Inflation, Productivity and Monetary Policy: from the Great Stagflation to the New Economy*

Andrea Tambalotti†
Federal Reserve Bank of New York

September 2003

Abstract

This paper investigates how the productivity slowdown and the systematic response of monetary policy to observed economic conditions contributed to the high inflation and low output growth of the seventies. Our main finding is that monetary policy, by responding to real time estimates of output deviations from trend as its main measure of economic slack, provided a crucial impetus to the propagation of the productivity shock to inflation. A central bank that had responded instead to a differenced measure of the output gap would have prevented inflation from rising, at the cost of only a marginal increase in output fluctuations. This kind of behavior is likely behind the much-improved macroeconomic performance of the late nineties.

*I wish to thank Pierre-Olivier Gourinchas and Michael Woodford for comments and advice and Giorgio Primiceri for extensive conversations. The views expressed in the paper are those of the author and are not necessarily reflective of views at the Federal Rserve Bank of New York or the Federal Reserve System.

†Research and Market Analysis Group, Federal Reserve Bank of New York, New York, NY 10045. Email: `Andrea.Tambalotti@ny.frb.org`.

1 Introduction

The attempt to account for the unusual behavior of US inflation and real activity in the seventies, what Blinder (1979) dubbed the “great stagflation”, has recently attracted a great deal of attention in the literature. This is partly the result of a growing consensus among monetary economists on the foundations of a theory of monetary policy, as articulated most recently by Woodford (2003). One of the tenets of this consensus is that the complexities behind the conduct of monetary policy can be usefully summarized, for both theoretical and empirical purposes, through a simple interest rate rule of the kind first proposed by Taylor (1993). This has then suggested a natural strategy to investigate the role of monetary policy in the starkly different macroeconomic performances of the US economy before and after Paul Volcker’s appointment as chairman of the Federal Reserve, namely to look for differences in estimated interest rate rules across the two periods.[1]

This strategy has been followed for example by Clarida, Galí, and Gertler (2000) (henceforth CGG), who find significant differences in the coefficients of interest rate rules estimated (roughly) before and after 1979.[2] In particular, they attribute the extreme fluctuations of the seventies to the failure of nominal interest rates to move more than proportionally in response to changes in inflation. According to the “Taylor principle” in fact, interest rate rules that do not guarantee a positive long-run response of real rates to increases in inflation give rise to a multiplicity of rational expectations equilibria, including, but not limited to, sunspot fluctuations. The observed swings in inflation and real activity of the seventies, and the remarkable stability of the last two decades, would then simply reflect the different emphasis on controlling inflation that characterized the two policy “regimes”.

An alternative to this view has been proposed by Orphanides (2000b; 2000a; 2001; 2002).[3] According to his interpretation of the data, the policy mistake of the seventies was spurred by an excessive reliance on imperfect real-time measures of the amount of “slack” in the economy. Not having realized that the growth rate of potential output had significantly slowed in the early seventies, the policy authorities strived to reach an overambitious notion of full employment, delivering high inflation instead . Orphanides supports this story by showing that interest rate rules estimated on the pre and post-Volcker samples display very similar coefficients, once a real-time measure of the output gap is included into the equation. In fact, the main difference between the two samples is not in the response to inflation, as argued by CGG, but rather in an attenuated response to the

[1] This research program falls under the rubric of Christiano and Fitzgerald’s (2003) “people’s view” of the great stagflation. There is of course an even more extensive literature on the “institution view” of that episode, started by Kydland and Prescott (1977) and Barro and Gordon (1983), which attributes the bad macroeconomic outcomes of the seventies to the lack of commitment in monetary policy.

[2] Taylor (1999), Judd and Rudebusch (1998) and, in a fully specified model, Lubik and Schorfheide (2002) and Boivin and Giannoni (2003) reach very similar conclusions. Sims (2001), Sims and Zha (2002) and Primiceri (2003) on the contrary find little evidence of a discontinuity in the policy regime. They emphasize instead the reduction in the volatility of the economic environment in the latter part of the sample.

[3] See also Lansing (2000), Bullard and Eusepi (2003) and Nelson and Nikolov (2002), who consider the case of the UK.

output gap after 1979.

This paper reconsiders the questions addressed by CGG and Orphanides from the perspective of the model of inflation and output determination developed in Tambalotti (2002). There we argue that to understand the macroeconomic effects of persistent changes in the growth rate of productivity, like the slowdown of the early seventies and the pickup of the early nineties, we need to take into account agents' uncertainty on the persistence of observed shocks. One of the distinguishing features of the model is that all economic agents, including the central bank, share this uncertainty. Even in this environment however, a policy authority that follows a Taylor rule could not have possibly caused the great stagflation. This is true independently of the exact rule specification, as long as it maintains determinacy of the rational expectations equilibrium and it responds to the "best " measure of the output gap available to the central bank.[4]

In this paper, we investigate the consequences of relaxing this last assumption. In particular, we consider a monetary authority that responds to an ad hoc measure of the output gap, the deviation of output from a recursively estimated linear trend. Although hard to justify from the perspective of the model, in which the central bank has access to a fairly accurate measure of the "true" output gap, this specification has a clear positive content. It captures the fact that, before Nelson and Plosser (1982) potential output was very often modeled as a deterministic linear trend. This assumption is consistent with the slow revision of the estimates of trend growth that followed the productivity slowdown. In our model, it also produces output gaps which are comparable with the real time measures used by Orphanides. One of the advantages of our approach is that the output gap that enters the central bank's reaction function is endogenous, since it depends on the equilibrium realizations of output. Given the forward looking nature of private agents' decisions, this allows us to model explicitly their expectations on the future stance of policy, as a function of their views on the central bank's future estimates of potential. Following the literature on bounded rationality and least squares learning, we assume that the private sector simply projects the current estimated trend into the future, without internalizing the effect of its own decisions on future trend estimates.[5] Further study of the role of expectations' formation on equilibrium outcomes is an important avenue for further research (see Bullard and Eusepi (2003) for some progress in this direction).

As for the CGG hypothesis, it is addressed here only partially. In particular, we consider policies characterized by a weak response of interest rates to inflation, similar to the ones estimated by CGG on the pre-Volcker period. However, given the technical difficulties involved in characterizing the associated multiple equilibria, we modify their specification just enough to avoid the region of indeterminacy. This allows us to verify if the volatility that characterized the great stagflation could be ascribed to an "inactive" monetary policy, but one that would not necessarily unleash the

[4] In the model, perfect observations on the "true" output gap are available to the central bank with a one period delay. The "best" measure of the output gap referred to in the text is then its one period ahead forecast.

