Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 54 (2007) 1311–1339
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Inflation persistence: Alternative interpretations and policy implications ☆

Argia M. Sbordone

Federal Reserve Bank of New York, USA

Received 27 April 2007; received in revised form 24 May 2007; accepted 6 June 2007
Available online 19 June 2007

Abstract

In this paper, I consider the policy implications of two alternative structural interpretations of observed inflation persistence, which correspond to two alternative specifications of the new Keynesian Phillips curve (NKPC). The first specification allows for some degree of intrinsic persistence by way of a lagged inflation term in the NKPC. The second is a purely forward-looking model, in which expectations farther into the future matter and coefficients are time-varying. In this specification, most of the observed inflation persistence is attributed to fluctuations in the underlying inflation trend, which are a consequence of monetary policy rather than a structural feature of the economy. With a simple quantitative exercise, I illustrate the consequences of implementing monetary policy, assuming a degree of intrinsic persistence that differs from the true one. The results suggest that the costs of implementing a stabilization policy when the policymaker overestimates the degree of intrinsic persistence are potentially higher than the costs of ignoring actual structural persistence; the result is more clear-cut when the policymaker minimizes a welfare-based loss function.

Keywords: Inflation persistence; Monetary policy

☆ Prepared for the November 2006 Carnegie–Rochester Conference on Monetary Policy held in Pittsburgh on November 10–11, 2006. I wish to thank Benneth McCallum and Ken West for their comments, Mike Woodford for numerous discussions, and Krishna Rao for excellent research assistance. I also wish to thank Tim Cogley, my collaborator on the prior work that I build upon here. The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.
E-mail address: argia.sbordone@ny.frb.org

0304-3932/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.jmoneco.2007.06.007

1. Introduction

Inflation persistence is defined as the time that it takes for an inflation shock to dissipate. It is extremely important that central banks, which are responsible for stabilizing inflation at low levels, fully understand the nature of this process. Univariate analyses, based on the size of the highest autoregressive root, and multivariate analyses, based on impulse response functions, each show that U.S. inflation is a highly persistent process. However, less agreement exists on whether inflation persistence is an inherent characteristic of the economy, or if it instead depends on the specific historical sample considered. Moreover, if the degree of persistence is sample-specific, then what are the factors that effect changes in the degree of persistence over time?

Recent work by Cogley and Sargent (2006) and Stock and Watson (2007) favor the view that there has been substantial variation in U.S. inflation persistence over time associated with changes in the monetary policy regime. Meanwhile, Pivetta and Reis (2007) argue that this kind of evidence is not statistically significant. Benati (2006) analyzes the evolution of inflation persistence across countries and historical monetary regimes and observes that the degree of inflation persistence appears to have varied significantly and to have been lower in periods in which there was a clearly defined nominal anchor.

To understand the source of inflation persistence, we require structural models. In several recent macro models used for policy analysis, inflation dynamics are derived from a discrete-time version of the Calvo price-setting model. In its baseline formulation, this model is purely forward-looking: inflation depends on real marginal costs and expected future inflation. To accommodate the observed persistence in inflation data, two main variants of the model have been proposed in the literature. Both variants require some ad hoc assumptions about the price-setting process to generate the dependence of inflation on past values of inflation. The first variant, with a simple modification of the Calvo model, was introduced by Gali and Gertler (1999). They assumed that a proportion of the firms randomly assigned to reoptimize their prices follow a rule-of-thumb: their prices are a weighted average of the optimal prices set in the previous period plus an adjustment for expected inflation, which is based on lagged inflation. An alternative modification, obtained by assuming that the firms not assigned to reoptimize their prices index their prices to the aggregate inflation of the previous period, was later introduced by Christiano et al. (2005), and then further modified to allow for only partial indexation.[1]

The reduced forms of these two variants of the Calvo model are similar in that both generate a backward-looking component in the equation. In the first case, the weight of this component depends on the proportion of rule-of-thumb firms, and in the second, it depends on the indexation coefficient.[2] However, using indexation as a modeling strategy is a less appealing method because doing so implies that prices are revised at every point in time, which contradicts empirical evidence that some prices are fixed for a certain amount of time (the reason why models with nominal rigidities were developed). Single-equation estimates of these augmented models identify a small but statistically significant coefficient on past inflation and find that the introduction of an inflation lag helps to fit the data

[1]For single-equation estimates of the Calvo model with indexation, see Eichenbaum and Fisher (2007) and Sbordone (2005), among others.

[2]In the specifications that assume a unit indexation coefficient, the equation is estimated as the rate of inflation growth, rather than the level of inflation.

better than do purely forward-looking models. The coefficient on past inflation is typically estimated to be about 0.2–0.3, depending on the other specifications in the model.[3]

In two recent papers (Cogley and Sbordone, 2005, 2006), we estimate a new Keynesian Phillips curve (NKPC), taking into account the existence of a slow-moving inflation trend. We derive a variant empirical version of the Calvo model by log-linearizing the model around a time-varying inflation trend and obtain a NKPC with more forward-looking dynamics than the baseline Calvo model. Expectations of inflation further into the future matter, and the coefficients of the NKPC depend on the trend in inflation: the sensitivity of inflation to marginal cost decreases for higher levels of trend inflation, while the weight on future expectations increases. Our estimates of this specification of the NKPC favor a purely forward-looking model for the inflation gap, which we define as deviation of inflation from the time-varying trend. The absence of a significant intrinsic persistence in the inflation gap implies that the persistence of overall inflation is driven by the persistence of its underlying trend—which is a consequence of monetary policy, rather than a structural feature of the economy.

We were not the first to estimate a long-run moving trend in inflation. Our analysis explores the implications for a structural analysis of the results obtained via a reduced-form analysis by Cogley and Sargent (2005): they applied models with time-varying coefficients to explore changes over time in the persistence and volatility of inflation, unemployment, and interest rates.

Indeed, there are now a number of small-scale general equilibrium models for policy analysis that depict inflation as evolving around a long-run trend; this trend in turn is identified with the inflation objective of the policymaker. These models, unlike ours, assume the existence of intrinsic inflation persistence by introducing both indexation to past inflation and a partial indexation to trend inflation to obtain an empirical form of a NKPC of the standard type, in which inflation depends on an inflation lag, the expected future value of inflation, and current marginal costs (for example, Smets and Wouters, 2003, 2005; Ireland, forthcoming). Although the dependent variable of the NKPC in these models is the deviation of inflation from trend, because of the assumed indexation the coefficients of such NKPC do not depend upon the level of trend inflation, as they do in our model. The indexation parameter to past inflation estimated in these models varies, and can be as high as 0.5, as seen in the Smets and Wouters model.[4]

The absence of intrinsic persistence as well as the attribution of inflation persistence to persistent movements in trend inflation square with several other results in the literature. Altissimo et al. (2006) summarize research conducted within the Eurosystem inflation persistence network. They find that in aggregate data inflation persistence is very high for samples spanning different decades, but it falls dramatically when one allows for time variation in the mean level of inflation. They further find that the timing of the breaks in mean inflation correspond to observed breaks in the monetary policy regime.

In light of this evidence and considering that policymakers most often base policy decisions on models that postulate a substantial degree of inflation persistence, I investigate the

[3]See Gali and Gertler (1999), Gali et al. (2001, 2005), and Sbordone (2005, 2006).

[4]There are other lines of research that show reasons for spurious estimates of a backward-looking component. For example, Milani (2007) estimates a model with shifting inflation trends and reports parameters on backward-looking terms close to zero. Kozicki and Tinsley (2002) find that shifts in the long-run inflation anchor of agents' expectations explain most, but not all, of the historical inflation persistence in the United States and Canada.

policy implications of assuming alternative structural interpretations of observed inflation persistence. In particular, I ask the following questions: What are the costs of accommodating inflation when inflation persistence is not intrinsic? Is there risk that an incorrect policy response may translate temporary shocks to inflation into more persistent fluctuations?[5]

To address these issues, I focus on two alternative specifications of inflation dynamics, which correspond to the two alternative interpretations of inflation persistence discussed: one is the rule-of-thumb model introduced by Gali and Gertler (1999) and the other is the model with varying trend inflation estimated by Cogley and Sbordone (2006).

In the first exercise, I incorporate the inflation dynamics of the Cogley–Sbordone model in a very stylized model of the economy and ask whether it is possible that intrinsic persistence is spuriously detected in the data. Specifically, I consider whether one would estimate a statistically significant coefficient on lagged inflation when fitting the NKPC to data generated from a calibrated economy in which there is no intrinsic persistence.

I then consider the two alternative models to conduct a quantitative evaluation of optimal monetary policy. In particular, I consider the implications of implementing optimal stabilization policy under assumptions about inflation persistence that differ from what the true model implies. To refine this exercise, I also characterize optimal stabilization policy when the policymaker accounts for uncertainty about the model of the economy.

The rest of the paper is organized as follows. In the next section, I discuss the characteristics of the two alternative models of inflation persistence, discussing in more detail the model with trend inflation, as it is less known in the literature. In Section 3, I discuss whether it is possible to misconstrue intrinsic persistence from data generated by an economy in which the NKPC has a time-varying inflation trend. In Section 4, I present an analysis of the optimal response of the economy to cost-push shocks: I first consider the case of optimization based on an ad hoc loss function and then the case of a welfare-based optimal policy. Section 5 analyzes optimal stabilization policy under model uncertainty, and Section 6 concludes.

2. Inflation persistence: alternative interpretations

In the *NKPC* derived from the standard discrete-time version of the Calvo model with random intervals between price changes, inflation depends upon current marginal costs s_t and expected future inflation[6]:

$$\pi_t = \zeta s_t + \beta \mathrm{E}_t \pi_{t+1}. \tag{1}$$

The coefficient of marginal cost ζ is a non-linear combination of structural parameters whose expression depends on the specific market structure assumed. At the minimum, it includes the probability of not changing prices, which I will denote throughout this paper by α, and the discount factor β: $\zeta = (1-\alpha)(1-\alpha\beta)/\alpha$. In models with some form of strategic complementarity the coefficient may also depend upon the elasticity of substitution among differentiated goods, and the elasticity of marginal cost to firms'

[5] This is the interpretation that Ireland (forthcoming) gives to his results.

[6] Lowercase letters denote logs. $\pi_t = \ln \Pi_t$, where Π_t denotes the gross inflation rate: $\Pi_t = P_t/P_{t-1}$. s_t is the log of real marginal cost.

output, parameters that I denote, respectively, by θ and ω. For example $\zeta = (1-\alpha)(1-\alpha\beta)/\alpha(1+\theta\omega)$.[7] This richer specification decouples the degree of nominal rigidity from estimates of the coefficient of marginal cost.[8]

The variant of the Calvo model introduced by Gali and Gertler (1999) implies that the model includes a lagged inflation term, to become

$$\pi_t = \widetilde{\zeta} s_t + \gamma_{\mathrm{f}} \mathrm{E}_t \pi_{t+1} + \gamma_{\mathrm{b}} \pi_{t-1} + u_t. \tag{2}$$

A lagged inflation term now appears in the equation because the authors assume that a fraction χ of 'rule-of-thumb' firms set prices as a weighted average of the optimal prices set in the previous period plus an adjustment for expected inflation, which is based on lagged inflation. Gali and Gertler (1999) do not allow for strategic complementarities: when the fraction $\chi \to 0$, the equation is identical to (1). The coefficients of the forward and backward-looking terms depend upon the fraction χ as well, and their sum is approximately equal to 1 (it is exactly 1 for $\beta = 1$), making the equation similar to other hybrid Phillips curve formulations in the literature (e.g. Fuhrer and Moore, 1995). Models with partial indexation to past inflation can also be written in the form of Eq. (2), where the coefficient of past inflation depends upon the degree of indexation, and again the coefficients of the forward and backward-looking terms sum to one when $\beta = 1$.

