		0.20	0.50	0.86	0.96	0.98
5		0.16	0.47	0.85	0.96	0.98

Autocorrelations for $\sigma_e^2=3, \gamma=.03, \mu=.35$							
1	2	3	4	5	6	7	8
0.59	0.37	0.25	0.18	0.14	0.11	0.09	0.08

Figure 2 illustrates the dominance of λ_s —and thus μ —in determining the first autocorrelation (see equations (2.2) and (2.4) above). The figure plots the stable root along with the first autocorrelation as μ rises from 0 to 0.65. The stable root rises from about 0.5 to almost 0.9 as μ varies from 0.35 to 0.5. Correspondingly, the first autocorrelation of inflation rises from about 0.55 to 0.85 over this range. From this figure, it is clear that μ is the critical determinant of the autocorrelation properties of inflation in the HNKPC and that small variations in μ will imply significant differences in the model's implications for the autocorrelation of inflation.

Figure 2
Effect of μ on stable root and first autocorrelation of hybrid model




How much “hybrid” do we need in the NKPC to be roughly data-consistent?

Do commonly-employed estimates of μ and γ , in conjunction with a data-consistent process for the driving variable, imply a data-consistent amount of persistence for inflation? Of course, it is difficult to know what the data imply about inherited versus intrinsic persistence—this requires structural identifying restrictions. But the reduced-form persistence of inflation is relatively simple to compute and provides a useful benchmark against which to judge the implications of the structural hybrid NKPC.

We begin with full-sample estimates of a simple three-variable vector autoregression in the inflation rate, the federal funds rate, and real marginal cost.¹⁵ The full sample extends from 1966:Q1 to 2003:Q4. The autocorrelation for inflation that is implied by the VAR is derived in the same manner as described above for generic linear rational expectations models. Seventy and 90 percent confidence intervals are displayed for the VAR’s autocorrelation function, where the confidence intervals are computed by assuming that the vector of OLS estimates of the VAR parameters is drawn from a multivariate normal distribution.

¹⁵ See the data appendix for variable definitions.

Figure 3 displays the theoretical autocorrelation for inflation implied by the pure and hybrid NKPCs at the parameter values indicated. As the figure suggests, at a somewhat generous estimate of $\gamma = 0.03$, and μ at the estimate for the U.S. developed in Galí, Gertler, and Lopez-Salido ($\mu=0$ for the pure NKPC), the implied autocorrelation for inflation lies outside the 90 percent confidence interval of the VAR's inflation autocorrelation for the first 15 quarters, at which point the theoretical autocorrelation is essentially zero, and the VAR-based autocorrelation is insignificantly different from zero.^{16 17} The heavy dashed line shows the implied inflation autocorrelation with μ raised to 0.6. This parameter setting puts the theoretical hybrid autocorrelation in the middle of the distribution of estimates from the VAR.




¹⁶ Using a somewhat different methodology, Rudd and Whelan (2001) develop estimates of γ that are often an order of magnitude smaller than this.

¹⁷ A similar comparison that sets the relative variance of inflation to 1 produces essentially the same result. While the theoretical autocorrelations are shifted upward somewhat, they still lie completely outside the 90 percent confidence intervals for the VAR.

As discussed below, it may be that the simple three-variable VAR misrepresents both the variance and autocorrelation of inflation, as it excludes the effects of large relative price movements for energy and non-oil imported goods. Figure 4 displays the same exercise for a five-variable VAR that includes the relative price of oil and the relative price of imported goods (again, see the data appendix for details).

The inclusion of these variables does little to change the basic contours of the inflation autocorrelation, although the autocorrelations decay a bit more quickly toward zero in the five-variable VAR. Still, the qualitative conclusion remains: the pure and hybrid versions of the NKPC are unable to match the VAR's implications for the autocorrelation of inflation.




Recent work by Levin and Piger (2003) and O'Reilly and Whelan (2004) emphasizes the potential for time-variation in the intercept for inflation, which may influence estimates of inflation persistence. The next figure addresses this concern, again estimating a three-variable

VAR, but only over the period since mid-1984, a point that many have identified as a breakpoint for the volatility of macroeconomic time series, including output and inflation.¹⁸ With these somewhat lower autocorrelations and wider confidence intervals, the hybrid model with $\mu=0.35$ begins to skirt the now-wider 70 and 90 percent confidence intervals around the inflation autocorrelation. Now, a value of $\mu=0.5$ implies an autocorrelation function squarely in the middle of the distribution of VAR autocorrelations. This computation emphasizes a point made above: The autocorrelation of inflation in the hybrid NKPC is very sensitive to relatively small changes in μ . The difference between $\mu=0.35$ and $\mu=0.5$ can move the implied inflation autocorrelation from outside the confidence intervals to the middle of the distribution.

Figure 5

Comparison of theoretical ACFs with VAR (ulc) 1984–2003

$\rho=0.9, \gamma=0.03, \mu=0.35, \sigma_{\epsilon} = 3$




¹⁸ Choosing the breakpoint differently, say, to correspond to the change in the Fed Chairman in July 1987, makes little difference to the conclusions drawn from the figure.

How much of the persistence in the hybrid specification comes from the driving variable?

While the analysis above demonstrates that the persistence of inflation in the HNKPC derives mostly from the lagged inflation term, the next exercise calibrates the remaining contribution of the driving variable. Figure 6 displays the theoretical autocorrelation functions for the hybrid model for pairs of parameter values. The pairs of lines highlight the contribution to persistence from the lag of inflation ($\mu = 0.35$, $\mu = 0.5$) versus the driving variable, through both contemporaneous and expected future effects ($\gamma = 0$, $\gamma = 0.03$). Of course when $\gamma = 0$, the driving variable has no effect on inflation.

As the figure indicates, almost all of the persistence imparted to inflation in the hybrid specification arises from the effects of μ . For parameter values near the baseline chosen from the literature (the solid and dashed lines), the incremental difference between the autocorrelation for zero or non-zero γ is not zero, but it is quite small. For a value of μ that is data-consistent (the dotted and dashed-dot lines), the incremental difference is essentially zero. For a given degree of “backward-looking” behavior in the specification, the incremental addition from including the driving variable’s persistence is very small.

Figure 6
How Much Persistence from the Driving Variable?