[5] See for example Sargent (1993) and Evans and Honkapohja (2001) for a general treatment and Sargent (1999) for an application to the study of American inflation.

destabilizing properties of sunspots. As we already pointed out, several empirical studies do not find much evidence of a deviation of monetary policy from the Taylor principle in the Pre-Volcker period, making our hypothesis of clear empirical relevance.

To complete our positive investigation of the effect of persistent growth shocks on inflation and real activity, and on the role of monetary policy in their transmission, we also extend the analysis to the recent episode of productivity acceleration, what has been referred to as the "new economy". Our main finding is that there is indeed a significant difference in the conduct of monetary policy during the seventies, as opposed to the second half of the nineties. In accordance with Orphanides' hypothesis, this difference has mainly to do with the way in which the Federal Reserve calculates its reference measure of economic "slack", rather than with the intensity with which it responds to its observations on inflation. Given our model of private behavior, an interest rate rule that responds to the deviations of output from a linear trend successfully accounts for the general profile of inflation during most of the seventies and early eigthies. Remarkably, this includes the entire extent of Volcker's disinflation, without the need for any change in the rule's coefficients, including the bank's inflation "target".[6] In the second half of the nineties however, such a rule would have resulted in a counterfactual decline in inflation, mirroring the sharp increase observed in the seventies. According to our analysis, the Fed avoided this outcome by reacting to a measure of the output gap that takes into account the non-stationary nature of potential output. Moreover, we show that such a policy would have significantly dampened the inflationary effects of the productivity slowdown, with only marginal costs in terms of added output volatility.

Finally, it is worth commenting here on the basic methodology followed in this study. Our approach consists of simulating the model, using as input the observed sequence of productivity growth rates. Every period, the private sector observes the current realization of productivity, filters its permanent from its temporary component, and takes decisions on prices, wages and demand for consumption goods, given the current and expected future stance of policy. The monetary authority on the other hand, having observed these decisions, updates its estimate of the output trend and sets the interest rate according to the specified interest rate rule. This rule is also the basis on which private agents forecast its future behavior.

The output of this procedure are artificial sequences of inflation and output growth rates, which we compare with the data. The objective is to investigate to what extent different policy rules can replicate the main features of the data. Admittedly, this is a very informal criterion for model comparison. However, we will see that the differences in the predictions of some of the models are so macroscopic that it would be hard to deny them any significance. Of course, the existence of such differences encourages us to pursue some more formal statistical work to try to discriminate between more subtle alternatives. However, even with the very blunt instrument of simple visual inspection, we think our analysis can contribute to the debate on the causes of the great stagflation, and on the policy measures that might prevent such an episode in the future.

[6] In two papers very similar to ours, both Lansing (2000) and Bullard and Eusepi (2003) need to include a "shock" to the inflation "target" to account for Volcker's disinflation.

The rest of the paper is structured as follows. Section 2 presents our model of the economy, which was derived from microfoundations in Tambalotti (2002). In its log-linear form, the model is composed of equations describing consumers' demand and wage and price setting, and closed by an interest rate feedback rule capturing the behavior of the monetary policy authority. Details on the specification of the rule are contained in section 2.2, while section 2.3 details the simulation strategy followed in the rest of the paper. Section 3 presents simulations for the twenty quarters between 1973:III and 1978:II, the same period analyzed in Tambalotti (2002). Section 4 extends this analysis to include Volcker's disinflation, while in section 5 the sample is further extended to include the second half of the sixties. In section 6 we conduct counterfactual experiments on the performance of several alternative policy rules, taking as given the observed sequence of productivity shocks. These exercises lead us to advance some normative conjectures on the relative merits of the set of rules considered. Finally, section 7 shifts the focus of the analysis to the second half of the nineties. Section 9 concludes.

2 The Model

This section presents the model adopted in the rest of the paper as the data generating process for simulating the effect of different policy rules on inflation and real activity. Its private sector equations are derived as log-linear approximations to the first order conditions of a dynamic stochastic general equilibrium model with sticky prices in Tambalotti (2002). There we also argue that, together with an empirical interest rate rule, the model provides a fairly accurate description of the dynamic response of the U.S. economy to an identified productivity shock. As in that positive exercise, here we also close the model with an interest rate feedback rule, whose exact specification is the main source of variation in the experiments conducted below, as detailed in section 2.2.

The choice of an interest rate rule to describe the behavior of the monetary policy authority introduces an important asymmetry with the private sector. Even if in the model's general equilibrium all variables are simultaneously determined, it is useful to think of the latter as *optimally* determining prices, wages and the level of demand (and therefore of output), given observations on the current and past interest rate, and forecasts of its future level. The monetary authority, on the other hand, *mechanically* sets the interest rate as a function of present and past observed economic conditions only, without any attempt at optimization. Following the rule therefore does not require the central bank any insight into the behavior of the private sector, or to forecast future economic conditions. All the bank needs to do instead is to observe some readily available economic indicators, like output and inflation, and to commit to set its instrument as a particular function of those measurements. To close the equilibrium loop then, this commitment forms the basis for the private sector's forecasts of the future stance of policy, which is one of the inputs into their current decisions.

This asymmetry, although undesirable in terms of theoretical consistency, simplifies the analysis significantly. In fact, by severing the equilibrium feedback channel from private agents' expectations to policy decisions, it allows us to disregard the strategic dimension of the interaction between the two sectors highlighted by the time inconsistency literature (e.g. Kydland and Prescott, 1977; Barro and Gordon, 1983). In particular, given the learning framework in which we assume policy is conducted, this simplification limits the hierarchy of reciprocal beliefs that would otherwise be relevant in equilibrium to the first order consideration of how the private sector thinks the central bank behaves, as we will see in section 2.2 below.

Let us then start the description of this economy from the set of equations pertaining to the behavior of the private sector. We refer the reader to Tambalotti (2002) for more details on the model's assumptions and for its derivation from first principles.

2.1 The Private Sector

The private sector is composed by a representative consumer, who takes spending decisions to maximize her intertemporal utility; by workers, who set wages given the level of demand for their labor, and by firms, which set prices to maximize the present discounted value of their profits. Both workers and firms are assumed to operate in monopolistically competitive markets and to face a "technological" impediment to their ability to set prices, which we model as a random Poisson arrival of price setting opportunities as in Calvo (1983). From this environment, given standard assumptions on tastes and technology, we can derive log-linear approximations to agents' optimal decision rules which characterize the dynamic behavior of prices, wages and output, for any given sequence of interest rates.