The *NKPC* in either form (1) or (2) is derived as a log-linear approximation to the exact non-linear inflation dynamics described by the Calvo model, where the log-linearization is taken around a steady state with zero inflation. This conventional approximation is useful for normative studies, since inflation should be close to zero under an optimal policy rule. But in the historical periods covered by empirical analyses—typically some subsample of the post-WWII period, inflation is often substantially above zero. This raises a question as to how accurate the log-linear approximation used in empirical work may be. Moreover, because the degree to which inflation exceeds zero has been subject to fairly persistent fluctuations, the approximation error may substantially affect the estimated degree of intrinsic persistence.

To address this problem, the variant of the Calvo model estimated in Cogley and Sbordone (2006) takes the following form[9]:

$$\widehat{\pi}_t = \widetilde{\varrho}_t \widehat{\pi}_{t-1} + \zeta_t \widehat{s}_t + b_{1t} \mathrm{E}_t \widehat{\pi}_{t+1} + b_{2t} \sum_{j=2}^{\infty} \varphi_{1t}^{j-1} \mathrm{E}_t \widehat{\pi}_{t+j} + u_t. \tag{3}$$

Unlike the previous equations,[10] this specification is derived by log-linearizing the non-linear equilibrium conditions of the Calvo model around a steady state with a time-varying trend inflation. The hat variables denote log-deviation of inflation from trend ($\widehat{\pi}_t = \pi_t - \overline{\pi}_t$), and marginal cost from trend ($\widehat{s}_t = s_t - \overline{s}_t$), where $\overline{\pi}_t$ and $\overline{s}_t$ indicate trend

[7]The *NKPC* model estimated by Cogley–Sbordone (2005, 2006) includes this form of strategic complementarities. For an analysis of other specifications of strategic complementarities, and their implications for monetary policy, see Levin et al. (2006).

[8]The component $1/(1+\theta\omega)$ in the coefficient ζ represents a measure of real rigidities.

[9]This form of the model is correct under the assumption of log-utility. Otherwise it includes some further terms in the expected value of output growth and the discount factor. We considered the restricted form (4) because of evidence that the coefficient of the extra terms were empirically insignificant.

[10]The error term is also more complex than in the previous equations, because it includes trend inflation innovations.

variables.[11] The coefficients are indexed by t because they depend upon trend inflation $\overline{\pi}_t$, as does the value of $\overline{s}_t$; and they also depend upon the primitives of the Calvo model, the probability of changing prices and the elasticity of demand, which are the parameters that we estimate. Our estimation procedure is based on the moment conditions derived by enforcing the restrictions that the model places on a *VAR* model for inflation and unit labor costs (the last variable proxies the unobservable marginal cost of the model, as in most variants of empirical *NKPCs*). The reduced form *VAR* is a model with time-varying coefficients and stochastic volatility, which we estimate jointly with the parameters of the Calvo model. We then use the estimated *VAR* model to compute the implied inflation trend.

The coefficient on lagged inflation $\widetilde{\varrho}_t$ depends on the indexation to past inflation, a feature that we include to allow for possible intrinsic persistence: its value is zero when the indexation parameter is zero. Our preferred specification, in fact, excludes lagged inflation, since we find that the indexation parameter is very small, and that a model without indexation is statistically preferred. We have then

$$\widehat{\pi}_t = \zeta_t \widehat{s}_t + b_{1t} \mathrm{E}_t \widehat{\pi}_{t+1} + b_{2t} \sum_{j=2}^{\infty} \varphi_{1t}^{j-1} \mathrm{E}_t \widehat{\pi}_{t+j} + u_t. \tag{4}$$

We conjecture that the contrast between our result and those that find, in the same data, a statistically significant role for lagged inflation arises because the latter may proxy for omitted terms in the *NKPC*. These are the additional forward-looking terms of our more precise approximation to the model, which obtains when trend inflation is non-zero.

We conclude that inflation deviations from trend do not show *intrinsic* persistence, and that the persistence of overall inflation is driven by persistent fluctuations in its underlying trend. In this interpretation inflation persistence, rather than a structural feature of the economy, appears to be a consequence of the way monetary policy has been conducted.

In this respect, although the source of persistence can be categorized as 'exogenous', it is different from letting serially correlated shocks capture inflation persistence, as some of the recent literature has it. This is because trend inflation, unlike the shocks, is ultimately under the control of the policy maker (as a general equilibrium framework would make clear), and therefore need not be taken as given when the policymaker sets her policy. One can argue that movements in trend inflation should in principle be endogenous[12] and perhaps rationalized as optimal behavior of the central bank. But even when endogenized, this source of inflation persistence is different from one 'intrinsic' to the price-setting process.

It is worth to report at this point two results from Cogley–Sbordone (2006). Fig. 1 graphs estimated inflation trend, actual inflation and average inflation for the period 1960–2003, all expressed at annual rates. As the figure shows, the trend is a quite persistent process, hovering around an annual average of 2.5% up to the early 1970s and after the mid-1980s, but rising to slightly above 7% in the late 1970s.

[11]For more details on the derivation of this equation, and the estimation procedure, see Cogley–Sbordone (2006).

[12]To my knowledge, the only author that attempts to endogenize the policymaker target is Ireland (forthcoming), who maintains a unit root in the process, but allows it to respond to other shocks in the model. He finds, however, weak statistical evidence of such an endogenous response.

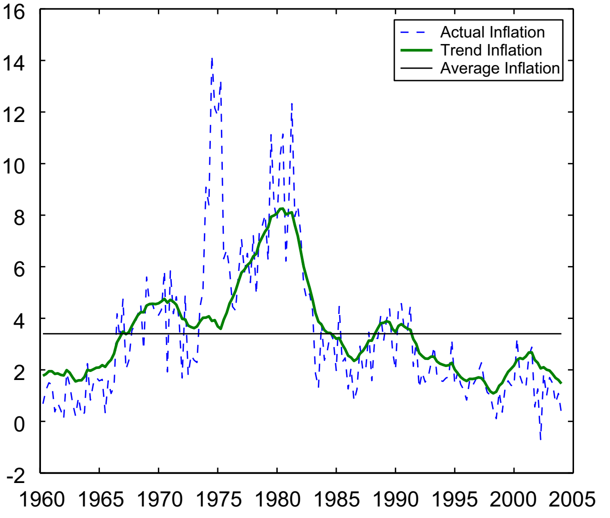

Fig. 1. Inflation, trend inflation and average inflation.

Table 1
Autocorrelation of the inflation gap

	1960–2003	1960–1983	1984–2003
Mean-based inflation gap	0.835	0.819	0.618
Trend-based inflation gap	0.632	0.664	−0.038

The figure suggests that the properties of the inflation gap depend to a large extent upon its measurement. A gap defined as deviation from a constant mean has a great deal more persistence than a gap measured as deviation from a time-varying trend. This is also shown in Table 1 (again taken from Cogley and Sbordone, 2006), which reports the serial correlation of two measures of inflation gap. The first assumes a constant trend equal to the sample average of inflation over the period, while the second, labeled trend-based inflation gap, is computed as deviation of inflation from the estimated trend, which we obtain, as described, by imposing the restrictions of the forward-looking model (4).

The table clearly suggests that the inflation gap measured as deviation from trend has substantially less persistence than the gap measured as deviation from the mean. The difference is particularly striking over the second sub-sample, when the trend-based gap is close to white noise. Both the figure and the table illustrate how the need for introducing a backward-looking component in structural models of inflation dynamics may derive from an inadequate measure of the inflation gap, which leads to overemphasize the persistence that such models are asked to explain.

3. Can intrinsic persistence be spurious?

The first question that I ask is whether spurious intrinsic persistence may be estimated in *NKPC* models that are log-linearized around a zero-inflation steady state. If the actual economy is instead characterized by a drifting inflation trend, a more appropriate specification of inflation dynamics is that of Eq. (4), supplemented by a law of motion for the inflation trend. To evaluate the extent to which such a spurious persistence may arise, I construct a sample economy, where the *NKPC* has the form (4), and the underlying inflation trend has the properties of the one estimated in Cogley–Sbordone (2006). I complete this economy with a very simple specification of the demand side, a policy rule of a kind common in the literature, and a number of shocks which generate temporary departures of inflation from trend. I then ask whether estimating, with a standard econometric technique, an equation of the kind (2) on data generated by this economy one would make a correct assessment of the degree of intrinsic inflation persistence.

For simulating the sample economy it is convenient to use a recursive representation for Eq. (4). This is obtained by defining an auxiliary variable $\widehat{d}_t$ that represents the further forward-looking terms in the equation.[13] The model economy is then composed of the following equations:

$$\begin{aligned}
\widehat{\pi}_t &= \zeta_t \widehat{s}_t + \phi_t \mathrm{E}_t \widehat{\pi}_{t+1} + \gamma_t \widehat{d}_t, \\
\widehat{d}_t &= g_{1t} \mathrm{E}_t \widehat{\pi}_{t+1} + g_{2t} \mathrm{E}_t \widehat{d}_{t+1}, \\
\widehat{s}_t &= \widetilde{\omega} \widehat{y}_t + \mu_t, \\
\widehat{y}_t &= \mathrm{E}_t \widehat{y}_{t+1} - \sigma(i_t - \mathrm{E}_t \pi_{t+1} - r_t^n), \\
i_t &= i_{t-1} + \phi_\pi \widehat{\pi}_t + \phi_y (\widehat{y}_t - \widehat{y}_{t-1}), \\
\overline{\pi}_t &= \overline{\pi}_{t-1} + \varepsilon_{\pi t}.
\end{aligned} \tag{5}$$

The first two equations represent the *NKPC* with trend inflation discussed before.[14] Note that when trend inflation is zero ($\overline{\Pi}_t = 1$), the coefficient of $\widehat{d}_t$ is zero, and the coefficients ζ_t and ϕ_t become time invariant, reducing the equation to the standard formulation (1). The third equation describes the relation between marginal cost and output,[15] and the fourth has the form of an intertemporal *IS* equation describing the evolution of output. A difference rule for the evolution of the interest rate[16] and a

[13]A detailed derivation of the recursive representation for a more general case is in an appendix available from the author. Parts are also reported in the appendix of Cogley–Sbordone (2006). A similar derivation, for the case of a constant inflation trend, can be found in Ascari and Ropele (2006).

[14]The coefficients are defined as follows:

$$\zeta_t = \frac{1 - \alpha\beta\overline{\Pi}_t^{\theta(1+\omega)}}{1+\theta\omega} \frac{1 - \alpha\overline{\Pi}_t^{\theta-1}}{\alpha\overline{\Pi}_t^{\theta-1}}, \quad \phi_t = \beta\overline{\Pi}_t^{1+\theta\omega}, \quad \gamma_t = \frac{1}{1+\theta\omega} \frac{1 - \alpha\overline{\Pi}_t^{\theta-1}}{\alpha\overline{\Pi}_t^{\theta-1}} (\overline{\Pi}_t^{1+\theta\omega} - 1),$$

$$g_{2t} = \alpha\beta\overline{\Pi}_t^{\theta-1} \quad \text{and} \quad g_{1t} = (\theta - 1) g_{2t},$$

where $\overline{\Pi}_t$ denotes the gross trend inflation rate (see Cogley–Sbordone, 2006 for further details).