Estimation of the hybrid model

A more direct way to compare the properties of the model with the data is to directly estimate the HNKPC specification. Maximum likelihood (ML) has been shown to have some attractive features for this class of Euler equation-based models (see, for example, Fuhrer and Rudebusch (2004) and Fuhrer and Olivei (2004)). In this section, both ML and conventional GMM estimates of the specification are presented.

The following table displays the ML estimates for the specification in equation (1.1), using either the output gap or real marginal cost as the driving variable.¹⁹ For this estimation, the sample is constrained to the Greenspan era, 1987:Q3 to the present. The discount rate is constrained to 0.98, and the remaining parameters are estimated freely. Two estimates of the standard error are presented, the first from the numerical Hessian of the optimization problem,

¹⁹ All data definitions appear in the data appendix. Note that the inflation data are annualized quarterly log changes.

and the second from the BHHH algorithm that uses only first-derivative information (Berndt, Hall, Hall, and Hausmann, 1974). The table also displays the ratio of the estimated shock variances.

Table 3			
Maximum Likelihood Estimates of Specification (1.1)			
Output Gap (CBO estimate of potential)			
Parameter	Estimate	SE	BHHH SE
μ	0.67	0.14	.071
β	0.98	-	-
γ	0.011	0.023	0.026
ρ	0.96	0.036	0.041
$\sigma_e^2 / \sigma_u^2 = 1.20$			
Real marginal cost			
Parameter	Estimate	SE	BHHH SE
μ	0.71	0.24	0.064
β	0.98	-	-
γ	0.001	0.058	0.031
ρ	0.91	0.059	0.053
$\sigma_e^2 / \sigma_u^2 = 0.65$			

As the table indicates, ML yields estimates that are consistent with the informal calibrations in Figures 3 and 5 above. The ML estimate of μ centers around 0.7, and is precisely estimated. It remains difficult to estimate a significant γ , and the point estimates are quite small. Replicating the exercise in Figure 6 around the ML estimates produces virtually identical results. At these estimates, the persistence inherited by inflation from the driving process is essentially nil. Note that for both driving variables, the estimate of ρ is quite high, so in principle inflation could inherit considerable persistence from the driving variable. But in the HNKPC specification, it does not.

The following table summarizes GMM estimates for a variety of samples and instrument sets. The instrument sets vary from “bare-bones” (three lags of inflation and marginal cost) to the “kitchen sink” (four lags of those two variables, plus an output gap, oil prices, and the federal funds rate). The baseline estimation sample spans the past 45 years. To examine the

stability of the estimates, the table provides results for subsets of those years that split at the current Fed chairman's term in mid-1987 and more recently.

The estimates of the “forward-looking” and “backward-looking” parameters vary considerably; in other work we address the difficulties in obtaining reliable estimates of these parameters via GMM as conventionally implemented (Fuhrer and Rudebusch (2004), Fuhrer and Olivei (2004)). The basic results for estimating γ are similar to those for ML: in no case is the estimated parameter on real marginal cost significantly different from zero. Only one estimate exceeds 0.03, and in general the estimates center on about 0.005 for this annualized-change inflation data. Two of the J -tests reject at conventional levels of significance for instrument set one. Lagging this instrument set one additional period raises the p -value for the J -statistic to .05 or above, leaving the parameter estimates and significance essentially unaffected. Thus it seems difficult to attribute the general result of a very small estimated γ , to inadequate exogeneity of the instruments.²⁰ A small estimated γ , generally 0.01 or smaller on *annualized* inflation rates, is the norm, regardless of estimation method, sample period, or instrument set.

<p style="text-align: center;">Table 3a GMM Estimates of Hybrid Specification Annualized inflation data $\pi_t = \mu\pi_{t-1} + (\beta - \mu)E_t\pi_{t+1} + \gamma y_t$</p>						
Estimation period	Instrument set	Estimated μ	Estimated $(\beta - \mu)$	Estimated γ	p -value of t -statistic for γ	p -value of J -statistic
1960:4-2003:4	1	0.61	0.38	0.0052	0.84	0.27
1960:4-2003:4	2	0.62	0.37	0.0081	0.79	0.48
1960:4-2003:4	3	0.48	0.52	-0.0015	0.95	0.73
1960:4-2003:4	4	0.52	0.48	0.000	0.99	0.66
1987:3-2003:4	1	-0.0044	1.002	0.020	0.55	0.013
1987:3-2003:4	2	0.87	0.14	0.011	0.71	0.55
1987:3-2003:4	3	0.17	0.83	0.0092	0.74	0.71

²⁰ These estimates were run in Eviews version 5.0, using a fixed 4-quarter Bartlett kernel, no pre-whitening, and simultaneous iteration of the parameter estimates and the weight matrix. A constant is included in each instrument list.

1987:3-2003:4	4	0.66	0.34	0.0058	0.79	0.53
1992:1-2003:4	1	0.059	0.94	0.018	0.64	0.015
1992:1-2003:4	2	0.73	0.27	0.0086	0.77	0.81
1992:1-2003:4	3	0.34	0.67	0.032	0.27	0.75
1992:1-2003:4	4	0.45	0.54	-0.0015	0.95	0.62
1960:4-1987:2	1	0.58	0.42	-0.0035	0.93	0.16
1960:4-1987:2	2	0.59	0.40	0.0026	0.95	0.38
1960:4-1987:2	3	0.49	0.51	-0.0098	0.80	0.73
1960:4-1987:2	4	0.54	0.46	-0.013	0.74	0.66
Instrument sets 1: 4 lags of inflation, real unit labor cost, output gap 2: 3 lags of inflation and real unit labor cost 3: 4 lags of inflation, funds rate, real unit labor cost, relative price of oil 4: 3 lags of inflation, funds rate, output gap, real unit labor cost, relative price of oil						

Autocorrelated inflation shocks

Many implementations of the NKPC, especially in fully articulated general equilibrium models such as Christiano, Eichenbaum, and Evans (2005) allow shocks to be autocorrelated, augmenting the behavioral dynamics of the model. This addition would obviously alter the model's implications for the autocorrelation of inflation. The appendix presents the key matrices A and S_0 for the case of the purely forward-looking model augmented with a serially correlated shock e_t . Not surprisingly, in this version of the model, the autocorrelations of inflation depend almost entirely on the size of the autocorrelation coefficient for the inflation shock e_t .