In particular, from the representative consumer's Euler equation, we can derive

$$y_t = E_{t-1}y_{t+1} - E_{t-1}\left[\hat{\imath}_t - \pi_{t+1}\right] \tag{1}$$

where y_t denotes the logarithm of output, π_t is inflation and $\hat{\imath}_t$ is the nominal interest rate, expressed in deviations from its steady state value in the absence of growth, $\bar{\imath} \equiv -\ln\beta$, with β the discount factor. Note that in this formulation the level of output is predetermined, or measurable with respect to I_{t-1} information, where I_{t-1} includes the history of all endogenous variables—prices, wages, output and interest rates—dated up to period t, together with the current realization of the level of productivity A_{t-1}.

Coming now to the price and wage setting decisions, we follow Erceg, Henderson, and Levin (2000) and treat workers and firms symmetrically, as facing frictions of the same nature. The solutions to their respective optimization problems yield therefore very similar expressions for the resulting rates of price and wage inflation, π_t and π_t^w respectively[7]

$$\pi_t = \left(1-\psi_p\right) E_{t-2}\pi_t + \psi_p E_{t-1}\left[\kappa_p x_t + \xi_p \hat{\omega}_{At} + \beta\pi_{t+1}\right] \quad (2a)$$

$$\tilde{\pi}_t^w = \left(1-\psi_w\right) E_{t-2}\tilde{\pi}_t^w + \psi_w E_{t-1}\left[\kappa_w x_t - \xi_w \hat{\omega}_{At} + \beta\tilde{\pi}_{t+1}^w\right] \quad (2b)$$

where x_t, the output gap, is the percentage deviation of actual output from its efficient level, and $\hat{\omega}_{At}$ similarly denotes the deviation of the real wage from the value that would prevail in the absence of any price-setting or informational friction. These efficient levels are proportional to the level of productivity and represent the main channel through which changes in productivity are transmitted to the equilibrium prices and quantities. So for example, everything else being equal, an increase in productivity would result in a decline of the wage gap $\hat{\omega}_{At}$, causing goods' prices to fall and wages to increase. In the long run, this adjustment will be such that the real wage moves in proportion to any permanent shift in productivity. Note that this cointegration property is imposed by assuming that nominal wages are automatically indexed to workers' forecasts of the growth rate of productivity, $\gamma_t^a \equiv \ln\left(A_t/A_{t-1}\right)$, so that $\tilde{\pi}_t^w \equiv \pi_t^w - E_{t-1}\gamma_t^a$. Moreover, as in Rotemberg and Woodford (1997) and Amato and Laubach (2003), new prices and wages are set on the basis of either I_{t-1} or I_{t-2} information, in proportions that depend on the ψ coefficients. Finally, κ's and ξ's, all positive coefficients, denote the elasticity of inflation to the output and wage gap and are convolutions of the structural parameters that describe tastes, technology and price-setting frictions.[8]

2.1.1 Productivity and Information

The model's main point of departure from the literature is in its treatment of productivity and information. Having been originally developed to study the optimal policy response to persistent changes in the growth rate of productivity, the model explicitly incorporates a non-stationary process for productivity. In particular, we assume that the growth rate of productivity γ_t^a is the sum of a trend component γ_t, and of a cyclical component ε_t^a, which for simplicity is restricted to be i.i.d.

$$\gamma_t^a = \gamma_t + \varepsilon_t^a \quad (3)$$

In Tambalotti (2002), we furthermore assumed that γ_t is a persistent, but stationary process around a mean γ. This allowed us in turn to log-linearize the model around this mean growth rate, being confident that the approximation would remain first-order accurate, given an appropriate bound on the shocks. However, this "precaution" was crucial only in the context of the welfare analysis conducted in that paper, which requires to evaluate the expected value of the representative agent's utility over the infinite future. Over that horizon, a non stationary γ_t would almost

[7] The model's underlying market structure is one of monopolistic competition in a continuum of markets for differentiated kinds of labor and goods. "Price" and "wage" inflation rates therefore refer to quarterly changes in the price indexes associated with the Dixit-Stiglitz aggregators of goods and labor services demanded by consumers and producers in those markets.

[8] As shown in Tambalotti (2002), the wage Phillips curve should include a correction term equal to the two step ahead forecast error of the growth rate of productivity. For reasons of expositional simplicity this term has been omitted here, but will be included in the experiments conducted in the following sections.

surely stray away from any constant growth rate, invalidating the first order characterization of the system's dynamics used in the welfare calculations. However, for the purposes of this paper, in which the analysis is limited to simulating the effect of certain sequences of shocks on the behavior of the model's endogenous variables, we need not worry about the trend growth rate being stationary, as long as we can verify ex-post that those shocks do not push it "too far" from the constant growth rate γ around which the model is still approximated.[9]

In light of the these considerations then, we assume for the remainder of this paper that the trend growth rate γ_t is itself a random walk,

$$\gamma_t = \gamma_{t-1} + \varepsilon_t^{\gamma} \tag{4}$$

or at least that this is a reasonable statistical representation of its true distribution, in the sense that it cannot be formally rejected given the available data. In other words, we are assuming that, just like many real world econometricians (see for example Stock and Watson, 1998; French, 2001; Roberts, 2001), the agents in our artificial economy model the trend growth rate of productivity on the basis of a parsimonious random walk specification, even if theoretical considerations might suggest that this process should in fact be stationary. Given our simulation approach, which involves taking the observed productivity series as an input and letting agents filter this data according to the statistical model (3)-(4), it becomes irrelevant whether this model is "true" in a positive sense, as long as it represents a good approximation to the way in which agents' form their forecasts (see Barsky and De Long (1993) for a similar argument).

The second defining characteristic of our model of productivity is that agents are assumed to observe the growth rate of productivity γ_t^a, but not its individual components γ_t and ε_t^a separately. As shown in Tambalotti (2002), besides being clearly realistic, this assumption is important to explain the observed negative correlation between inflation and real activity, conditional on persistent productivity shocks. The basic mechanism at play in the model is fairly simple. Given their lack of information on the state of productivity, agents use a Kalman filter to estimate its trend based on the sequence of past growth rates, $\{\gamma_s^a\}_{s\leq t}$. Given our distributional assumptions, this estimate, denoted by $\gamma_{t|t}$, takes the form

$$\gamma_{t|t} = K\gamma_t^a + (1-K)\,\gamma_{t|t-1} \tag{5}$$

where K, the Kalman gain, is a small constant related to the signal to noise ratio $\mathrm{Var}(\varepsilon_t^{\gamma})\,/\mathrm{Var}(\varepsilon_t^a)$.[10]

In this environment, following a trend shock, agents update their estimate of the long run growth rate, and therefore their wage demands, only gradually. This in turn introduces a wedge between real wages and realized productivity, inducing a decline in marginal costs and hence in

[9] In practice, we choose this value as the average growth rate of productivity in the period preceding that of the simulations. So for example, when focusing on the period 1973:III to 1978:II, γ is the average productivity growth between 1948:I and 1973:II.