[15]Although the coefficient $\widetilde{\omega}$ may vary with trend inflation as well, for simplicity I assume it constant in this exercise.

[16]Orphanides and Williams (2002) advocate this kind of difference rule for monetary policy (where the short-term nominal interest rate is changed from its existing level in response to inflation and changes in economic activity) to deal with the uncertainty in estimates of the natural rate. Gali (2003) derives a similar interest rate process from the money market equilibrium, where money supply follows a unit root process. The coefficient of inflation is then interpreted as the inverse of the interest rate (semi) elasticity, and that on output is the ratio of output and interest elasticities of money demand. The choice of this form of interest rate rule assures a determinate equilibrium for all the values of trend inflation considered in the model. A standard Taylor rule results instead in indeterminacy of equilibrium for high values of trend inflation, as discussed by Ascari (2004).

stochastic process for trend inflation close the model. Hat variables denote, as previously, deviations from steady state. The steady state of this economy is characterized by slowly evolving trend inflation and trend labor share (the proxy for the theoretical marginal cost). The values for trends $\overline{\pi}_t$ and $\overline{s}_t$ are those of the estimated series in Cogley–Sbordone (2006), as are the values calibrated for underlying parameters of the Phillips curve, the probability of not changing prices α and the elasticity of demand θ which are set at their respective posterior means ($\alpha = .55$ and $\theta = 12.3$). The coefficients of the *NKPC* are non-linear functions of these and other calibrated parameters, and are also function of the inflation trend: they are computed accordingly. The dynamics of the economy is driven by shocks to the marginal cost, μ_t, and by natural rate shocks r_t^n. Both disturbances are assumed to follow autoregressive processes, with serial correlation, respectively, of $\varrho_r = .8$ and $\varrho_\mu = .2$.[17] The other parameters are calibrated to values used elsewhere in the literature: $\sigma = 6.25$, the value estimated in Rotemberg–Woodford (1997), and $\phi_\pi = \phi_y = .03$.[18]

I use this economy to create 100 samples of length equal to the inflation series used in Cogley–Sbordone (2006). Initial values are chosen to represent the economy at the beginning of the sample, which is 1960:Q1, and the economy is simulated forward for 176 periods, the length of the period for which trend inflation was estimated. The coefficients of the *NKPC* depend upon $\overline{\Pi}_t$, and vary with it. I take as initial value for the trend the estimated value of trend inflation for the first period ($\overline{\Pi}_0 = \overline{\Pi}_{60:q1}$).

Collecting all the parameters in a vector ψ, we have that at any time t, given $\psi_t = (\psi, \overline{\Pi}_t)$ there is a unique solution to the dynamic system that describes the evolution of $\widehat{\pi}_t, i_t$ and $\widehat{y}_t$, as function of state variables and shocks. From this solution I compute the one step forward value for the endogenous variables as the next realization, and repeat the same steps for the number of desired observations.

I therefore obtain series for inflation, marginal cost, output and interest rate generated by a trend inflation economy. On these series I then estimate a constant parameter specification of the type introduced by Gali and Gertler (1999), and since then successfully estimated for various countries in different time periods.

Specifically, I fit to the simulated series a hybrid model of the form (2), where π_t is a mean-based inflation gap. Denoting by Z_{t-1} a set of variables dated $t-1$ and earlier, the assumption of rational expectations and the assumption that the error term u_t is an i.i.d. process imply the following orthogonality conditions:

$$\mathrm{E}_{t-1}\{(\pi_t - \gamma_\mathrm{b}\pi_{t-1} + \gamma_\mathrm{f}\pi_{t+1} + \zeta s_t)Z_{t-1}\} = 0.$$

I exploit these conditions for estimation, using a parsimonious set of instruments, which include two lags of real marginal cost, output, interest rate and inflation. As Gali et al. (2005), I consider both an unconstrained estimate, and one where the coefficients of future and past inflation are constrained to sum to 1.[19] The table reports two sets of results, one for the whole sample I created, the other, for comparison, for the shorter sample 1960:Q1–1997:Q4, as in the original work of Gali and Gertler (1999). For each coefficient,

[17] I assume that the natural rate follows a process $r_t^n = (1-\varrho_r)r + \varrho_r r_{t-1}^n + \varepsilon_{r,t}$ and calibrate the mean r to the average value of the real interest rate in the sample. The variance of the shock μ_t is calibrated from the variance of the disturbance to the *NKPC* estimated in Cogley–Sbordone (2006).

[18] The value for σ is reported in Table 5.1 of Woodford (2003, p. 341); ϕ_π and ϕ_y correspond to the case of a money-growth target, under the assumption of a money demand function with an income elasticity of 1 and an interest semi-elasticity of 7, as in the semi-logarithmic model proposed by Lucas (2000).

[19] This is imposed by estimating the model as $\pi_t - \pi_{t-1} = \gamma_\mathrm{f}(\pi_{t+1} - \pi_{t-1}) + \zeta s_t$.

Table 2
Estimate of a hybrid standard *NKPC*

	ζ	γ_b	γ_f	*J*-stat: *p*-value $<.05$ (%)
Full sample				
Unrestricted	[.017,.052]	[.123,.226]	[.536,.793]	29
$\gamma_b + \gamma_f = 1$	[.010,.016]	–	[.772,.942]	39
1960:1–1997:4				
Unrestricted	[.016,.053]	[.105,.219]	[.535,.801]	37
$\gamma_b + \gamma_f = 1$	[.007,.015]	–	[.761,.929]	15

I report in square brackets a 90% confidence interval; for the J-statistic, I report instead the percentage of J-statistics with a *p*-value less than 5%.

The coefficient of interest here is the measure of intrinsic persistence γ_b. Although the true model has no intrinsic persistence, the estimation uncovers a positive coefficient: with 90% confidence we would not reject the hypothesis that there is a significant source of intrinsic persistence in inflation dynamics. The intuition for this spurious result is that past inflation is correlated with future inflation terms, an omitted variable in the empirical specification, creating upward bias in the coefficient of past inflation. This correlation is itself due to the serial correlation of the generated series, driven by the high persistence in the inflation trend. It is also interesting to note that with 95% confidence the J-statistics (reported in the last column) would fail to reject the cross-equation restrictions of the model in about 60–70% of the cases, on average across the various specifications (Table 2).[20]

4. Policy implications: optimal response to cost-push shocks

I now turn to consider the implications for monetary policy of the two structural interpretations of inflation persistence discussed in Section 3. One is the hybrid *NKPC*, and the other is the purely forward-looking model that takes into account the dependence of *NKPC* coefficients on inflation trend, now identified with the inflation objective in the policy loss function.

I conduct a quantitative exercise to evaluate the optimal stabilization policy in response to small cost-push shocks. I consider first the case where the objective of the policymaker is represented by an ad hoc quadratic loss function, and then the case of a welfare-based loss function. In both cases I focus on the optimal response to cost push shocks when the policymaker believes, alternatively, in one of the two different models. Assuming a certain form of the policy rule, I evaluate the equilibrium outcome of the optimal stabilization policy in terms of the implied paths of inflation, output and interest rate. The same ad hoc loss function is assumed for both models, but the welfare-based loss function is necessarily different for each model.

Since my objective in this exercise is to discuss the role played by the two different assumptions about inflation persistence as represented by the different forms of the *NKPC*,

[20] One should however note that the failure to reject may be due to the low power of the J-test in small samples.

I construct two model economies that differ only in their supply side assumptions. The rest of the economy is, in both cases, described by an intertemporal IS equation and a simple Taylor rule, where the interest rate responds to output gaps, and to deviations of inflation from target. In addition, the intercept of the Taylor rule is a function of the cost-push shocks, in a way that will be explained below.

To evaluate the cost of implementing monetary policy under wrong assumptions about the model economy, I compare the equilibrium paths of output, inflation and interest rate under the hypothesis that the true model of the economy is the same as that used by the policymaker with the paths that develop when the policymaker uses instead a different model. In this case the equilibrium paths of output and inflation will be different from the optimal paths, and the cost of using a 'wrong' model for policy—specifically of over or underestimating the degree of intrinsic persistence—can be computed by comparing the present value of the cumulative loss in the two cases. The specific metric for comparison is discussed later.

I calibrate the *NKPC* curves in the two models to the parameter values estimated in the single equation models discussed previously: Cogley and Sbordone (2006) model for the forward-looking model with time-varying inflation trend, and Gali et al. (2005) model for the rule-of-thumb model.

This analysis builds on existing analyses in the literature. Steinsson (2003) analyzes optimal monetary policy in the presence of inflation inertia in the context of a generalized rule-of-thumb model. He compares optimal responses of output and inflation to a cost push shock under different assumptions about the policymaker's loss function, whether ad hoc or welfare-based, and different assumptions about the degree of commitment. He considers i.i.d. cost push shocks, derived as a combination of shocks to the elasticity of demand and shocks to the tax code.[21] In this analysis I use the original specification of Gali and Gertler (1999), and consider both the case of i.i.d. shocks and the case of mildly serially correlated shocks.

Ascari and Ropele (2006) analyze optimal monetary policy in a purely forward-looking model with non-zero trend inflation. They find that in the case of positive trend inflation the optimal response to cost-push shocks is an aggressive deflation, and a persistent adjustment of the output gap, engineered through an increase in the interest rate; the higher is the level of trend inflation, the smaller is the optimal response of the output gap and the interest rate.

The key factors that shape the response to a cost push shock in these calibrated models are the parameters of the loss function, namely the relative weight on the output gap, the persistence of the shock, and the degree of intrinsic inertia. Without inertia, the response of inflation to i.i.d. shocks has its maximum at impact, and may be followed by a short period of deflation; output also has maximum decline at impact. The slope of the Phillips curve matters as well: higher sensitivity of inflation to marginal cost, for a given proportionality of the latter to output gap, reduces the 'optimal' response of output consistent with the inflation response.

[21] Among other results Steinsson finds that in the case of a traditional loss function an increase in the backward-looking component of the *NKPC* implies that the optimal response of inflation is lower in the period of the shock and more persistent, and that increasing the backward-looking component reduces the impact response of the output gap, and makes it less persistent.

4.1. Ad hoc loss function

The first economy I consider is described by the following three equations:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \pi_{t+1} + \kappa x_t + u_t, \tag{6}$$

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1}), \tag{7}$$

$$i_t = \bar{\imath}_t + \phi_\pi(\pi_t - \bar{\pi}_t) + \phi_x x_t. \tag{8}$$

The first describes the inflation dynamics derived by a rule-of-thumb *NKPC* model, where x_t is a measure of the output gap, and u_t is a cost-push shock; in this baseline specification I assume that u_t is an i.i.d. process, and consider later the case of a small serial correlation, setting $u_t = \varrho_u u_{t-1} + \varepsilon_t$, where ε_t is white noise, and $\varrho_u = .2$. The second equation is an intertemporal *IS* equation, and the third is the policy rule. $\bar{\pi}_t$ is the long run target of the policymaker, and the time-varying intercept of the Taylor rule, $\bar{\imath}_t = \bar{\imath}(u_t, u_{t-1,} u_{t-2}, \ldots)$, embeds the optimal stabilization policy: $\bar{\imath}_t$ is in fact the response to cost-push shocks which implements optimal paths of inflation and output gap in response to the shock u_t, according to the model used by the central bank.[22] This specification of the policy rule assures that, if the model to which the policymaker conditions her policy is correct, the equilibrium paths of output gap and inflation are exactly those that are optimal under the postulated loss function.