That certain types of inflation shocks—"cost-push" shocks, for example, from large changes in relative prices—might be autocorrelated is not controversial. But how autocorrelated are such shocks, and how much persistence do they contribute to inflation in the U.S.? We can get a feel for the degree of serial correlation that might plausibly be added to e_t by examining the autocorrelation of the relative oil price and import price series used in the VARs above. Interestingly, the autocorrelation of the change in the relative oil price, both over the full sample

and limited to the decade of the 1970s, is essentially zero. The autocorrelation of the change in relative import prices for the same two samples, is about 0.5. Adding an autocorrelated shock with relatively low persistence to the pure forward-looking NKPC would not qualitatively change the conclusions about the model.²¹

III. Some Extensions

Adding explicit monetary policy to the model

There are good reasons to believe that the persistence of inflation should be affected by the systematic component of monetary policy. For example, Fuhrer and Moore (1995a) show that in a data-consistent, forward-looking model, policy rules that respond more or less aggressively to inflation and output imply corresponding changes in the persistence of output and inflation. Could the addition of inertial interest-rate policy save the purely forward-looking NKPC?

In the models examined below, changes in the systematic component of monetary policy *do* alter the properties of the driving process and of inflation. But the intuition from the discussion above remains: Monetary policy affects inflation in this model through its effect on the current and expected values of the driving variable. While more inertial or aggressive monetary policy generally alters the persistence of the driving process, in the purely forward-looking model or in the hybrid model with modest μ , inflation is relatively unaffected by these changes.

To demonstrate this result, a simple inertial policy rule is added to the model without lagged inflation:

²¹ For example, using the autocorrelations derived for the model with equations (6.9) in the appendix, at the baseline parameter settings $\gamma=0.03$, $\sigma_\epsilon^2=3$, $\rho=0.9$, and $a=5$, the first autocorrelation of inflation is 0.51, a bit lower than the autocorrelation for the hybrid model with $\mu=0.35$ in Figure 3.

$$\begin{aligned}
\pi_t &= \beta E_t \pi_{t+1} + \gamma y_t + e_t \\
y_t &= \rho y_{t-1} - a(i_{t-1} - \pi_t) + u_t \\
i_t &= s i_{t-1} + (1-s)(a_\pi \pi_t + a_y y_t)
\end{aligned} \tag{3.1}$$

Can the addition of inertial monetary policy qualitatively change the conclusions about inherited versus intrinsic persistence in inflation?

While the algebra becomes more tedious, numerical examples serve to illustrate the point well. As Figures 7 and 8 demonstrate, without significant intrinsic persistence in the inflation process, the presence of inertial monetary policy does little to change the implications from the simpler model without monetary policy. Regardless of the size of γ , σ_e^2 , s , or the vigor with which monetary policy responds to inflation and output, and thus regardless of the persistence of y , inflation inherits quite little of the persistence of the driving process. When compared with the persistence implied by the full-sample VAR, the autocorrelations fall well outside the 90 percent confidence interval.

Figure 8 displays the results for the hybrid model. With relatively limited intrinsic persistence, the hybrid model cannot replicate the autocorrelation properties of inflation. Only setting $\mu=0.6$ (the lighter dashed line) puts the autocorrelation into the confidence region for the full-sample VAR autocorrelation function. If we perform the same comparison for a VAR estimated beginning in 1984 (not shown), some of the cases lie between the 70th and 90th percentiles of the distribution. But qualitatively, the results are the same. A data-consistent representation of inflation, even with inertial monetary policy, requires a significant weight on lagged inflation.

Figure 7
Inflation autocorrelation with endogenous monetary policy
Forward-looking model, various parameter settings





Figure 8
Inflation autocorrelation with endogenous monetary policy
Hybrid model, various parameter settings




Has the persistence of the driving variables changed over the past four decades?

Figure 5 suggests that the estimated persistence of inflation may have declined over the past two decades. If so, is this the result of a decline in the persistence in the driving variable, which could be in turn the result of a change in monetary policy (or a change in any other factor that influences the reduced-form persistence of output or marginal cost)?

In short, the answer is: No. Figure 9 displays a crude measure of persistence, the sample autocorrelations of an output gap, and a unit labor cost measure for three subsamples. While volatility of inflation and output have declined (as documented by many, including McConnell and Perez-Quiros 2000), the persistence of the key driving variables for the Phillips curve has remained just as it was in earlier decades.²² This observation suggests that one must look not to monetary policy or other changes in the driving process, but to changes to the intrinsic persistence in inflation to explain recent declines in inflation persistence.²³

²² Of course, these sample autocorrelations implicitly allow for changes in the intercept of the series at the indicated breakpoints. This has been a significant element of the debate over the possibility of changes in inflation persistence in recent data.

²³ The sum of the lag coefficients in a univariate autoregression for these series varies from 0.89 to 0.91 for the three samples indicated.



A slightly more sophisticated test of change in persistence may be obtained by performing an unknown (multiple) breakpoint test, using the methodology of Bai (1999). The test regression is

$$\Delta y_t = \alpha y_{t-1} + \sum_{i=1}^k \beta_i \Delta y_{t-i} + e_t \quad (3.2)$$

where α is an estimate of the sum of the lag coefficients in the univariate autoregression for y . The full sample begins in 1966:Q1 and ends in 2003:Q4. The smallest admissible subsample is set to ten percent of the sample size, and the critical value for rejection of n breaks in favor of $n-1$ breaks is 0.05. A value of $k=2$ appears to be sufficient for these two series, although the results do not depend importantly on the choice of k . The test for the output gap cannot reject a single break in favor of no breaks. The test for real unit labor cost finds a single break in 2000:Q1, at which point the α is found to have increased from -0.08 to -0.04, that is, the sum of the lag coefficients for the level of unit labor costs has increased from 0.92 to 0.96, so that the persistence of real unit labor cost has *increased*.