[10] An interesting property of this filter is that it has the same form as the constant gain algorithm that agents would use to adaptively learn the value of a "constant" γ_t in representation (3). Under this interpretation, agent's forecasting model would closely resemble the one for dividend growth proposed by Barsky and De Long (1993).

newly optimized prices. This mechanism then has the potential to explain the negative correlation between inflation and real activity that characterized the best part of the 1970s, as well as the second half of the 1990s, both periods that followed significant and extremely persistent productivity shocks, although of opposite sign. Of course, to fully explore this hypothesis, we need to close the model with a description of the behavior of the policy authority. This task is taken up in the next section.

2.2 Monetary Policy

Since the work of Taylor (1993), it has become very common in monetary policy research to summarize the behavior of central banks by way of an interest rate feedback rule. So, for example, several recent accounts of the "great inflation" of the 1970s, and of the much improved macroeconomic performance of the last two decades, are based on the premise that the Federal Reserve sets the federal funds rate according to a rule of this kind (e.g. Clarida, Galí, and Gertler, 2000; Orphanides, 2000a; Orphanides, 2001; Lansing, 2000; Bullard and Eusepi, 2003; Christiano and Gust, 1999). In fact, most of those studies explain the differences in the macroeconomic outcomes of the two periods as the result of the rule's specification, attributing the failures of the seventies either to an insufficient reaction of nominal interest rates to changes in inflation, CGG's hypothesis, or to an overoptimistic reliance on a mismeasured indicator of inflationary pressures, Orphanides' hypothesis. Since one of the objectives of this paper is to explore whether these explanations of the "great inflation" are robust to the modeling environment introduced above, we also assume that monetary policy is set according to a simple interest rate rule of the form

$$\hat{\imath}_t = \phi_i \hat{\imath}_{t-1} + (1-\phi_i)\left[\phi_c + \phi_n \bar{r}_t^n + \phi_\pi (\pi_t - \pi^*) + \frac{1}{4}\phi_x x_t^{CB}\right] \tag{6}$$

where x_t^{CB} and $\bar{r}_t^n$ are the bank's estimates of the output gap and of the "natural" real rate of interest, which is defined as the rate that would prevail in equilibrium in the absence of price rigidities.

While the general form of this rule is entirely standard, the definition of some of its arguments is not. First, inspired by Orphanides' hypothesis, we consider rules in which the central bank responds to an imperfect measure of the output gap, rather than to the "true" output gap x_t that appears in the price and wage setting equations. In particular, we assume that x_t^{CB} is constructed as the deviation of output from a linear trend, either in levels, or in first differences (in which case the gap will be denoted by $\overline{\Delta x_t}$). More precisely, under the "level specification", $x_t^{CB} = \bar{x}_t \equiv y_t - \bar{y}_t$, where $\bar{y}_t$ is a linear trend whose constant and slope coefficients, collected in the vector c_t, are estimated recursively as

$$\begin{aligned} c_t &= c_{t-1} + kR_t^{-1} z_t \left(y_{t-1} - z_{t-1}' c_{t-1}\right) \\ R_t &= R_{t-1} + k\left(z_t z_t' - R_{t-1}\right) \end{aligned}$$

with $z_t' \equiv \begin{bmatrix} 1 & t \end{bmatrix}$ and k a constant gain. In words, under the level specification, policy makers "estimate" potential output as a simple linear trend, and update their estimates of the trend's coefficients by a recursive least squares algorithm. Moreover, they insure themselves against the possibility of breaks in those coefficients by exponentially discounting past observations through the constant gain k.[11]

Note that, in the model, $x_t \equiv y_t - a_t - \bar{y}_A$, where a_t is the logarithm of the level of productivity and $\bar{y}_A$ is a constant. Therefore, in the absence of a direct measure of the output gap, the level specification for $\bar{x}_t$ would represent a reasonable approximation of the true output gap if the low frequency properties of productivity, and therefore of output, were well described by a linear trend. We can then think of this specification as a description of how policymakers constructed their estimate of potential output in the sixties and seventies, a period in which output was widely believed to be trend stationary (see for example Nelson and Plosser (1982) and the references therein). Note however that even if trend productivity were in fact linear, $\bar{x}_t$ would overstate the high frequency movements in x_t, since the latter reflects deviations of y_t from the actual level of productivity, and not only from its trend. Whether this discrepancy is relevant from a welfare perspective is not among the objects of our current investigation.[12]

Alternatively, we also consider a definition of the output gap that explicitly recognizes the integrated nature of the productivity process, and therefore of potential output. Under this "first difference specification", $x_t^{CB} = \overline{\Delta x}_t \equiv \Delta y_t - \overline{\Delta y}_t$, where

$$\overline{\Delta y}_t = \overline{\Delta y}_{t-1} + k \left(\Delta y_{t-1} - \overline{\Delta y}_{t-1} \right) \tag{7}$$

is an "average" growth rate computed recursively as a distributed lag on past observed growth rates. This specification has similar features to the nominal income growth rate rules advocated by Orphanides (2003), and can be thought of as a first differenced version of the level specification. Heuristically, this specification is desirable because it prevents autocorrelated errors in the measurement of the output gap from accumulating, as it is instead the case in the level specification. A bit more precisely, its main advantage is that it imposes on output the true unit root restriction, rather than the false trend stationary representation that underlies the level specification. Of course, the fact that, as we will see, this rule performs "well" given the particular shock realizations in a given historical period, is only suggestive of its normative content, and should be subject to more careful scrutiny before establishing its broader theoretical desirability.

The second non standard feature of interest rate rule (6) is its intercept. In the benchmark specification, we define $\phi_c \equiv \gamma * 1_{\{\phi_n \neq 0\}}$, so that the intercept is equal to γ when $\phi_n = 0$. This assumption "centers" the rule around the value of the interest rate that is compatible with a

[11] See Evans and Honkapohja (2001) for a general introduction to least squares learning and Sargent (1999), Noah Williams and Sargent (2002), Orphanides and Williams (2002) and Bullard and Eusepi (2003) for some recent applications of constant gain learning to monetary policy.