The policymaker problem is to choose the sequence $\{\bar{\imath}_t\}$ that minimizes the following discounted loss function:

$$E_0 \sum_{t=0}^{\infty} \beta^t L_t,$$

where the instantaneous loss is specified as

$$L_t = \tfrac{1}{2}((\pi_t - \bar{\pi}_t)^2 + \lambda x_t^2),$$

subject to the constraint that Eqs. (6)–(8) be satisfied. This problem can be solved by looking for the paths of inflation and output gap $\{\pi_t, x_t\}$ that solve the constrained optimization problem.[23] These optimal paths recover, through Eq. (7), the interest rate path consistent with them, and then determine the intercept $\bar{\imath}_t$. Once $\bar{\imath}_t$ is determined, the equilibrium paths of output gap, inflation, and interest rate are determined as the solution to the system of equations that represent the model of the economy (the same system of Eqs. (6)–(8), if the policy model is correct).

To calibrate the *NKPC* parameters I chose among the values estimated by Gali et al. (2005). First, I impose that the weights on the forward and backward-looking components sum to 1: $\gamma_f + \gamma_b = 1$, and then consider as baseline value for the forward-looking parameter the value $\gamma_f = .65$. As an alternative, I also consider the case of a higher value $\gamma_f = .89$, which corresponds to a smaller degree of intrinsic persistence. Note that this equation does not appear in the exact same form as that estimated by Gali et al., where

[22]Denoting with the superscript o the response in the optimal equilibrium according to the central bank's model, $\bar{\imath}_t = i_t^o - \phi_\pi(\pi_t^o - \bar{\pi}_t) - \phi_x x_t^o$, where i_t^o is the interest rate derived from the *IS* equation under the optimal path.

[23]In practice, the *IS* constraint is not binding, so it can be ignored in this step of the calculation of optimal policy.

the driving variable is instead the labor share, a proxy for the theoretical marginal cost of labor. I obtain Eq. (6) by assuming proportionality between marginal cost and output gap (here both variables are in deviation from their steady state values), as in the model used to construct the sample economy of the previous section. The slope of the curve is then the product of the proportionality factor, which I set to .63 and the estimated parameter of the marginal cost in Gali et al., which is equal to .013, giving $\kappa = .0082$. Finally, I set $\sigma = 6.25$, and the parameters of the Taylor rule as $\phi_\pi = 1.5$, $\phi_x = .125$. For the policymaker preferences I put, as a benchmark, equal weights on the two objectives of inflation and output stabilization. Since inflation here is expressed in quarterly rather than annualized rate, the normalized weight on output in the case of equal weights in the calculations is $\frac{1}{16}$, or $\lambda = .0625$. I will show later some sensitivity analysis to the value of λ.

The second economy has the same specification for the *IS* curve and the Taylor rule, and the same calibration of the relative parameters, but the *NKPC* takes the form of a curve with trend inflation as in the model estimated by Cogley–Sbordone discussed above. The model is represented in recursive form by the following two equations:

$$\begin{aligned}\widehat{\pi}_t &= \widetilde{\kappa}_t\widehat{x}_t + \phi_t \mathrm{E}_t\widehat{\pi}_{t+1} + \gamma_t\widehat{d}_t + u_t,\\ \widehat{d}_t &= g_{1t}\mathrm{E}_t\widehat{\pi}_{t+1} + g_{2t}\mathrm{E}_t\widehat{d}_{t+1}\end{aligned} \tag{9}$$

which are obtained by compacting the first three equations of the model described in Eq. (5). I use again a proportionality factor of .63 to calibrate a value for the parameter $\widetilde{\kappa}_t$, since the model estimated in Cogley–Sbordone has the labor share as driving variable. The tilde over κ_t indicates that the 'slope' of this *NKPC* is different (indeed higher) than the one calibrated in the previous model following the estimates of Gali et al. (2005).

The coefficients of the *NKPC* in (9) are non-linear combination of the inflation trend and the underlying parameters of the pricing model, as defined in Section 3. I use Cogley and Sbordone's estimates of these underlying parameters, the probability of not changing prices and the demand elasticity. For trend inflation I consider two levels of $\overline{\pi}$ which correspond to two annualized rates of inflation, respectively, 1.5% and 2%, most frequently discussed in the policy debate. From Fig. 1, this range roughly covers the average trend inflation estimated for the past decade. The values of the coefficients in (9) depend upon these values: for the lower trend inflation case, the values are $\widetilde{\kappa} = .0313$, $\phi = 1.013$, $\gamma = .0028$, $g_1 = 6.415$ and $g_2 = 0.568$; for a trend inflation of 2%: $\widetilde{\kappa} = .029$, $\phi = 1.021$, $\gamma = .0036$, $g_1 = 6.651$ and $g_2 = 0.576$. One can see that for higher levels of trend inflation, as already discussed, the responsiveness of inflation gaps to output gaps is lower, while the importance of forward-looking terms is enhanced.

4.1.1. Results

Tables 3 and 4 contain a first set of results of this quantitative exercise. Each cell of the table corresponds to a combination of a policy model and a 'true' model of the supply side. In the cells I report the value of the discounted loss function for that particular combination, where I approximate the infinite sum by the first 64 terms. Since the parameters for each policy are calibrated, as discussed, to the empirical estimates of Gali–Gertler and Cogley–Sbordone, the initials GG and CS identify the models; next to the initials I indicate which particular value for the intrinsic persistence (γ_b) characterizes the GG model, and which annual rate of target inflation ($\overline{\pi}^A$) is associated with the CS model. The first table reports the results for the case of i.i.d. cost push shocks, and the

Table 3
Cumulative loss function—equal policy weights, i.i.d. shocks

Policy model	True model GG, $\gamma_b = .11$	GG, $\gamma_b = .35$	CS, $\bar{\pi}^A = 1.5\%$	CS, $\bar{\pi}^A = 2\%$
GG, $\gamma_b = .11$	1.221		1.036	1.096
GG, $\gamma_b = .35$		**2.862**	2.118	**2.248**
CS, $\bar{\pi}^A = 1.5\%$	1.467	4.163	.829	
CS, $\bar{\pi}^A = 2\%$	1.480	**4.184**		**.812**

Table 4
Cumulative loss function—equal policy weights, persistent shocks ($\varrho_u = .2$)

Policy model	True model GG, $\gamma_b = .11$	GG, $\gamma_b = .35$	CS, $\bar{\pi}^A = 1.5\%$	CS, $\bar{\pi}^A = 2\%$
GG, $\gamma_b = .11$	2.038		1.702	1.862
GG, $\gamma_b = .35$		**5.476**	3.888	**4.210**
CS, $\bar{\pi}^A = 1.5\%$	2.551	7.994	1.255	
CS, $\bar{\pi}^A = 2\%$	2.576	**8.033**		**1.226**

second considers the case of a mild persistence in the cost push shocks, represented by an autoregressive coefficient of 0.2. Boldface number refer to two baseline policy/true model pairs.

Consider first the case of i.i.d. shocks. The question I want to ask is what is the cost of implementing an optimal stabilization policy conditioning on a wrong model of the economy. This cost can be evaluated by comparing the cumulative loss of that policy with the loss that would have been incurred had an optimal policy for the right model of the economy been implemented. For example, to evaluate the cost of conducting policy on the basis of a model with intrinsic persistence when the true model is purely forward-looking, one should compare the numbers in the upper right 2×2 quadrant of the table with the numbers on the bottom right quadrant. To compute instead the cost of ignoring true intrinsic persistence one should compare the numbers on the bottom left 2×2 quadrant with those in the quadrant above.

The metric I choose for this comparison is the arithmetic difference between losses.[24] This metric measures the incremental cost that I want to capture, and has a decision theoretic justification in the notion of 'regret' introduced by Brock et al. (2006) for the comparison of simple versus optimal rules. In their formulation "the regret associated with a policy and a model measures the loss incurred by the policy relative to what would have been incurred had the optimal policy been based on the model." (p. 10)

[24]One could choose, of course, different approaches, and may get different conclusions. A more conservative approach such as minimax, for example, compares absolute losses. In the case of a choice between the two baseline policies that I discuss, for example, this metric would lead to choose the policy that allows for intrinsic persistence.

Consider first the cost of a policy which assumes a significant amount of persistence (the policy labeled GG, $\gamma_b = .35$ in the table) when the true model is instead purely forward-looking, and where the policymaker has a target inflation of 1.5% (in the table the 'true' model for this case is 'CS, $\overline{\pi}^A = 1.5\%$'). The regret of such a policy is measured by the increase in loss from .829, the loss that would occur if the policymaker used the correct model of the economy, to 2.118, a regret of 1.289. The regret is somewhat higher, 1.436, in the case of the higher inflation target of 2%.

For the opposite type of policy mistake—that of ignoring the degree of intrinsic inflation persistence that characterizes the economy, the cost is not particularly sensitive to the value of the inflation target. The loss, again computed for the case of high intrinsic persistence, increases from 2.862 to 4.163, in the case of a CS policy with the low target inflation, giving a regret of 1.301. When target inflation is higher, the regret of the CS policy increases to 1.322. From this comparison, a policymaker that wants to minimize the regret of the two policies would choose to ignore intrinsic persistence when target inflation is 2%, but would marginally choose the other policy when target inflation is lower.

This result is sensitive, however, to the weight assumed for the objectives of output and inflation stabilization in the loss function. As I said, the table is constructed for a loss function that assigns equal weight to the objectives of output and inflation stabilization: a case often used in policy discussions, but somewhat extreme. Fig. 2 compares instead the regrets of the two policies for a range of values of the relative weight in the loss function λ that assign more weight to the inflation stabilization objective ($\lambda \leqslant 0.0625$). The left graph in the figure shows the case with trend inflation at $\overline{\pi}^A = 1.5\%$, the right one the case where $\overline{\pi}^A = 2\%$. It is evident from the figures that in the case of higher trend inflation a policy which ignores intrinsic persistence is less costly for every value of λ, while when $\overline{\pi}^A = 1.5\%$ such a policy is less costly for values of λ which assign just a little bit more weight to inflation stabilization (the two lines in the left graph cross for $\lambda \simeq .057$).