The same test performed on the inflation series used in this paper develops two breakpoints, one in 1972:Q4 and one in 1981:Q1. The estimated value of α rises from 0.49 in the

pre-1972 period, to 0.69 in the 1972-81 period, and to 0.77 in the post-1980 period. These estimates of persistence are lower than the full-sample estimate (0.93), perhaps because of shifts in the intercept, as suggested by the authors cited above.²⁴

IV. Empirical Evidence on the “Inflation Shock” and the Small Coefficient on the Driving Process

In this section, we examine empirical evidence bearing on the two deterrents to inheriting inflation persistence, the size of the shock to the inflation process, and the small estimated coefficient on the driving process. Properties of both unconstrained and constrained models of inflation are examined to explore further the source of the problems with inheriting persistence in the NKPC.

What is the shock to inflation?

In the model in which the driving variable is real marginal cost, many candidate interpretations of the shock—“supply shocks” such as large changes in the relative price of oil or non-oil imported goods, or shocks to trend productivity that shift the supply relation, or “mark-up” shocks of price over unit labor cost—are ruled out, as these are incorporated in the measure of marginal cost, and thus should appear as part of the shocks in the driving process.²⁵ Such shocks may well be autocorrelated, but because they perturb only the driving process, they would still constitute a source of inherited, not intrinsic persistence. There is no doubt that some measurement error distorts the measures of real marginal cost commonly used in the specification; if such error were autocorrelated, this would appear in the inflation shock. If the inflation shock were fairly small, measurement error might be a reasonable interpretation. The next section examines the size of the inflation shock.

²⁴ Note that an alternative interpretation of these results is that the lack of correspondence between changes in the persistence of inflation and the persistence of the driving process could mean that the NKPC model fits the recent data better than it does the data for the 1960s to 1980s.

²⁵ Shocks to the desired markup, which enters as an element in the nonlinear combination of parameters that pre-multiplies real marginal cost in the fully-articulated NKPC, would show up as shocks to the inflation equation.

How big is the variance of the shock to the inflation process?

Central to the discussion above about how much of inflation's persistence is inherited versus intrinsic is the size of σ_ϵ^2 , the variance of the inflation shock. Two approaches are used to measure the relative size of σ_ϵ^2 . The first looks at estimated variances from simple VARs, computing relative variances for the reduced-form errors. Of course, because of the well-known difficulties in associating reduced-form VAR errors with any underlying structural disturbance, this should only be done with some trepidation. Interestingly, the reduced-form errors are approximately orthogonal. This reduces somewhat the concern that the shock in the VAR's inflation equation is a linear combination of other underlying shocks.

The second approach employs the three structural models of inflation from Sections I, II and III (the NKPC, the HNKPC, and the HNKPC with explicit monetary policy, equations (1.0), (1.1), and (3.1), respectively). The U.S. data described in the data appendix are used, solving each of the models for the structural (or pseudo-structural) shocks for a variety of parameter values. Then, the ratio of the variance of the inflation shock to that of the driving variable is computed for each case. It is important to note that, in the first two cases, the identification of the driving process is suspect, as the simple AR(1) process likely serves as a reduced form for a more fully articulated aggregate demand relation and monetary policy rule. Only in the case of the HNKPC with explicit monetary policy can one claim to have identified underlying structural shocks.

Table 4									
How big is the inflation shock?									
VAR specifications									
Ratio of reduced-form inflation shock to driving variable shock									
Estimation range	π, r, y		π, r, ulc		π, r, y, p^o, p^{no}		π, r, ulc, p^o, p^{no}		
1966:Q1-2003	1.98		2.39		1.66		1.99		
1984:Q3-2003	5.45		1.16		4.34		1.07		
	Structural Models								
	Ratio of identified inflation shock to driving variable shock								
Parameter sets	Pure forward-looking model			Hybrid model		Hybrid with monetary policy			
1	7.41			1.08		1.17			
2	8.67			0.68		0.59			
3	8.93			1.31		1.22			
4	4.55			3.68		1.04			
	Parameter sets								
	β	γ	μ	γ	M	γ	S		
1	.98	.03	0.35	.03	0.35	.03	.8		
2	.98	.01	0.5	.03	0.5	.03	.8		
3	.50	.03	0.35	.01	0.35	.01	.8		
4	.98	.03	0.15	.03	0.35	.03	.0		

As Table 4 indicates, it is rather uncommon for the variance of inflation to be less than that of its driving process. For the VARs, the variance is about twice as large on average as the variance of the driving process. This finding is relatively invariant to the set of conditioning variables in the VAR. One might have assumed that partialling out the variation that arises from oil or non-oil import prices might significantly reduce the variance of the shock to inflation, but this is not the case.

In the “pure” NKPC model, the estimated variance is five to nine times greater than the driving process,²⁶ depending on the parameter values chosen. The hybrid model reduces the relative variance, as the presence of lagged inflation absorbs much of the autocorrelation that remains in the “pure” model’s errors. Still, the variance of the identified inflation shock is on

²⁶ Of course, some of this blow-up in variance arises from the significant autocorrelation left in the inflation shock for this model.

average about as large as that of the driving process. Adding explicit monetary policy leaves this conclusion unchanged.

While it is difficult to put a compelling economic interpretation on this shock, it is nearly impossible to relegate it to a small nuisance, perhaps attributable to the measurement error that no doubt plagues the standard proxy for real marginal cost. If the estimates above are of the right order of magnitude, there would have to be at least as large a variation in the measurement error as there is in the shock to the driving process. That seems implausible. Consequently, it appears that the “inflation shock” is central to the inflation process, and central to the debate over how much of inflation’s persistence is inherited versus intrinsic. What the inflation shock is remains an important challenge for inflation modeling.

Reduced-form evidence on the relationship between inflation and the driving process

In the purely forward-looking model, the restrictions imposed by the rational expectations assumption could force the estimated coefficient on the driving process towards zero in the following way.²⁷ When the model is purely forward-looking, a shock to the driving process will imply a “jump” in the response of inflation. To the extent that such jump behavior is at odds with the relatively inertial behavior of inflation in the data, estimates of γ may tend towards zero to minimize this counter-factual dynamic implication of the model.²⁸

This section examines variance decompositions of unconstrained VARS, which provide a reduced-form confirmation of the relatively small influence of the driving processes on inflation. The first step is to examine simple variance decompositions from VARs that incorporate the key variables in the NKPC system. Then, VARs that include other variables, including energy and import prices, are examined to determine the extent to which these reduced-form implications are stable across the full sample. Finally, the combination of

²⁷ Estimation methodology can also affect the estimated coefficients in this class of models, but this lies outside the scope of the current paper.