[12] Tambalotti (2002) presents a more careful analysis of the welfare consequences of different monetary policies in the context of this model.

constant growth rate of output γ, the point around which we log-linearize the model. In some of the experiments however, we will consider the possibility that policy makers try to track the evolution of the natural rate of interest, setting $\phi_n = 1$ (see for example Nelson and Nikolov, 2002). In the model, given the absence of stochastic shocks other than productivity, the true natural rate is simply the two-steps ahead forecast of productivity growth, $r_t^n \equiv E_{t-1}\gamma_{t+1}^a$. As above, we assume that its value is estimated by the policy authority according to the constant gain algorithm

$$\bar{r}_{t+1}^n = k\gamma_t^a + (1-k)\,\bar{r}_t^n$$

This formulation is particularly interesting in the context of a non-stationary model since it allows the central bank to adjust its "target" interest rate in response to low-frequency movements in productivity growth, which translate into persistent changes in the equilibrium interest rate.

Besides those regarding the specification of the arguments of the policy rule, we will also conduct a series of experiments inspired by CGG's hypothesis. According to their empirical findings, there is a significant difference in the coefficients of an interest rate rule similar to (6) when estimated on the Pre-Volcker and the Volcker-Greenspan periods. Moreover, in the Pre-Volcker period, the estimated rule is found not to conform with the Taylor principle, with an estimated coefficient on inflation significantly below one. CGG argue therefore that the macroeconomic instability observed in the seventies was the result of sunspot fluctuations associated with the indeterminacy of rational expectations equilibrium induced by such a policy rule. Even though a detailed investigation of this hypothesis in the context of our model is not within the scope of this paper, it is interesting to compare the explanatory power of the "Orphanides hypothesis", with a weak version of CGG's story, in which a suboptimal (though not outright unstable) choice of coefficients for the policy rule might be the culprit of the seventies' instability.[13] To study this hypothesis then, we will compare the macroeconomic outcomes associated with four main parameter configurations for the interest rate rule, chosen to match CGG's and Orphanides's (2001)estimates, as illustrated in table 1. Note from the table that the only departure from those values, which reflects our choice to avoid discussing issues of multiplicity of equilibria, is in the calibration of CGG's Pre-Volcker rule, in which we substituted the unstable estimated coefficient 0.83 with the stable 1.1.

A further important consideration on the specification of the interest rate rule regards the inflation "target" π^*, which we normalize to zero throughout all our experiments. In the simulations, this is accounted for by comparing the model with demeaned data.[14] To understand this choice, start from observing that, with a vertical Phillips curve, average inflation —which need not coincide with steady state inflation— is equal to π^*.[15] This implies that, as long as we are free to choose a value for the inflation "target", the model does not impose any restriction on average inflation.

[13] Lubik and Schorfheide (2002) recently provided the tools to carry out this analysis.

[14] Following the same convention applied to productivity growth, we define the "mean" level of inflation as its average over the period preceding that of the simulations.

[15] As long as $\beta \neq 1$, the model's long run Phillips curve is in fact not vertical, so that average inflation is not exactly equal to the "target". However, this effect is quantitatively small and can be safely ignored for the purpose of this argument.

Demeaning the data is then equivalent to identifying their mean as the policy "target". Of course, we could always treat π^* as the other parameters of the model and identify it instead with its empirical counterpart, as estimated for example by CGG. However, since the inflation target is not identified in (6) separately from the rest of the intercept, this approach requires first to fix a value for the interest rate "target".[16] But, as explained above, we are in fact interested in considering changes in this target, and would therefore rather fix π^*. This is then accomplished by simply identifying π^* with "average" inflation and expressing actual and simulated data as deviations from that value.[17]

Given the forward looking nature of the model's private sector, describing the way in which the monetary authority sets the interest rate in every period is of course not sufficient to completely characterize the equilibrium. From equation (1) we can in fact observe that the level of demand depends not only on current interest rates, but also on the public's expectations of their entire future path. Moreover, to the extent that changes in demand are reflected in changes in the output gap, future expected interest rates will also be among the determinants of current prices and wages. In a standard rational expectations setting, forecasting future interest rates simply requires model-consistent forecasts of the arguments of the policy rule, inflation and the output gap. Under our specification though, the central bank in general does not respond to model-consistent measurements of the output gap, but rather to measurements that are based on an approximate statistical model of potential output. This then raises the question of what is the model of central bank's behavior adopted by the private sector. One possibility would be to assume that private agents not only understand the way in which the bank estimates trend output, but are also sophisticated about the effect of their present and future behavior on the bank's future estimates. The simpler alternative adopted in this paper is instead to assume, as is standard in the learning literature (e.g. Evans and Honkapohja, 2001), that agents naively project the current estimated trend (or constant growth rate) into the future, as their forecast of the bank's future estimate of potential output. More formally, we write $E_t\bar{y}_{t+s} = a_t + b_t\,(t+s)$ and $E_t\overline{\Delta y}_{t+s} = \overline{\Delta y}_t$, where a_t and b_t are the trend coefficients estimated with data up to time $t-1$. Another reason for choosing this specification of agents' forecasting behavior is that it is arguably very close to the way in which trend-stationary processes (with i.i.d. errors) would be forecasted in practice, by first estimating the coefficients on the available sample and then projecting forward the resulting fitted trend. An important question which is neglected by this approach is to distinguish the effect of the current setting of the policy instrument from that of agents' expectations on its future path. We leave the investigation of this question for future research.

A final important remark about the representation of monetary policy adopted in this paper

[16] CGG choose the average real rate over the entire post-WWII sample, so that in fact their estimate of π^* should correspond to average inflation.

[17] Yet another alternative would be to exploit the degree of freedom provided by π^* and choose it so that the simulations match the average inflation in the data. However, since our experiments cover different time periods and different specifications of the interest rate rule, this would result in widely different values of π^* across experiments, obscuring some important differences across rule specifications.

is that all the policy rules described so far are in a sense entirely arbitrary, and were suggested by positive considerations alone. The main objective of the exercise is in fact just to find that particular representation of policy that, together with the model's description of private behavior, and the observed sequence of productivity growth rates, can best account for the observed evolution of inflation and real activity during a particular historical period. For example, given observations on output and productivity, which should be readily available to them according to the model, policymakers could devise "better" policies than the ones they are here assumed to be following. For example, they could choose to make their interest rate feedback rule a function of the true output gap, rather than of the ad-hoc measures described above. Nevertheless, this choice would not make the rule any less arbitrary than the ones proposed here, in the sense that such a rule would still be far from optimal according to the model's welfare criterion (see Tambalotti, 2002). Following this approach then, we are entirely disregarding why, in particular historical periods, policymakers might have chosen to follow certain policy rules—or rather a complex behavior that we can usefully summarize through a feedback rule for interest rates. Nevertheless, as we hope the rest of the paper demonstrates, it is useful to reconstruct ex-post what policy behavior might have lead to certain notable historical outcomes, and what alternative behaviors might have possibly produced more desirable outcomes than those in fact observed.