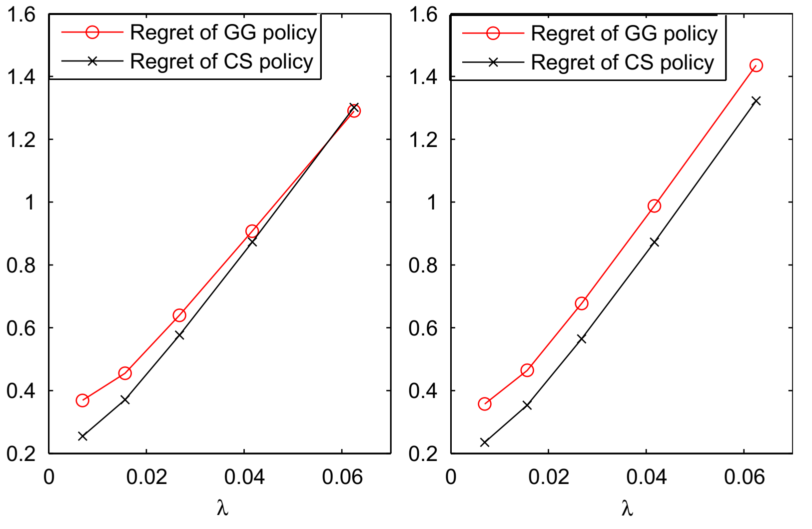

Fig. 2. Regrets assuming high persistence—left graph: $\overline{\pi}^A = 1.5\%$, right graph: $\overline{\pi}^A = 2\%$.

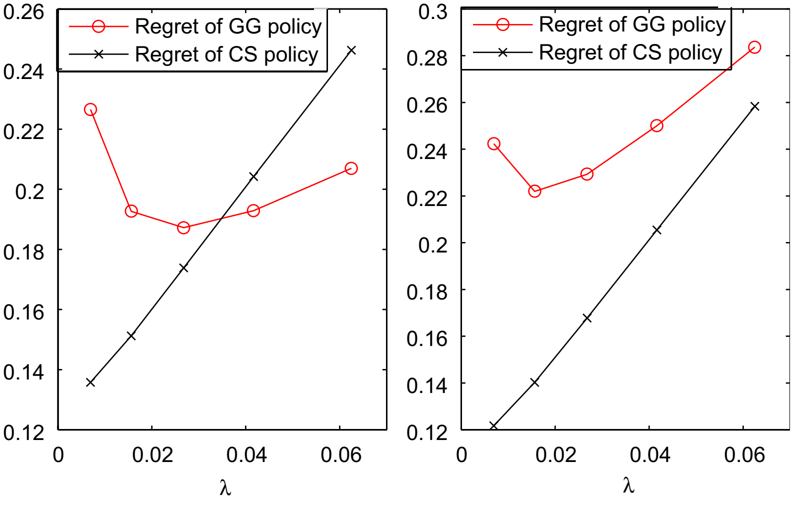

Fig. 3. Regrets assuming low persistence—left graph: $\overline{\pi}^A = 1.5\%$, right graph: $\overline{\pi}^A = 2\%$.

Similar regret comparisons can be made as well for the case of a policy model with a more moderate degree of intrinsic persistence (the losses under this policy are on the first row of the table). In this case the regret of such a policy when the true model of the economy has no intrinsic persistence is measured by the increase in loss from .829 to 1.036, a regret of .207. The regret is slightly higher, .284, in the case of a higher inflation target. For the mistake of instead ignoring a small degree of intrinsic inflation persistence the measured regrets are .246 (an increase in loss from 1.221 to 1.467) and .259, respectively, for low and high target inflation. Fig. 3 shows that when the policy assumes low intrinsic persistence the cutting point between the regret of the two policies occurs at a lower weight ($\lambda \simeq .037$) when trend inflation is low (left graph); it remains true, however, that when trend inflation is higher the cost of a policy that disregards possible intrinsic persistence is lower for all values of λ (graph on the right).

The story is only marginally different for the case of serially correlated shocks, as the numbers in Table 4 show. Although the loss is in absolute value higher for all the cases considered, the regret of a high intrinsic persistence policy is very close to the previous case. In the case of a low persistence policy, compared to the case with i.i.d. shocks, there are more values of λ for which the CS policy is less costly (it is enough that $\lambda \leqslant .054$).

Another cut at these results is presented in Figs. 4 and 5: here I look at the equilibrium paths of inflation, output and interest rate in response to a cost-push shock in three different models,[25] when the policymaker implements a stabilization policy that is optimal from the point of view of the model that she uses.

In each figure the top two panels show the equilibrium paths of output and inflation. The panel on the lower left (the intercept of the Taylor rule $\overline{\imath}_t$) is the policy response to the shock, and the last panel is the equilibrium interest rate i_t. In all graphs the horizontal axis

[25]I assume a shock of 1% at annual rate, so that the responses should be interpreted in terms of percentage variations.

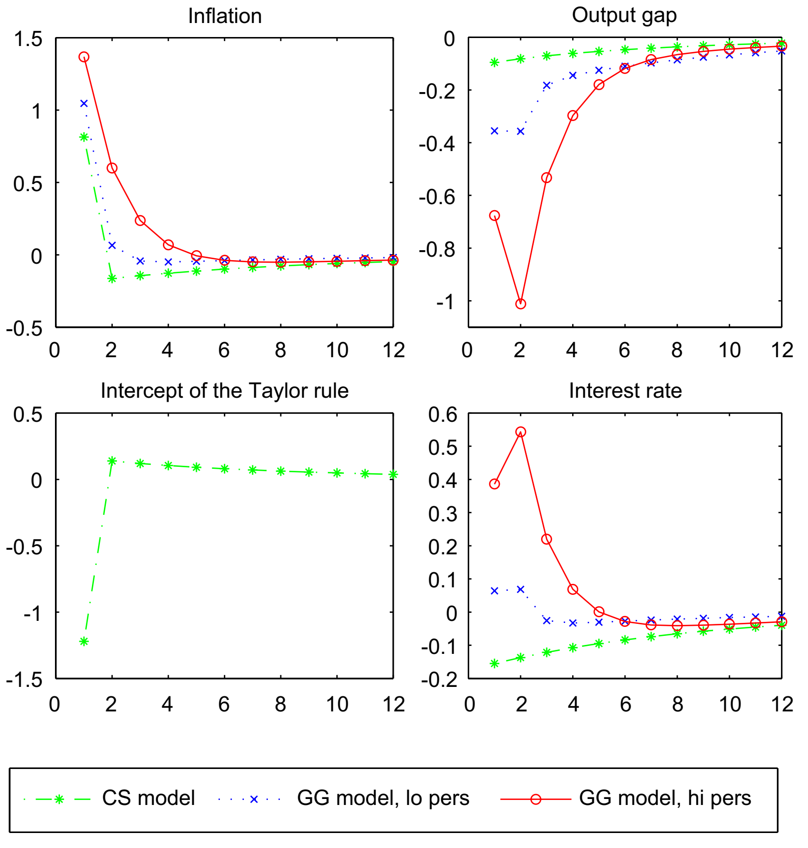

Fig. 4. CS optimal policy, i.i.d. shocks (ad hoc loss function).

indicates the periods after the shock, and the responses represent deviations of the variables from their initial steady state; in the case in which the policy model is the correct model of the economy, the equilibrium responses of output and inflation are also the responses that minimize the policymaker's loss function.

Fig. 4 considers the case of a policymaker implementing an optimal stabilization policy under the assumption that the correct model of the economy is a purely forward-looking model, which incorporates a 2% inflation target.[26] The graphs report, clockwise, starting from the lower left panel, the optimal policy and the equilibrium path of inflation, output and interest rate when the policy model is indeed the correct model of the economy. In Fig. 4 this is the CS model, and the relative paths are indicated with starred green lines.[27]

As the shock occurs, inflation increases, then declines sharply falling below target in the second period, and recovers slowly in the following quarters. The optimal response of

[26] I assume a 2% inflation target in all the figures of this exercise.

[27] Throughout the figures, I use green lines with stars to label paths under the CS model, and red lines with circle to label paths under the GG model with high persistence. The paths under the GG model with low persistence are indicated with blue dotted lines.

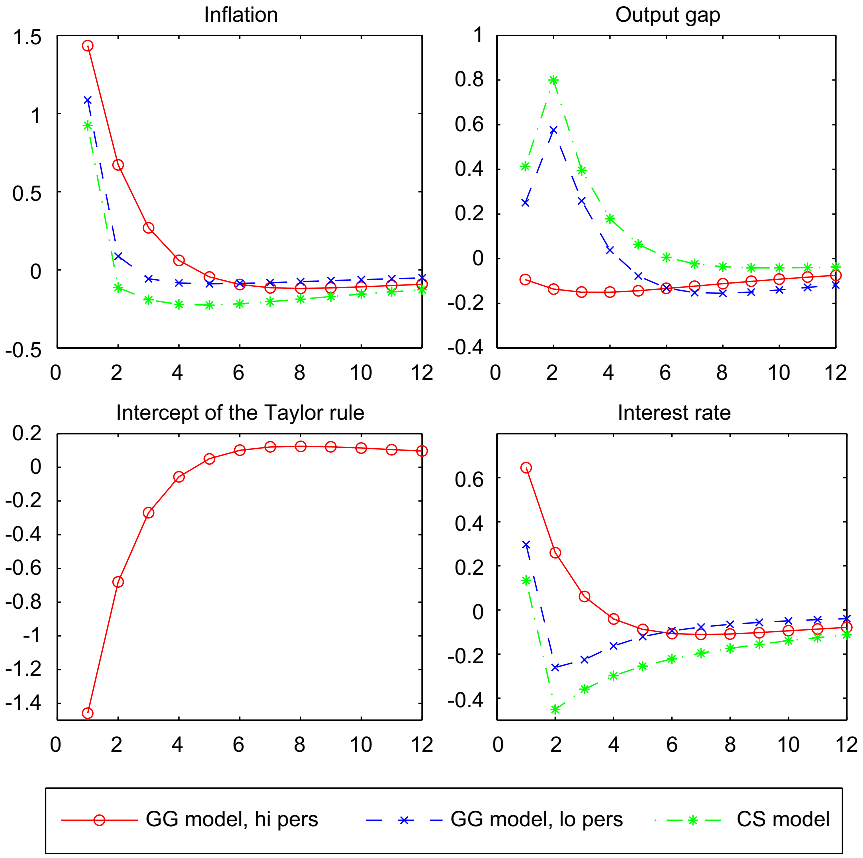

Fig. 5. GG optimal policy, i.i.d. shocks (ad hoc loss function).

output to the shock is a moderate decline at impact, and a monotonic return to steady state. These responses are engineered by the policymaker via a downward shift in the intercept of the Taylor rule at impact followed by an increase above steady state in the second period. As output gradually converges to steady state from below, the Taylor rule intercept declines to steady state from above and inflation returns to target. In equilibrium, the interest rate declines at impact, and returns monotonically to zero in the following quarters.

Although not shown in the figure, the higher the target level of inflation, the smaller is the impact effect of the shock on inflation, and the more pronounced is the second period deflation; higher target inflation also implies a greater reduction in the equilibrium interest rate, and a smaller contraction in output.

The same policy, however, delivers quite different outcomes under the alternative models. In the figure I plot the results obtained under the two parametrizations of the GG model considered previously, one with low intrinsic persistence (setting the backward-looking term in the GG model $\gamma_b = .11$), as a blue dotted line, and the other with higher intrinsic persistence ($\gamma_b = .35$), the red line with circles. If the degree of persistence in the true model is relatively low, the equilibrium path of inflation is not too different from the optimal path under the rule considered: the response to the shock is slightly higher at impact, and goes only marginally below target. The negative response of output is, however, larger and more persistent, mirrored by the path of the equilibrium interest rate, which remains above steady state for two quarters. These patterns are much more pronounced in the case in which the economy has a higher degree of intrinsic persistence. In this case the policy generates a sharp negative output gap that increases further in the second period, and a spike in the equilibrium interest rate. As the interest rate returns monotonically to steady state, the output gap slowly closes. As the figures show, an optimal policy based on a completely forward-looking model has a larger cost in terms of lost output the higher is intrinsic inflation persistence. The attempt to bring inflation down abruptly in this case does not succeed, causing instead output to contract because of the sharp increase in the interest rate.