²⁸ See Estrella and Fuhrer (2002) for more analysis of this counterfactual “jump” implication of many dynamic models.

structural parameters that is most consistent with the variance decomposition evidence is determined. Not surprisingly, the conclusions here mirror those above.

Variance decompositions


The variance decomposition from a VAR computes the contribution to inflation's variance from the driving process for a particular ordering and orthogonalization of the errors. For the VARs considered, the covariance matrix of the residuals is close enough to diagonal that, in general, the ordering does not affect the variance decompositions.²⁹ To preview the conclusions, regardless of the ordering used in the VAR, or of the driving process, it will be the case that the fraction of the variance of inflation accounted for by the driving variable is quite small. The fraction accounted for by the own-lags of inflation is quite large. Thus, it will be difficult to estimate larger coefficients on the driving process, as to do so would imply an unrealistically large contribution of the driving process to the variance of inflation, as the autocorrelation functions above suggest.

First, consider two VARs in inflation, a driving variable (a measure of the output gap or real marginal cost), and the federal funds rate. The VARs are estimated for two sample periods, 1966:Q1 to 2003:Q4 (corresponding to the estimation period in Galí and Gertler (1999) and others), and 1984:Q3 to 2003:Q4, in recognition of the decline in the volatility of all variables and in the persistence of inflation that appears to have occurred over the past twenty years or so. Four lags of each variable are included, and both the conventional ordering (inflation and output respond to the funds rate only with a lag, and inflation also responds to output with a lag) are used, as well as alternative orderings.

Figure 10 displays the variance decomposition for the VAR that includes the output gap. As the figure shows, the variance of inflation attributed to the output gap is never larger than 12 percent. The contribution of the "inflation shock" drops slowly to 72 percent after 30 quarters.


²⁹ In the five-variable VAR with real marginal costs, there is some mildly significant correlation between the funds rate and oil prices, and between inflation and relative import prices.

Figure 10



The next figure displays the same variance decomposition for the VAR in which real marginal cost (here proxied by real unit labor cost) is the driving variable. The results are similar, but not identical, for this VAR. For the first two years, the contribution to the variance is below 20 percent. But by five to eight years out, real marginal cost accounts for slightly more than 30 percent of inflation’s variance. Still, the share of the variance attributable to the “inflation shock” never dips below 61 percent and falls below 75 percent only after 12 quarters.


Figure 11



One explanation for the relatively small contribution of the driving process—or, equivalently, the large contribution from the “inflation shock”—could be the prevalence of “supply shocks” in the U.S. data—food and energy, non-oil imported goods prices, etc. The following chart

depicts the variance decomposition for a VAR in which we explicitly account for the variation in the relative price of oil and the relative price of non-oil imported goods. Such shocks, if imperfectly reflected in measures of marginal cost, might contaminate the shock that appears to disturb the inflation equation.


Figure 12



The inclusion of these variables does not change the qualitative features of the variance decomposition. The inflation shock still accounts for the largest share of variance, and the output gap for a relatively small share.


Conditioning on important relative price changes has a more significant effect for the system with real marginal cost. As Figure 13 indicates, the contribution of real marginal cost now never exceeds 18 percent, and it looks very similar to the decomposition with the output gap in Figure 11. Some of the inflation shock is absorbed by the import price and oil price variables, so that its contribution diminishes to 40 percent over 30 quarters. But the contribution from the driving variable remains quite small.

Figure 13



Because we know that persistence and volatility have declined for many variables in recent years, estimation of the latter VARs is repeated over the period 1984:Q3–2003. Once again, the basic picture is unchanged. Regardless of the driving process included in the VAR, its contribution to inflation is quite small. Thus, it will be difficult to estimate a large and significant parameter on the driving process in the NKPC.

Figure 14



What parameter settings allow the model to match the evidence in the variance decompositions?

One can readily compute the variance decomposition for the structural model of equation (1.1). As demonstrated above, the solution for the model has an upper-triangular contemporaneous coefficient matrix with an (assumed) diagonal shock covariance matrix. Utilizing standard techniques, the variance decomposition of inflation for the structural model is computed for the following three parameter settings:

- a. “Baseline” case: $\mu = 0.6, \gamma = 0.03, \sigma^2_\epsilon=3$;
- b. “Low- μ hybrid” case: $\mu = 0.35, \gamma = 0.03, \sigma^2_\epsilon=1$;
- c. “Pure forward-looking” case: $\mu = 0, \gamma = 0.03, \sigma^2_\epsilon=0.1$;

The first parameter setting most closely matches the autocorrelation patterns in the reference VAR. The second corresponds to conventional estimates in the literature that estimates hybrid NKPC models. The third is a baseline forward-looking case, with no weight on lagged inflation, and very small variance for the inflation shock. Note that in all cases the variance decomposition for the driving variable is uninteresting (always 100 percent attributed to the driving variable, 0 percent to inflation).

As the figure indicates, the “baseline” parameter setting most closely matches the qualitative features of the reference VARs’ variance decompositions.³⁰ Just as the autocorrelation functions are quite sensitive to modestly lower μ , so too are the variance decompositions. The pure forward-looking case has little chance of matching the data-consistent variance decomposition, and the low- μ hybrid model with conventional parameter settings implies much too large a fraction of inflation variance attributed to the driving variable, and much too little attributed to inflation’s own shock.