2.3 Calibration and Simulation Strategy

The calibration of the model's parameters follows Tambalotti (2002), to which we refer the reader for more details. Here, we restrict the discussion to the values chosen for the gains determining the speed with which the private sector and the central bank update their estimates of trend productivity and trend output respectively, K and k. The Kalman gain K in equation (5) is directly derived from Roberts' (2001) estimates of an hidden state model for labor productivity of the same form as (3). He finds a signal to noise ratio of 0.035, which results in a value of $K = 0.08$. Experiments with values for K between 0.05 and 0.2 did not produce any significant change in the results reported below. As for k, the constant gain in the central bank's recursive least squares algorithm, we chose a benchmark value of 0.1, mainly because it produced a "reasonably" fast decline in the estimate of the slope of the estimated output trend following the seventies' slowdown. Given the crucial importance of this parameter for the results, section 8 presents a fairly broad set of robustness checks with respect to changes in its value.

We conclude this section with a detailed description of the simulation strategy adopted in this paper. As already anticipated, the input of the simulations is the sequence of labor productivity growth rates observed during the period of interest, as measured by the Bureau of Labor Statistics.[18] These growth rates are filtered by the agents in the model according to equation (5), in order to

[18] In this exercise, we are employing the latest available vintage of productivity data. We are therefore neglecting the possibly significant effect of contemporaneous mismeasurement of productivity, an issue that would certainly merit further investigation.

obtain an estimate of the current and future expected state of trend productivity, $\gamma_{t|t}$. Given this estimate, and the current and forecasted stance of monetary policy, agents take decisions on prices, wages and the level of output, according to equations (1), (2a) and (2b). These values, together with the new central bank's estimates of trend output, will then produce, through the monetary policy rule, a new value for the interest rate. The sequence of inflation rates, growth rates of output and interest rates generated by this recursion is reported in the graphs below. Note that this strategy is fairly different from that followed in the case studies reported in Tambalotti (2002). There we generated the simulated paths for the variables of interest as impulse responses to a one time shock to the trend growth rate of productivity, assumed to decay as an AR(1) process. Here instead we are feeding the model with the observed sequence of productivity growth rates, which of course includes not only the hypothesized initial shock to trend productivity (say the 1973 "slowdown" in the case of the great inflation), but also all subsequent persistent and transitory shocks. In order to smooth some of the resulting higher frequency gyrations in the simulated sequences, we report year over year values of the growth rates and the corresponding four quarter moving average of the Federal Funds rate. These transformations are applied to both the data and the simulated output.

3 The Great Stagflation: 1973 to 1978

We start our analysis where we had left it off in Tambalotti (2002), trying to account for the evolution of inflation and real activity in the twenty quarters between 1973:III and 1978:II. The choice of this particular period is to explore the possibility that the first inflation run-up of the seventies, which peaked in the first quarter of 1975, might be explained by the concomitant productivity slowdown, whose onset is traditionally dated sometime in 1973. As a first step in the analysis, this section limits the time window of the simulation to only five years, thereby isolating the first episode of truly high inflation of the period. The next two sections will then broaden the horizon to include the second high inflation episode of the late seventies and the disinflation of the early eighties, as well as the less dramatic but steady increase of inflation that started in the mid 1960s.

For purposes of comparison with Tambalotti (2002), figure 1 reports impulse responses of inflation and real activity to a one time shock to trend productivity. As in the exercises conducted in that paper, the shock is assumed to decay as an AR(1) process with a coefficient of 0.93, and is calibrated to match the long run impact on the level of GDP found in the data. Monetary policy follows the Volcker-Greenspan rule estimated by CGG (see Table 1), with the output gap correctly measured by x_t. The only difference with figure 7 Tambalotti (2002) is then the fact that the model of trend productivity on which the private sector now bases its forecasts is that in equation (4), rather than the correct stationary model. As it is clear from the picture, this assumption alone does not have a significant impact on the resulting simulations. Inflation in particular, although rising slightly for the first half of the sample, is never anywhere near the data.

Several features of the monetary policy rule behind figure 1 might explain this negative result.

Two in particular will be investigated in what follows, the way in which the central bank is assumed to measure the output gap and the value of the rule's coefficients. Before addressing this issue though, we should verify that the change in the simulation strategy adopted here might not alone be able to improve on this result. Figure 2 therefore presents the evolution of the two variables of interest, under the same Volcker-Greenspan rule adopted above, but now using as input in the simulations the actual values of labor productivity growth, as described in the previous section. As we can see, the new simulation approach narrows significantly the discrepancy between the simulated output growth rate and the data. This is not too surprising though, since the model's production function is of the form $Y_t = A_t f(h_t)$, where h_t denotes hours. Given that equilibrium hours are fairly smooth in the model, as they are in the data, fluctuations in productivity are almost directly transmitted to output. Nevertheless, this represents a reassuring check on the model. For this reason however, we will usually de-emphasize the performance of the model in replicating the data on output, focusing instead on its ability to match inflation. In this respect, the use of the actual productivity series moves simulated inflation closer to the data, but only by a negligible amount. Still, on a qualitative level, inflation starts to display some of the features of the data, increasing in the first few quarters to reach a peak around 1975:I, then declining to a through sometime in 1976 and recovering again towards the end of the sample. This in turn suggests that movements in productivity might in fact help explain changes in inflation, but that the transmission mechanism embedded in figure 2 is too weak. In the rest of the paper we will investigate if changes in the way monetary policy responds to economic conditions can contribute to amplify those movements.

Output Gap Mismeasurement Our first experiment on the effect of the monetary rule specification is inspired by the work of Orphanides. The central thesis of his work is that the undesirable policy outcomes of the seventies can be ascribed to policymakers' overly optimistic reliance on imperfect measures of the output gap as guides for setting interest rates. To explore this possibility, figure 3 reports simulations for the case of a central bank that responds to inflation and deviations of output from an estimated linear trend, with the coefficients estimated by Orphanides (2001) for the Pre-Volcker period. As argued in section 2.2 above, this formulation has a clear positive content, in the sense of well approximating the way in which policymakers actually computed their measure of the output gap in real time. Moreover, it captures Orphanides' notion that this measure might be particularly misleading around periods of significant structural change, which in our framework takes the form of an unusually "big" shock to trend productivity.