Fig. 5 is constructed in similar way to Fig. 4, and shows the reverse situation. It assumes that policy is implemented on the basis of a model with a relatively high degree of intrinsic persistence (the backward-looking term in the GG model has a coefficient $\gamma_b = .35$). The red lines with circles indicate the optimal policy for this case (lower left panel) and the equilibrium effects of this policy (other three panels) when the true model of the economy has in fact a relatively high degree of intrinsic persistence.

The optimal response to a cost-push shock in this case is to let inflation remain above target for about four quarters, and to maintain the intercept of the Taylor rule below steady state throughout that period. Under this policy output declines only mildly below steady state and the equilibrium interest rate raises at impact, and then declines slowly back to steady state. In this experiment, as the degree of assumed persistence increases, the return of inflation to target takes longer time. During that period policy remains accommodative, and output and interest rate approach the steady state more and more gradually.

The other lines show the effects of this stabilization policy for the model with low inflation persistence and the purely forward-looking one considered in the previous figure. If the true model of the economy is purely forward-looking, equilibrium output departs significantly from the optimal path. Indeed, as we saw in the previous figure, when the model is forward-looking the optimal output response to the shock is a mild decline at impact and convergence to steady state from below. Here the more protracted period of accommodative policy determines, in equilibrium, a jump of output above steady state and of the interest rate below steady state: both effects peak in the second quarter. The equilibrium outcomes for the case of low inflation persistence are qualitatively similar to this case, albeit quantitatively smaller.

Figs. 6 and 7 illustrate the same experiment of the previous two figures, but allow the cost-push shocks to be mildly serially correlated. I calibrate the autoregressive coefficient ϱ_u to 0.2, a value often estimated for this composite shock process. The most noticeable effect of assuming serial correlation in the shock process is to smooth somewhat the optimal response of inflation to the shock in the case of both policies. For example, in the case of a policy implemented in a forward-looking model (compare the upper left panels of Figs. 6 and 4) inflation declines below target only in the third period, and in a policy model that allows for high intrinsic persistence (compare the upper left panel of Figs. 7 and 5) it takes longer time for inflation to decline to target than in the case of i.i.d. shocks.

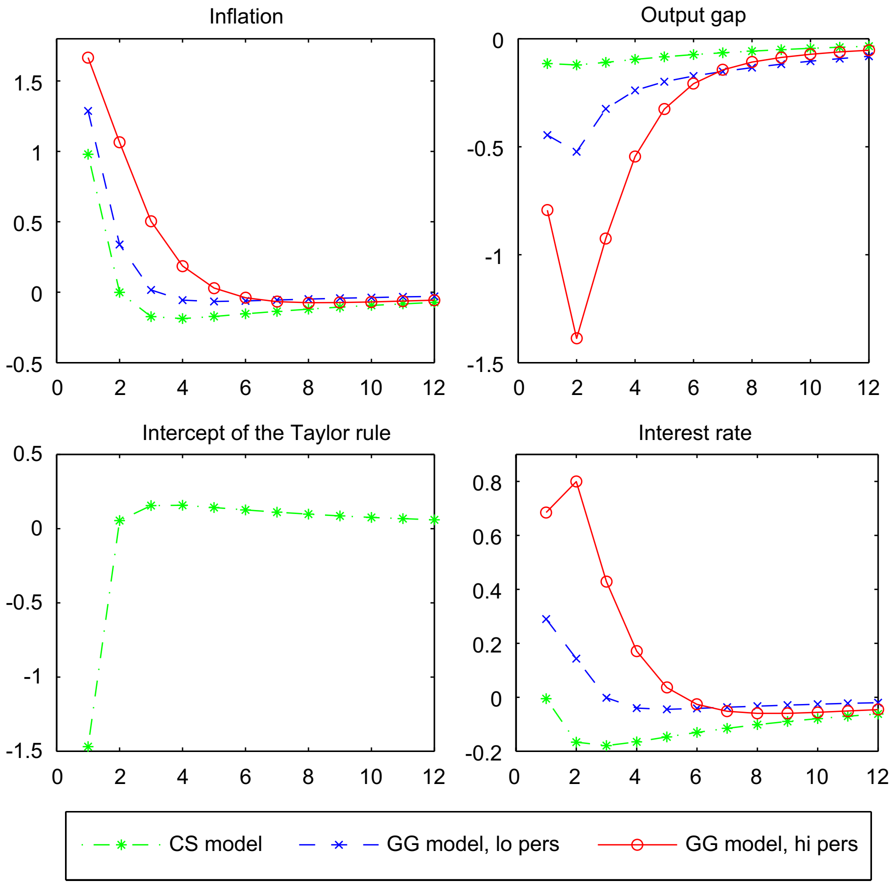

Fig. 6. CS optimal policy, serially correlated shocks.

4.2. Welfare-based loss function

In this section I repeat the previous analysis for the case of a policymaker's loss function defined in terms of the welfare implied by the micro-foundations of each model. While in the previous exercise I assumed a desired inflation target, in this case I do not have such a choice, as the optimal target inflation follows from the welfare analysis. As shown in Benigno and Woodford (2005), the optimal rate of inflation for this class of models is zero, hence it makes sense to consider an approximation of the inflation dynamics around a steady state with zero inflation. The approximation to the welfare function in the case of a model without any intrinsic persistence is therefore the one that obtains for the case of the standard purely forward-looking model, since the model here does not feature any distortion other than nominal price rigidity. Based on the result in Benigno–Woodford (2005) a welfare-based loss function for the baseline Calvo model has the same quadratic form of the ad hoc loss function considered before, but its weight depends on the model parameters, specifically: $\lambda = \kappa/\theta$. According to the

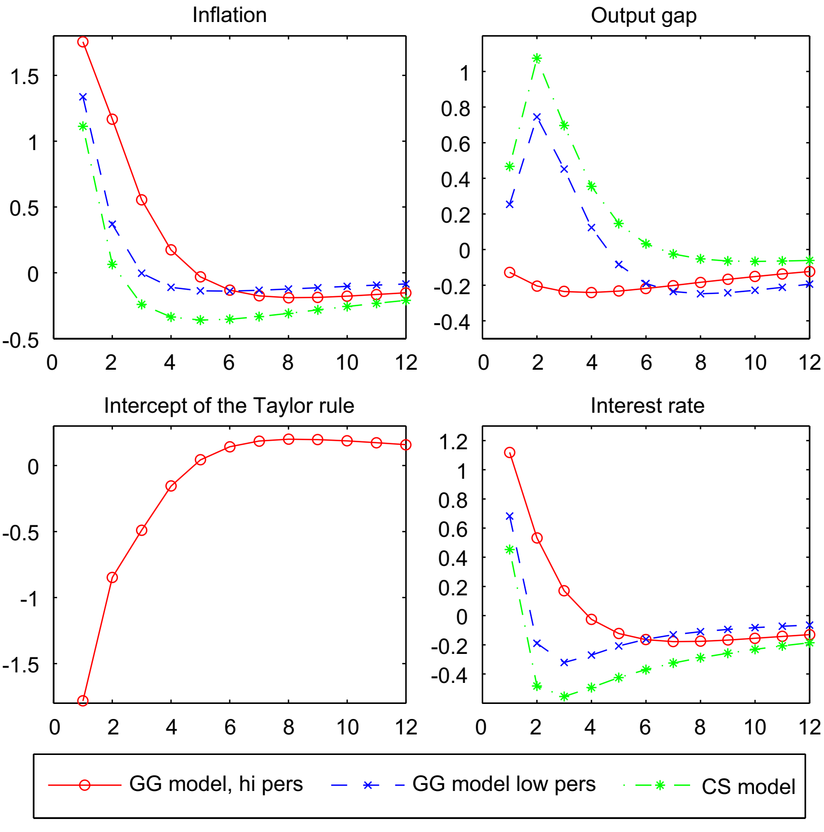

Fig. 7. G optimal policy, serially correlated shocks.

parametrization of Cogley–Sbordone (2006), setting trend inflation equal to zero, this gives $\lambda = 0.003$.

For the case of a rule-of-thumb *NKPC*, Steinsson (2003) derives a welfare-based loss function for a model where the supply side is a slight generalization of the rule-of-thumb model of Gali and Gertler (1999), and is otherwise similar to the one considered here. Based on his derivation, the approximate welfare function for the original model of Gali and Gertler (1999) is proportional to the following loss function:

$$L_t = \pi_t^2 + \lambda_1 x_t^2 + \lambda_2 \Delta\pi_t^2,$$

where the coefficients are, respectively,

$$\lambda_1 = \frac{\widetilde{\omega}(1-\alpha)(1-\alpha\beta)}{\alpha\theta(1+\theta\omega)},$$

$$\lambda_2 = \frac{\chi}{\alpha(1-\chi)}.$$

Here, the symbol χ indicates, as before, the proportion of firms that use a rule-of-thumb when resetting prices, and the other symbols have the interpretation given to them previously.

The value of χ can be backed out from the value calibrated for $\gamma_{\rm f}$ given that the two parameters satisfy the following relationship

$$\gamma_{\rm f} = \frac{\alpha}{\alpha + \chi(1 - \alpha(1 - \beta))}.$$

For the case of low persistence, $\gamma_{\rm f} = .89$ implies $\chi = .0684$, while for the higher persistence case the value of $\gamma_{\rm f}$ implies $\chi = .2978$. To get round numbers, I set $\chi_{\rm L} = .07$, and $\chi_{\rm H} = .3$, and compute the weights of the loss function on the basis of these values.

4.2.1. Results

Table 5 reports the values of the cumulative discounted loss when optimal policy is chosen in each case by minimizing the welfare-based loss functions described above.

Since in this case I consider an optimal inflation target of zero, I report the CS case only for that value. As expected, the values in the table are much lower than the corresponding values of Table 3. But a ranking of the relative losses gives more clear-cut results in this case: this is not surprising, since any welfare-based criterion puts quite a low weight on the output gap in the loss function. Implementing monetary policy under the wrong assumption that the economy is characterized by some degree of intrinsic inflation persistence increases the welfare loss more than threefold, the exact value depending on the degree of persistence assumed. Following a policy model that ignores existing persistence increases instead the welfare loss by a relatively small margin. Therefore, in the case of a welfare-based loss function, a comparison of costs across policies would always induce the policymaker to choose to ignore intrinsic inflation inertia.

As in the case of the ad hoc loss function, the next two figures analyze the effects of optimal policy under different models. Fig. 8 shows the effect of optimal policy implemented under the purely forward-looking model. Comparing this figure with Fig. 4, one sees quite clearly what is the effect of changing the relative weights in the loss function. Recall that for this model the ‘ad hoc’ loss function has the same form (to a second order approximation) of the welfare-based loss function, but the weight λ in the latter is much smaller. The result is that under an optimal policy output declines much more at impact, while the equilibrium interest rate rises less in the first period and approaches the steady state monotonically from above. Interestingly, if such a policy were implemented in an economy characterized by some degree of intrinsic persistence, the equilibrium interest rate would be much higher (see the dotted blue lines and the red lines with circles in the lower right graph of the figure) and the output loss would be slightly larger. Inflation would remain above steady state for about five quarters.