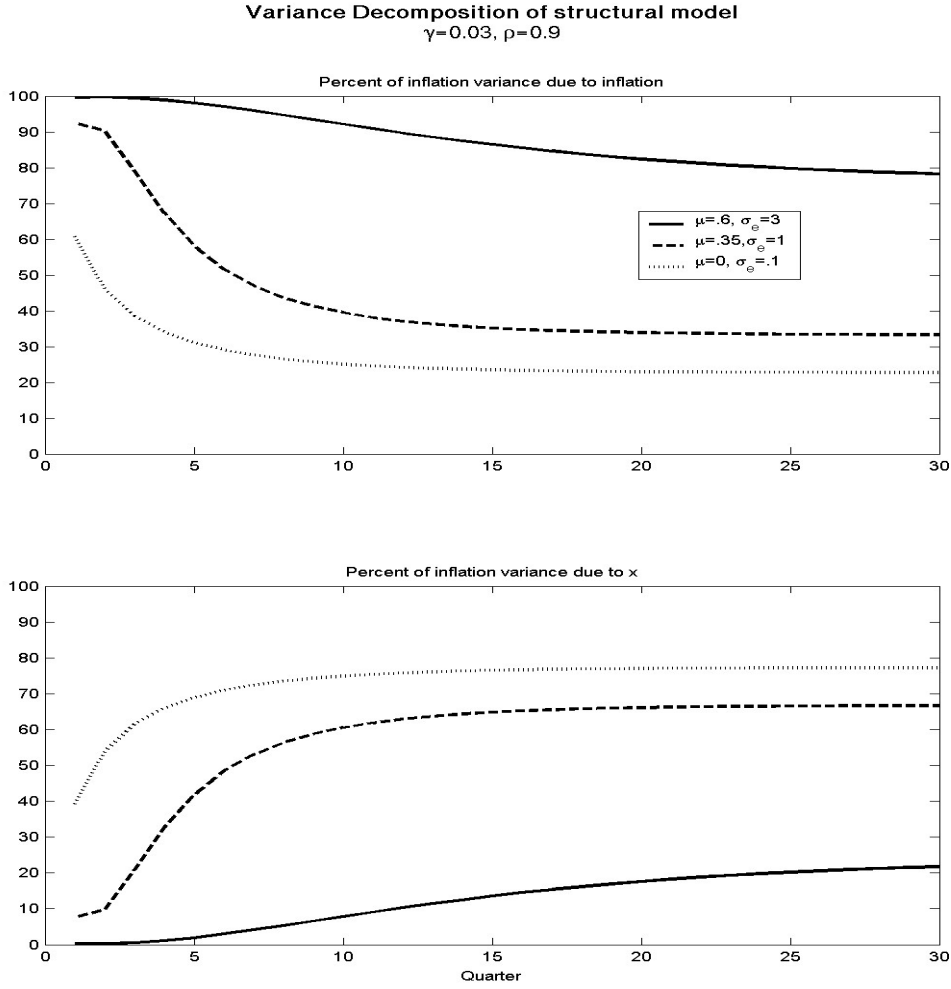
These variance decomposition comparisons provide a partial explanation for the difficulty in estimating a sizable and significant coefficient on the driving process. In the data, the fraction of inflation’s variance that may be attributed to either the output gap or marginal

³⁰ This result should come as no surprise, as the variance decompositions contain the same information as the vector autocovariance function, presented from a different perspective.

cost is quite small. The larger is the coefficient on the driving variable, the larger will be the implied fraction of inflation's variance attributable to the driving variable. For values of γ much larger than .01, the variance decomposition is strongly at odds with the decomposition found in the data-consistent VARs.

One competing hypothesis for the presence of a small coefficient on the driving process is that our proxies for real marginal cost are imperfect, and thus subject to measurement error. If the measurement error is classical (additive), then the coefficient on the driving process will be biased downward for the conventional reason. This will be true in both the structural model and the VAR (variance decomposition) exercises below. However, in the structural models examined below, to explain the order-of-magnitude shortfall of the estimated coefficient relative to one that would impart significant persistence from the driving process, the size of the measurement error would have to be quite large.

Figure 15



V. Conclusions

Finding a data-consistent, optimizing, rational expectations model of price-setting has been an important goal in macroeconomics for decades. An emerging consensus suggests that the New Keynesian Phillips curve, augmented by modest frictions of one flavor or another, is a good benchmark model for price-setting in dynamic stochastic general equilibrium (DSGE) models usable for macroeconomic analysis. When the driving process is assumed to be real marginal cost, the parameter on the driving process can be estimated with the correct sign, and, in principle, inflation should inherit considerable persistence from this variable.

This paper reaches conclusions that differ markedly from the prevailing wisdom. It suggests that:

- (1) Using conventional parameter estimates, inflation in the hybrid NKPC inherits little if any persistence from the driving process.
- (2) In part, this lack of inherited persistence derives from the presence of a large inflation shock whose variance is typically one to five times as large as the shock that perturbs the driving process.
- (3) The lack of inherited persistence also derives from a rather small estimated coefficient on the driving process.
- (4) The predominant source of inflation persistence in the NKPC is the lagged inflation term. The amount of persistence imparted by the lag is quite sensitive to the size of the lag, with significant differences in persistence implied by an increase in μ from 0.3 to 0.6.
- (5) As several papers have noted, the persistence of inflation appears to have declined in recent years. If that is true, this paper suggests that the reason for that decline in persistence is unlikely to be related to a decline in the persistence of the driving process.³¹ First, the standard candidates for the driving process have nearly the same persistence today as they did two decades ago. Second, to a first approximation, the NKPC as conventionally implemented does not allow important changes in the persistence of the driving process to affect the persistence of inflation.
- (6) Because monetary policy in the standard models acts through its effect on output and marginal cost, it becomes more difficult to attribute recent changes in inflation persistence to changes in monetary policy. This does not necessarily imply that monetary policy has had no such effects, but it does suggest that the current crop of models will have difficulty in attributing such changes to monetary policy.

³¹ There is considerable debate surrounding this observation, much of it methodological. Recent discussions at the European Central Bank's Inflation Persistence Network conference highlight the issues. See especially Session I at <http://www.ecb.de/events/conferences/html/inflationpersistence.en.html>.

These conclusions have other important implications for price modeling in DSGE models. They suggest that the optimizing foundations in the standard specifications are nearly unrelated to the dynamics observed in the data for inflation and real marginal cost. That is, lagged inflation is not a second-order add-on to the optimizing model, it *is* the model. One may motivate price-setting behavior from these optimizing foundations, but in practice, they tell us little about why inflation behaves the way it does.

The conclusions also imply that in order to understand inflation dynamics, we will need to identify the economic source of the large inflation shock in the specification. In turn, the findings in this paper imply either that this identified shock is itself highly autocorrelated, or that we require a micro-founded mechanism that generates substantial intrinsic persistence in inflation.