From the picture it is evident that, by responding to a mismeasured output gap, the central bank contributes to bringing about much higher levels of inflation than under CGG's rule. In fact, inflation surges now to around 10% in the middle of 1975, partially retrenches to a trough in 1976:IV and starts surging again towards the end of the sample, remaining consistently above the data for all but the first year of the simulation. Note also that this dramatic change in the inflation profile does not imply any major change in the simulated path of output, which still experiences a deep and prolonged recession, with a negative growth rate of almost 10% at the end of 1974. These results then strikingly confirm the potentially dire consequences of a monetary policy that relies on an imprecise measure of the output gap to guide its interest rate policy.

To provide a quantitative sense of the extent of this mismeasurement, and of its consequences for the stance of monetary policy, figure 4 reports data and simulations on the Federal Funds rate and on the true and estimated output gap, x_t and $\bar{x}_t$ respectively. This comparison clearly illustrates the significant and persistent discrepancy between the two measures of the output gap, with the central bank's estimate dipping to values around -8% for the entire year 1975, in correspondence with true values of only -2%. As in the data, the central bank's real time perception of a much deeper recession than that implied by the true output gap does not translate into a reduction of the nominal interest rates, due to the concomitant surge in inflation. In fact, the simulated Fed funds rate increases to a level very close to that observed in the data, even though with a delay of about one year. It then partially retrenches in the second half of the period, following the decrease in inflation after 1975, though maintaining a positive differential with the data of around 2%. Consider however that, during the same period, the model also predicts higher inflation than what observed, so that in real terms predicted and actual interest rates are in fact fairly close.

It is interesting to compare the estimated output gaps endogenously generated by the model, with the real time output gaps computed by the Congressional Budget Office and reported by Orphanides (2000b) in his figure 4. According to the CBO's calculations, the output gap first turned negative around 1970, and returned to positive values only around 1988—ready to dip back into negative territory in correspondence with the recession of the early nineties! Even more striking, at the trough of the recession in 1975, the CBO estimated an output gap of -15%, which was then revised all the way back to around -5% on the basis of the final data reported by Orphanides (2000b). Given the obviously scant quantitative reliability of these estimates, it is hard to draw any firm conclusion from comparing them with those generated by the model. The best we can do is then simply to note that the latter are safely within the "confidence band" generated by the CBO's real time and revised estimates. This provides at least some evidence in favor of the hypothesis that the model's definition of the output gap is a reasonable approximation to some of the real time measures commonly referred to in the literature. Moreover, the fact that the model's output gap, even if providing a considerably more pessimistic picture than the model's truth, or than the CBO's revised estimates, is still well above the CBO's real time estimates, might explain the negative interest rate differential between the data and the simulations. If in fact, as argued by Orphanides (2000b), Arthur Burns' Federal Reserve did respond to assessments of the depth of the recession as pessimistic as those of the CBO, this could account for the difference between the observed drop in the nominal interest rate of approximately 5% during the year 1975, and the more moderate and delayed policy loosening generated by the model.

Before proceeding to the examination of alternative theories of the great inflation, it is worth checking if the dramatic surge in inflation registered under Orphanides' policy rule might not simply depend from the difference between his estimated coefficients and those used in figure 2, rather than from the mismeasurement of the output gap built into $\bar{x}_t$. Note however that according to Orphanides' hypothesis, the source of the seventies' policy mistake is an excessively pessimistic view of real economic conditions, and therefore of the output gap. This translates into coefficient estimates, even for the pre-Volcker period, that are in fact not very far from those found by CGG for the last two decades, so that we would in fact not expect this difference to significantly affect the results in figure 2. This conjecture is indeed confirmed by figure 5. If anything, when Orphanides' estimated coefficients are coupled with a correctly specified output gap, this produces even less inflation than under CGG's Volcker-Greenspan rule, with even a fair amount of deflation in the second half of the sample.

Changing the Rule's Coefficients A possible alternative to Orphanides' account of the great inflation is that put forth, among others, by CGG. They claim that the "great inflation" might be explained by the fact that the Federal Reserve under Arthur Burns adopted a flawed interest rate rule, which did not raise nominal rates enough in response to surges in inflation. As already pointed out, it is not among the objectives of this paper to fully investigate CGG's hypothesis, which ultimately ascribes the high inflation of the seventies to the sunspot fluctuations associated with indeterminacy of the rational expectations equilibrium. What we will do instead is to consider a configuration of the rule coefficients that is identical to that estimated by CGG on the Pre-Volcker period, except for the response to inflation, which we (arbitrarily) assume to be equal to 1.1, and therefore within the range of values that assure determinacy of the rational expectations equilibrium. This is meant to capture the idea that, even in the absence of sunspot fluctuations, a weak response of interest rates to inflation is suboptimal and could potentially give rise to episodes of high inflation. Our experiments with this rule are then devised to verify if, in light of our model, the "great inflation" of the seventies could be viewed as one such episode.

The results of this experiment are displayed in figure 6. Quite surprisingly, the model predicts a sizable drop in inflation, accompanied by an even sharper recession than before, and in general a visible deterioration of the GDP growth predictions. The interesting explanation for this puzzling behavior lies in the role of the intercept of the interest rate rule. Under the baseline specification, this intercept coincides with the pre-sample average of productivity growth, which is equal to the steady state equilibrium interest rate in a stationary model. In our economy however, the natural rate of interest, which is equal to the private sector's forecast of future productivity growth, tends to decrease because of the observed slowdown in productivity, which is slowly built into agent's forecasts. This in turn imparts a contractionary bias to monetary policy, which is now "targeting" too high an equilibrium rate, and one that the private sector does not expect to be updated. This, interacted with the week response of policy to the ensuing deflation, explains the surprising prediction of the model.

This explanation is confirmed by figure 7, in which the monetary authority is assumed to be tracking its estimate of the natural rate of interest, $\bar{r}_t^n$ (i.e. $\phi_c = 0$ in equation (6). The adjustment in the intercept of the policy rule is enough to reverse the puzzling inflation decline of figure 6. In fact, this policy rule does a pretty good job at tracking the surge in inflation in the first half of the sample, though overstating a bit the subsequent decline. Also, GDP's fit is back to similar levels as under Orphanides' rule. This suggests then that a weak response to inflation in the Taylor rule, even if still within the region of determinacy, could be behind the run up in inflation of the early seventies. According to the model though, this explanation crucially relies on the Fed having adjusted the intercept of the rule to account for the productivity slowdown and the subsequent reduction in the equilibrium real rate.