Table 5
Cumulative loss function, i.i.d. shocks

	True model		
Policy model	GG, $\gamma_{\rm b} = .11$	GG, $\gamma_{\rm b} = .35$	CS, $\bar{\pi} = 0$
GG, $\gamma_{\rm b} = .11$	1.482	1.987	1.785
GG, $\gamma_{\rm b} = .35$	1.701	**2.023**	**1.970**
CS, $\bar{\pi} = 0$	1.233	**2.330**	**.514**

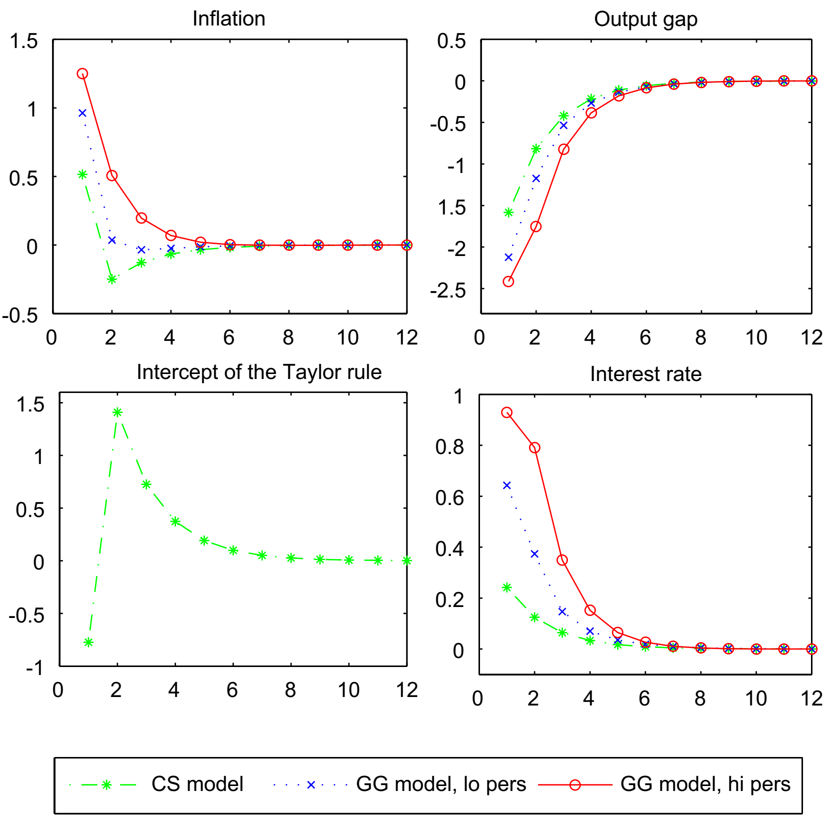

Fig. 8. CS welfare-based optimal policy, i.i.d. shocks.

Finally, Fig. 9 illustrates the implementation of a welfare-based optimal stabilization policy when policymakers assume a model with high intrinsic inflation persistence. Here, the optimal policy requires further tightening after the shock to prevent inflation to become ingrained in the economy. The tightening brings a mild deflation, and a noticeable contraction in output. However, if the economy is not characterized by intrinsic persistence, such a policy would generate instead a quite strong deflation accompanied by a period below steady state interest rates, while output would be only mildly above the optimal path (see the green lines in the graphs). In the case of a welfare-based optimal policy it is therefore much clearer which risk of model misspecification is more dangerous for the conduct of monetary policy.

5. Accounting for model uncertainty

So far I have assumed that the policymaker computes optimal policy using only one of the two alternative specifications of the *NKPC*. However, when it is not clear which model best describes the economy, it may be reasonable to adopt a policy strategy that is robust to model uncertainty. Several recent papers in the literature address the problem of integrating model uncertainty into policy evaluation (e.g. Levin and Williams, 2003; Levin et al., 1999; Brock et al., 2007). Here, I assume that the policymaker follows a Bayesian

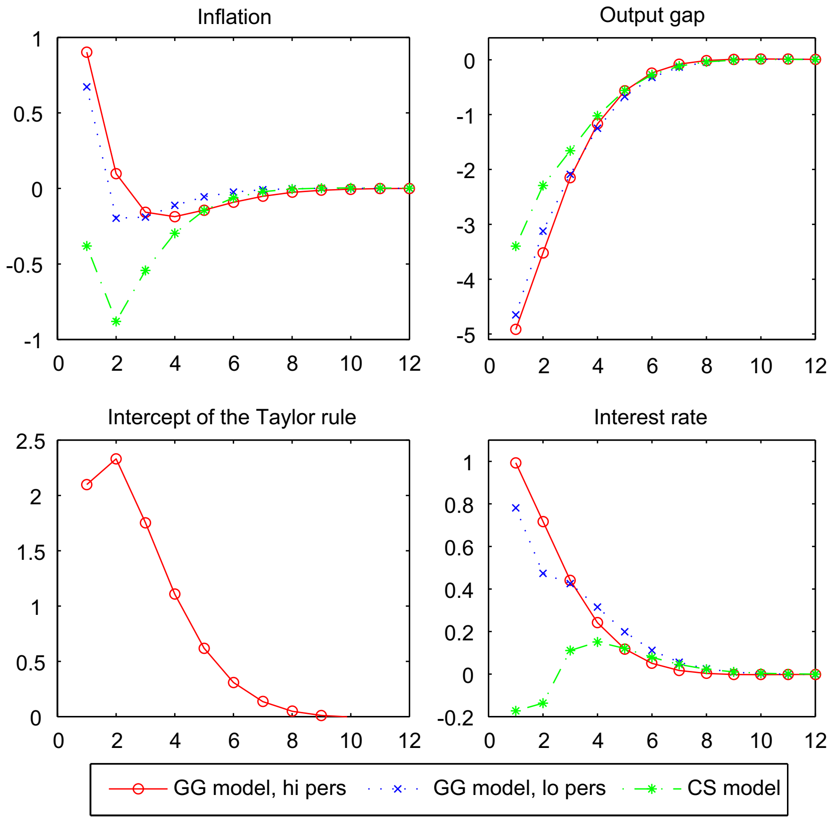

Fig. 9. GG welfare-based optimal policy, i.i.d. shocks.

approach and looks for a policy that minimizes the expected loss across models. Specifically, denoting by L_t^{GG} and L_t^{CS} the loss functions under the two alternative models of the economy, the objective function to minimize is

$$\mathscr{L} = \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t (wL_t^{\mathrm{GG}} + (1-w)L_t^{\mathrm{CS}}),$$

where the weight w represents the prior attached to the model with intrinsic inflation persistence.

Fig. 10 shows the optimal policy as a solid line in the bottom left quadrant of the figure. The other graphs plot equilibrium inflation, output gap and interest rate paths that obtain under each model (starred green line for the CS model, circled red line for the GG model) when the policymaker's objective function assigns equal weight to the two models. From a Bayesian point of view, this case corresponds to having flat priors on which of the two competing models represents a more accurate description of the economy. The lower left quadrant of the figure shows, as before, the policy. The loss functions L_t^{GG} and L_t^{CS} are the welfare-based loss functions discussed in the previous section.[28]

[28]To weight appropriately the arguments of the loss functions, however, for this computation L_t^{GG} and L_t^{CS} are no longer normalized by the weight on inflation as in the previous calculations.

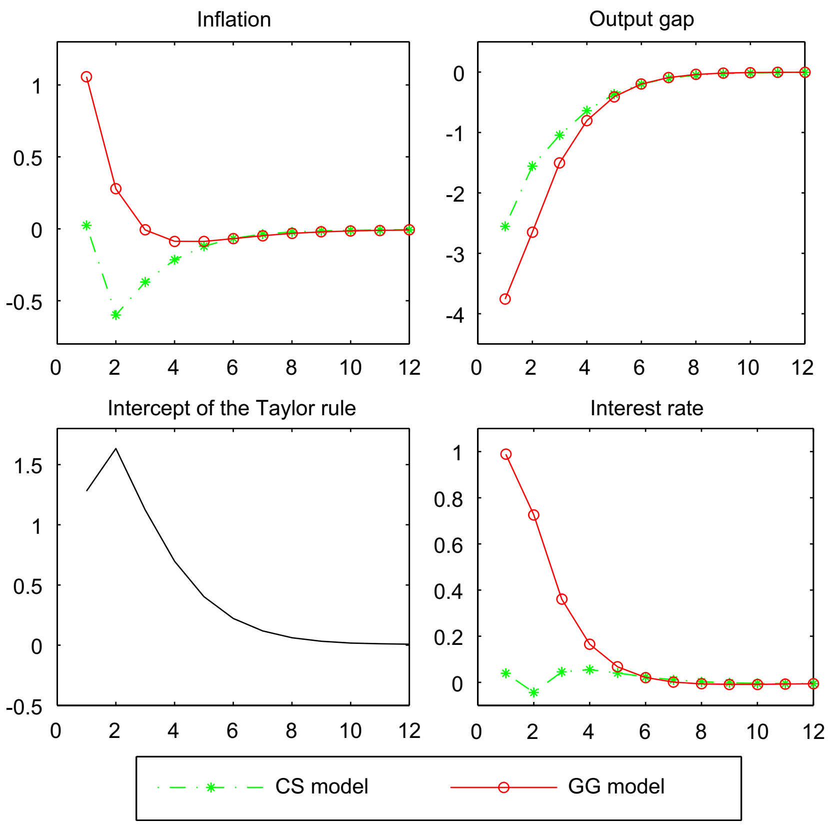

Fig. 10. Optimal policy under model uncertainty—welfare-based loss functions, $\omega = .5$.

Not surprisingly, a policy robust to model uncertainty is less aggressive than the one that would be optimal when all the weight is on the model with intrinsic persistence (compare the intercept of the Taylor rule in Fig. 10 with that in Fig. 9). As a result, compared to that case, if the true model of the economy is CS, the deflation following the shock would be milder and inflation would return faster to target; the negative output gap would be smaller and the equilibrium interest rate more stable.[29] Similarly, if the true model of the economy is GG, comparing the outcome under this policy with the one in which all the weight was put on the CS model, inflation returns to target faster, although the output loss is slightly larger (compare Fig. 10 with Fig. 8). Minimizing the expected loss brings the cumulative discounted loss to 1.556, a significant reduction relative to the cases of inappropriate policies shown in Table 4.

For the purpose of illustrating the consequence of holding stronger beliefs about one particular model, Figs. 11 and 12 replicate Fig. 10 for different weights, respectively, a larger weight on the forward-looking model the first ($w = .2$), and a larger weight on the model with intrinsic inflation persistence the second ($w = .8$).

[29]To see this compare the red line in the upper left panel of Fig. 10 with the green line in the upper left panel of Fig. 9.