References

- Anderson, Gary and George R. Moore (1985): "A linear algebraic procedure for solving linear perfect foresight models," *Economics Letters* 17(3): 247-252.
- Bai, Jushan (1999): "Likelihood Ratio Tests for Multiple Structural Changes," *Journal of Econometrics* 91(2): 299-323.
- Berndt, Ernest, Bronwyn Hall, Robert Hall, and Jerry Hausman (1974): "Estimation and inference in nonlinear structural models." *Annals of Economic and Social Measurement* 3: 653-665.
- Calvo, Guillermo (1983): "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics* 12: 383-398.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." Forthcoming, *Journal of Political Economy*.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics* 46 (October): 281-313.
- Estrella, Arturo and Jeffrey Fuhrer (2002): "Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational Expectations Models," *American Economic Review* 92(4) : 1013-1028 .
- Fischer, Stanley (1977): "Long-Term Contracts, Rational Expectations and the Optimal Money Supply Rule," *Journal of Political Economy* (February): 191-206.
- Fuhrer, Jeffrey C. (2002): Discussion of "Monetary Policy for Inattentive Economies," Ball, Laurence, Mankiw, N.G., and Ricardo Reis. Unpublished comments from a joint Swiss National Bank/*Journal of Monetary Economics* conference, October 11-12. See link at <http://www.bos.frb.org/economic/econbios/fuhrer/jme111102.pdf> .
- ┐ and George R. Moore (1995): "Inflation Persistence," *Quarterly Journal of Economics* 440 (February):127-159.
- ┐ and George R. Moore (1995a): "Monetary Policy Trade-Offs and the Correlation Between Nominal Interest Rates and Real Output," *American Economic Review* 85, (March): 219-239.

┘ and Glenn D. Rudebusch (2004), "Estimating the Euler Equation for Output," *Journal of Monetary Economics* 51: 1133-1153.

┘ and Giovanni Olivei (2004), "Estimating Forward-Looking Euler Equations with GMM and Maximum Likelihood Estimators: An Optimal Instruments Approach," in Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley. Faust, J., Orphanides, A. and D. Reifschneider, eds., Board of Governors of the Federal Reserve System.

Galí, Jordi (2005): "Has the Inflation Process Changed? A Comment." Based on the discussion of Cecchetti and Debelle's paper "Has the Inflation Process Changed?" presented at the Third BIS Annual Conference, Brunnen, 18-19 June 2004.

Galí, Jordi (2003): "New Perspectives on Monetary Policy, Inflation, and the Business Cycle" in: *Advances in Economic Theory*, edited by: M. Dewatripont, L. Hansen, and S. Turnovsky, vol. III, 151-197, Cambridge University Press.

┘ Gertler, Mark and David Lopez-Salido (2001): "European Inflation Dynamics," *European Economic Review* 45,(7): 1237-1270.

┘ and Mark Gertler (1999): "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44(2): 195-222.

Goodfriend, Marvin, and Robert G. King (1997): "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual* 1997: 231-283.

Gordon, Robert J. (1985): "Understanding Inflation in the 1980s," *Brookings Papers on Economic Activity* 16(1): 263-99.

Gray, Jo Anna (1978): "On Indexation and Contract Length," *Journal of Political Economy* February.

Ireland, Peter N. (2004): "Technology Shocks in the New Keynesian Model," *Review of Economics and Statistics* 86(4): 923-936.

Levin, Andrew and Jeremy Piger, (2003): "Is Inflation Persistence Intrinsic in Industrial Economies?" ECB Working paper no. 334, May.

Margaret M. McConnell, and Gabriel Perez-Quiros (2000): "Output Fluctuations in the United States: What has Changed Since the Early 1980s?" *American Economic Review* December.

Lucas, Robert E., Jr. (1972): "Econometric Testing of the Natural Rate Hypothesis," O. Eckstein, *The Econometrics of Price Determination Conference*, Washington, D.C.: Board of Governors of the Federal Reserve System.

Mankiw, N. Gregory and Ricardo Reis (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117 (4).

O'Reilly, Gerard and Karl Whelan (2004): "Has Euro-Area Inflation Persistence Changed Over Time?," European Central Bank Working Paper Series No. 335, April, forthcoming in the *Review of Economics and Statistics*.

Phillips, A. W. (1958): "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica* 25 (November):283-299.

Roberts, John M. (1995): "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking* 27: 975-984.

Rotemberg, Julio (1983): "Aggregate Consequences of Fixed Costs of Price Adjustment," *American Economic Review* 73(3): 433-436.

____ and Michael Woodford (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* 1997: 297-346.

Rudd, Jeremy B. and Karl Whelan (2001): "New Tests of the New-Keynesian Phillips Curve" (June). Board of Governors of the Federal Reserve Working Paper No. 2001-30, forthcoming in the *Journal of Monetary Economics*.

Sbordone, Argia. 2002: "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49(2): 265-292.

Taylor, John B. (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88(1): 1-24.

Woodford, Michael (2003): Interest and Prices, Princeton University Press, Princeton, NJ.

Appendix A

Algebraic derivations

The Purely Forward-Looking Model

For the NKPC with $\mu=0$, the matrix A in equation (1.2) takes a particularly simple form. It is the coefficient matrix in the reduced-form solution to the model, which may be expressed as³²

$$\begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \begin{bmatrix} 0 & \frac{\rho\gamma}{(1-\rho\beta)} \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} \quad (6.1)$$

We can use this solution to substitute for $E_t\pi_{t+1}$ in equation (1.1) to obtain the matrix S , which has partitions S_0 (the contemporaneous block) and S_1 (the lagged block):

$$\begin{aligned} S_0 \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} &= S_1 \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} e_t \\ u_t \end{bmatrix} \\ \begin{bmatrix} 1 & \frac{-\gamma}{1-\rho\beta} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} &= \begin{bmatrix} 0 & 0 \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} e_t \\ u_t \end{bmatrix} \end{aligned} \quad (6.2)$$

Under the assumption that the covariance matrix of the errors is diagonal, and normalizing the variance of u to 1 and denoting the variance of e by σ_e^2 , we can derive the unconditional variance for the vector process as

$$V \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \begin{bmatrix} \frac{\gamma^2}{(1-\rho\beta)^2(1-\rho^2)} + \sigma_e^2 & \frac{\gamma}{(1-\rho\beta)(1-\rho^2)} \\ \frac{\gamma}{(1-\rho\beta)(1-\rho^2)} & \frac{1}{1-\rho^2} \end{bmatrix} \quad (6.3)$$

Then the autocovariances C_i and autocorrelations Γ_i may be derived from the recursive equations

³² See Anderson and Moore (1985) for a derivation of the solution coefficient matrix.

$$C_i = A C_{i-1}$$

$$\Gamma_i(j, k) = \frac{C_i(j, k)}{\sqrt{V(j)V(k)}} \quad (6.4)$$

where C_0 is initialized as V .