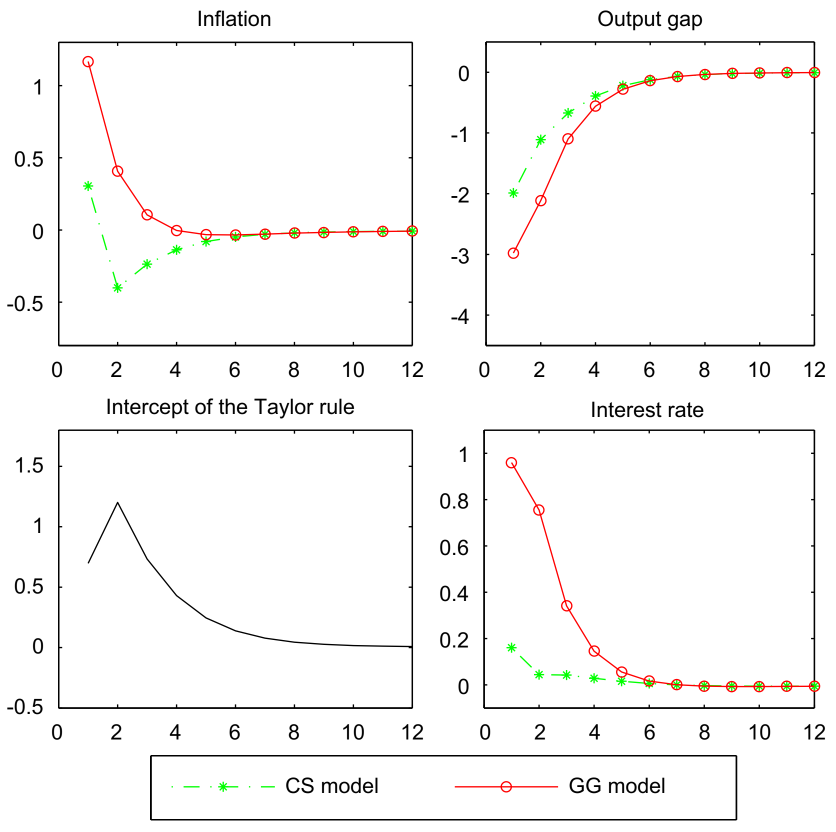

Fig. 11. Optimal policy under model uncertainty—welfare-based loss functions, $\omega = .2$.

6. Conclusion

In this paper I focus on two alternative interpretations of the observed persistence in inflation, that correspond to two alternative specifications of the New Keynesian Phillips curve. The first allows for some degree of intrinsic persistence, in the form of a term in lagged inflation in the *NKPC*. The second is a purely forward-looking model, where expectations farther into the future matter, and where the coefficients depend on the inflation trend. This specification attributes most of the observed persistence to persistent fluctuations in the underlying inflation trend; it therefore interprets inflation persistence as a consequence of monetary policy, which is ultimately responsible for the long-run level of inflation, rather than a structural feature of the economy.

I first discuss the reasons why the empirical evidence presented for structural inflation persistence in the literature may be misleading. First, I show that measures of inflation relative to trend have very different persistence properties, depending on whether one allows for a time-varying trend or not. Second, I present a simple example of how econometric inference may be misleading if the possibility of a time-varying trend inflation rate is neglected in estimating a log-linear *NKPC*.

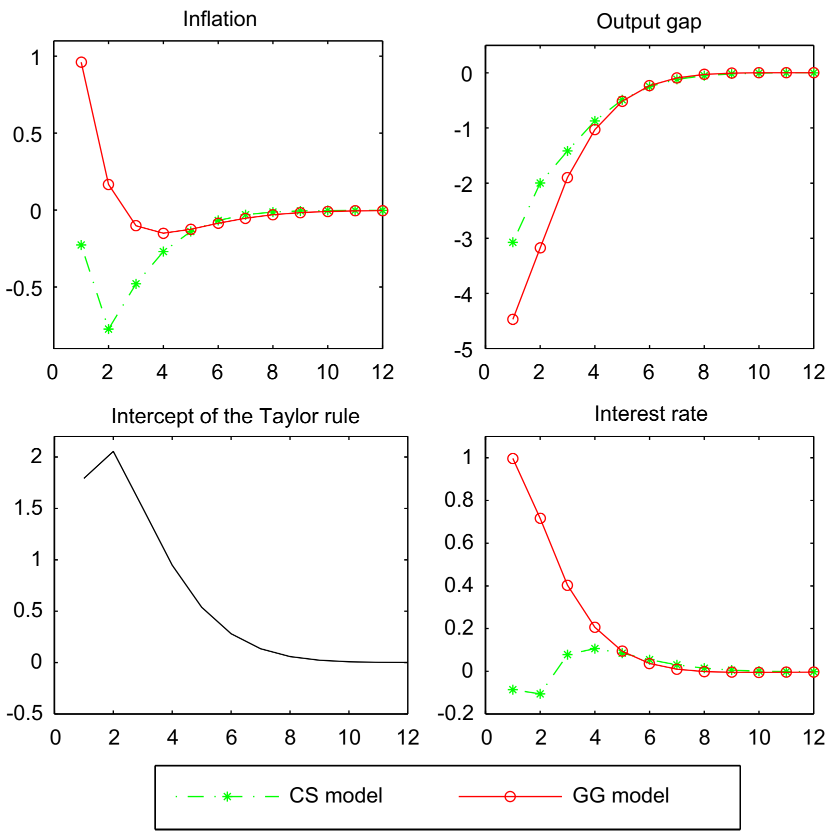

Fig. 12. Optimal policy under model uncertainty—welfare-based loss functions, $\omega = .8$.

I then analyze the consequences of optimal stabilization policy in response to cost-push shocks, when the policymaker misinterprets the degree of intrinsic inflation persistence. For this analysis I consider two very stylized economic models which differ only in their specification of the inflation dynamics. I compare the relative cost of implementing optimal policy conditioning on a wrong model of the economy, using alternative loss function specifications. I illustrate the response of inflation, output and interest rate to a cost-push shock under optimal policy, and compare them to the responses that obtain when the model economy differs from that assumed by the policymaker.

The results suggest that the cost of implementing a stabilization policy overestimating the degree of intrinsic persistence in inflation is in general higher than the costs of ignoring structural persistence; the result is more clear-cut when the inflation stabilization objective has relatively more weight in the policymaker's preferences, as it is the case when the loss function is welfare-based.

I finally discuss the same optimal stabilization exercise when the policymaker acknowledges model uncertainty and implements a policy that minimizes the expected loss across alternative models.

References

Altissimo, F., Bilke, L., Levin, A., Matha, T., Mojon, B., 2006. Sectoral and aggregate inflation dynamics in the Euro area. Journal of the European Economic Association 4 (2–3), 585–593.

Ascari, G., 2004. Staggered prices and trend inflation: some nuisances. Review of Economic Dynamics 7, 642–667.

Ascari, G., Ropele, T., 2006. Optimal Monetary Policy under Low Trend Inflation, unpublished.

Benati, L., 2006. Investigating Inflation Persistence Across Monetary Regimes, unpublished.

Benigno, P., Woodford, M., 2005. Inflation Stabilization and Welfare: The Case of a Distorted Steady State. Journal of the European Economic Association 3 (6), 1185–1236.

Brock, W.A., Durlauf, S.N., West, K.D., 2007. Model uncertainty and policy evaluation: some theory and empirics. Journal of Econometrics 136 (2), 629–664.

Brock, W.A., Durlauf, S.N., Nason, J.M., Rondina, G., 2006. Is the original Taylor rule enough? Simple versus optimal rules as guides to monetary policy, unpublished.

Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1–45.

Cogley, T., Sargent, T., 2005. Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. Review of Economic Dynamics 2, 262–302.

Cogley, T., Sargent, T., 2006. Inflation-gap persistence in the U.S., unpublished.

Cogley, T., Sbordone, A.M., 2005. A search for a structural Phillips curve. Federal Reserve Bank of New York Staff Reports No. 203.

Cogley, T., Sbordone, A.M., 2006. Trend inflation and inflation persistence in the new Keynesian Phillips curve. Federal Reserve Bank of New York Staff Reports No. 270.

Eichenbaum, M., Fisher, J., 2007. Estimating the frequency of price re-optimization in Calvo-style models. Journal of Monetary Economics, in press, doi:10.1016/j.jmoneco.2006.07.004.

Fuhrer, J.C., Moore, G.R., 1995. Inflation Persistence. Quarterly Journal of Economics 440, 127–159.

Gali, J., 2003. New perspectives on monetary policy, inflation and the business cycle. In: Dewatripont, M., Hansen, L.P., Turnovsky, S.J. (Eds.), Advances in Economics and Econometrics. Theory and Applications, Eighth World Congress, vol. III. Cambridge University Press, Cambridge, MA, pp. 151–197.

Gali, J., Gertler, M., 1999. Inflation dynamics: a structural econometric approach. Journal of Monetary Economics 44 (2), 195–222.

Gali, J., Gertler, M., Lopez-Salido, D., 2001. European inflation dynamics. European Economic Review 45 (7), 1237–1270.

Gali, J., Gertler, M., Lopez-Salido, D., 2005. Robustness of the estimates of the hybrid new Keynesian Phillips curve. Journal of Monetary Economics 52 (6), 1107–1118.

Ireland, P. Changes in the Federal Reserve's inflation target: causes and consequences. Journal of Money, Credit and Banking, forthcoming.

Levin, A.T., Williams, J.C., 2003. Robust monetary policy with competing reference models. Journal of Monetary Economics 50 (5), 945–975.

Levin, A.T., Lopez-Salido, D, Yun, T., 2006. Strategic Complementarities and Optimal Monetary Policy, unpublished.

Levin, A.T., Wieland, V., Williams, J.C., 1999. Robustness of simple monetary policy rules under model uncertainty. In: Taylor, J.B. (Ed.), Monetary Policy Rules. University of Chicago Press.

Lucas Jr, R.E., 2000. Inflation and welfare. Econometrica 68 (2), 247–274.

Milani, F., 2007. Expectations, learning and macroeconomic persistence. Journal of Monetary Economics, in press, doi:10.1016/j.jmoneco.2006.11.007.

Orphanides, A., Williams, J.C., 2002. Robust monetary policy rules with unknown natural rates. Brookings Papers on Economic Activity, vol. 2, pp. 63–118.

Pivetta, F., Reis, R., 2007. The persistence of inflation in the United States. Journal of Economic Dynamics and Control 31 (4), 1326–1358.

Sbordone, A.M., 2005. Do expected future marginal cost drive inflation dynamics? Journal of Monetary Economics 52 (6), 1183–1197.

Sbordone, A.M., 2006. U.S. wage and price dynamics: a limited information approach. International Journal of Central Banking 2 (3), 155–191.

Smets, F., Wouters, R., 2003. An estimated dynamic stochastic general equilibrium model of the Euro area. Journal of the European Economic Association 1 (3), 1123–1175.

Smets, F., Wouters, R., 2005. Comparing shocks and frictions in US and euro area business cycles. Journal of Applied Econometrics 20 (2), 161–183.

Steinsson, J., 2003. Optimal monetary policy in an economy with inflation persistence. Journal of Monetary Economics 50 (7), 1425–1456.

Stock, J., Watson, M., 2007. Why has U.S. inflation become harder to forecast? Journal of Money, Credit and Banking 39 (1), 3–34.

Woodford, M., 2003. Interest and Prices. Foundations of a Theory of Monetary Policy. Princeton University Press, Princeton, NJ.