The Hybrid Model

Now the matrix A from the reduced-form perfect-foresight solution to the model may be expressed as

$$\begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \begin{bmatrix} \lambda_s & \frac{\gamma\rho}{(\beta-\mu)(\lambda_b-\rho)} \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} \quad (6.5)$$

where λ_s and λ_b are the “small” and “big” roots (or stable and explosive, with moduli less than and greater than one, respectively) of the transition matrix for the model. It is important to note that λ_b and λ_s depend only on β and μ , and are independent of the parameters governing the y_t process or its interaction with π_t .

$$\lambda_s = \frac{1 - \sqrt{(1 - 4\mu\beta + 4\mu^2)}}{2(\beta - \mu)}$$

$$\lambda_b = \frac{1 + \sqrt{(1 - 4\mu\beta + 4\mu^2)}}{2(\beta - \mu)} \quad (6.6)$$

$$\lambda_b \lambda_s = \frac{\mu}{\beta - \mu}$$

We can use this solution to substitute for $E_t \pi_{t+1}$ in equation (1.1) to obtain

$$\begin{aligned}
S_0 \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} &= S_1 \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} e_t \\ u_t \end{bmatrix} \\
\begin{bmatrix} \frac{\mu}{\lambda_s} & \frac{-\gamma\lambda_b}{\lambda_b - \rho} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} &= \begin{bmatrix} \mu & 0 \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} e_t \\ u_t \end{bmatrix}
\end{aligned} \tag{6.7}$$

We can derive the unconditional variance for the vector process as

$$V \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \begin{bmatrix} \frac{V_y \frac{\gamma^2}{(\rho - \lambda_b)^2} [1 + \frac{2\mu\rho}{(\beta - \mu)\lambda_b - \rho\mu}] + \frac{\sigma_e^2}{\lambda_b^2}}{(\beta - \mu)^2 - \frac{\mu^2}{\lambda_b^2}} & \frac{\gamma\lambda_b}{(\rho - \lambda_b)[\lambda_b(\mu - \beta) + \rho\mu]} V_y \\ \frac{\gamma\lambda_b}{(\rho - \lambda_b)[\lambda_b(\mu - \beta) + \rho\mu]} V_y & \frac{1}{1 - \rho^2} \end{bmatrix} \tag{6.8}$$

where V_y is $V(y)$, i.e. $\frac{1}{1 - \rho^2}$. Then the autocorrelations may be derived as above, using the transition matrix in equation (6.5).

The Forward-looking Model with Autocorrelated Errors

The model is augmented to include an “inflation shock” that follows an AR(1) process:

$$\begin{aligned}
\pi_t &= (\beta - \mu)E_t\pi_{t+1} + \mu\pi_{t-1} + \gamma y_t + e_t \\
y_t &= \rho y_{t-1} + u_t \\
e_t &= ae_{t-1} + \varepsilon_t
\end{aligned} \tag{6.9}$$

For this model, the key matrices are

$$A = \begin{bmatrix} 0 & \frac{\rho\gamma}{(1 - \rho\beta)} & \frac{a}{1 - a\beta} \\ 0 & \rho & 0 \\ 0 & 0 & a \end{bmatrix}; S_0^{-1} = \begin{bmatrix} 1 & \frac{\gamma}{1 - \rho\beta} & \frac{1}{1 - a\beta} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{6.10}$$

and the unconditional variance for the vector process is

$$V = \begin{bmatrix} \frac{\gamma^2 \sigma_u^2}{(1-\rho\beta)^2(1-\rho)^2} + \frac{\sigma_e^2}{(1-a\beta)^2(1-a)^2} & \frac{\gamma\sigma_u^2}{(1-\rho\beta)(1-\rho)^2} & \frac{\sigma_e^2}{(1-a\beta)(1-a)^2} \\ \frac{\gamma\sigma_u^2}{(1-\rho\beta)(1-\rho)^2} & \frac{\sigma_u^2}{(1-\rho)^2} & 0 \\ \frac{\sigma_e^2}{(1-a\beta)(1-a)^2} & 0 & \frac{\sigma_e^2}{(1-a)^2} \end{bmatrix} \quad (6.11)$$

and from these one can derive the autocorrelation function for inflation.³³ The first autocorrelation in this case is

$$\Gamma_1 = \frac{\rho(a^2-1)(1-a\beta)^2\gamma^2\sigma_u^2 + a(\rho^2-1)(1-\rho\beta)^2\sigma_e^2}{(a^2-1)(1-a\beta)^2\gamma^2\sigma_u^2 + (\rho^2-1)(1-\rho\beta)^2\sigma_e^2} = \frac{\rho\delta_1\gamma^2\sigma_u^2 + a\delta_2\sigma_e^2}{\delta_1\gamma^2\sigma_u^2 + \delta_2\sigma_e^2} \quad (6.12)$$

As the text in Section II suggests, the autocorrelations are dominated by a , the autocorrelation parameter on the shock. In essence, this version of the model holds the same implications as the hybrid model: Here, the correlation of the shock term does all the work in the model, whereas in the hybrid model, the lagged inflation term plays the same role.

Data Appendix

Inflation: 400 times the log change in the GDP chain-type price index.

Output gap: 100 times the log difference between chain-weighted real GDP and the Congressional Budget Office's estimate of potential GDP.

Real marginal cost: proxied by real unit labor costs, i.e 100 times nominal unit labor costs (log of nonfarm compensation less the log of nonfarm output per hour) less the log of the implicit price deflator for the nonfarm business sector.

Relative price of oil: the log of West Texas intermediate oil price per barrel less the log of the GDP chain-type price index.

Relative price of imports: the log of the chain-type import price index less the log of the GDP price index.

³³ The roots in the model are particularly simple: $[\frac{1}{\beta}, \rho, a]$.