Another notable feature of figure 7, which is also related to the specification of the intercept in the policy rule, is the fact that most of the discrepancy between simulated and actual data can be ascribed to the initial conditions. In other words, if we considered the change in inflation from the beginning of the episode, we would find that the model closely replicates the data, with the possible exception of a few quarters around the beginning of 1977. This discrepancy in the first moments could of course be eliminated if we exploited the degree of freedom in the specification of the policy rule represented by the inflation "target" π^*, which we instead restricted to zero, as pointed out in section 2.2. Our preferred interpretation of this discrepancy then is that it reflects the steady increase in inflation that started around 1965, and that by 1973 had already lead it well above the average of the previous twenty-five years. This then suggests that, if we wish to explain the seventies inflation as a productivity phenomenon, we should consider the possibility that the productivity slowdown might have started much earlier than 1973, as argued for example by Bai, Lumsdaine, and Stock (1998). Section 5 explores this possibility in depth.

Before doing that though, we wish to develop a bit further our search for a positive description of the evolution of inflation in the 1970s, and possibly beyond. We start this search in the next section by broadening the horizon of our simulations to include the second spike in inflation of the late seventies and the subsequent abrupt disinflation, which has traditionally been attributed to a policy switch associated with chairman Volcker's arrival at the Federal Reserve.

4 Beyond the Seventies: Volcker's Disinflation

Given the positive performance of the model in accounting for the first peak in inflation of the 1970's, this section expands the period under consideration adding the five years between 1979 and 1983, which include both the second oil price shock and the swift disinflation of the early eighties.

Simulations of the equilibria corresponding to Orphanides' Pre-Volcker rule with output gap mismeasurement are reported in figure 8. As we can see, the model not only reproduces the two peaks in inflation of 1975 and 1981, but also accounts for the entire quantitative extent of the disinflation that followed. On a somewhat less favorable note, the model is perhaps even too "generous" in predicting high levels of inflation, especially in correspondence with the second peak of the early eighties. However, this result is fairly sensitive to the choice of the coefficient on the output gap in the policy rule. For example, substituting Orphanides' baseline coefficient of 0.57 with a coefficient of 0.4, closer to CGG's estimate of 0.27, and comfortably within both estimates' 95% confidence intervals, we obtain the impressive fit for inflation reported in figure **??**. This is just to remark that our baseline choice of coefficients for the rule directly follows some of the recent empirical work in the literature, and was not directed at improving the ability of the model to match the data. Still, this match could be significantly improved by exploiting the degrees of freedom represented by those coefficients, even within the boundaries of the statistical uncertainty of the point estimates.

The plausibility of Orphanides' hypothesis is further supported by figure 9, in which we report the evolution of the Federal Funds rate and of the true and estimated output gap. The model matches fairly well the dynamics of the nominal interest rate, with an initial hump around 1975 and the following slow increase towards a maximum in the early eighties. However, while the first simulated peak comes with around one year of delay with respect to that observed in the data, the second is a couple of quarters early and significantly short in quantitative terms of the more than 12% rate reached by the Federal Funds in late 1981.

This suggests in turn an interesting consideration on the mechanism bringing about the disinflation. Traditionally, this is attributed to a sudden shift in the Federal Reserve's attitude towards inflation, which is in turn associated with Paul Volcker's appointment as chairman.[19] Looking at the interest rate series alone, our model would arguably benefit from such an exogenous shift, which might help the simulated series to match the observed spike of 1981.[20] Such a spike though is not needed to reproduce the observed disinflation. In the model, this is instead mostly due to the concomitant sharp recovery in productivity, which we can see reflected in the steep surge of GDP growth, especially in the last few quarters of the sample. Note that here we are not assuming any change in the Fed's inflation target; in fact the simulations even maintain the Pre-Volcker specification of the rule's coefficients. This finding thus corroborates Orphanides' contention that, once we take into consideration the persistent measurement error in the output gap series available to the Fed in real time, there is no significant difference in the coefficients of the policy rule before and after Volcker's appointment to the Board. In fact, according to our simulations, the Pre-Volcker rule is compatible with the entire quantitative extent of the disinflation of the early eighties, as well as with the inflation experience of the 1970s.

Finally, we return to the bottom panel of figure 9 to observe the evolution of the central bank's

[19] Most existing models of the disinflation need to include an exogenous shift in the Federal Reserve's inflation "objective" to be able to capture the quantitative extent of the episode. Notable examples include Bullard and Eusepi (2003) and Lansing (2000), but see Primiceri (2003) for an exception.

[20] Adding to the simulations an unexpected shock to the level of the interest rate that lasted for only a few quarters around 1981 would most likely not alter significantly the simulated inflation series, due to the forward looking nature of the price setting process. Price setters's decisions are mainly influenced by their expectations of the future policy stance, rather than by the current level of the interest rate.

estimates of the output gap produced by the model. These estimates deteriorate consistently after reaching a local maximum in 1977, until the trough of the recession of 1980/81, which is followed by a swift recovery around the end of the sample. Once again, this evolution is broadly in line with that of the real time CBO estimates favored by Orphanides (2000b), which exhibit two deep troughs in mid-1980 and in late 1982 with negative output gaps of approximately -8% and -11% respectively, followed by a sharp recovery starting in the year 1983. Nothing in other words to make us doubt the reconstruction of the historical events proposed so far.

5 An Early Productivity Slowdown?

As pointed out in section 3, one of the less satisfactory properties of the inflation simulations is their discrepancy with the data in the first few periods of the sample. As discussed in that occasion, this discrepancy depends in part from our choice to normalize the model's average inflation to zero, and to subtract from the data series their pre-sample averages. The residual initial divergence between series and simulations then, just illustrates the well-known fact that inflation had started to rise from its stable and extremely low levels of the early sixties towards the double digit values of 1975 well before the occurrence of the productivity slowdown, at least according to its traditional dating in 1973. This in turn suggests to investigate the possibility that the steady build up of inflation in the second half of the sixties might be explained by an earlier occurrence of the productivity slowdown than normally thought, perhaps even as early as 1965. Note that this possibility is not completely ruled out by the available empirical evidence. Bai, Lumsdaine, and Stock (1998) for example, using a multivariate approach to inference on break dates, find evidence of a break in the growth rate of output, consumption and investment in 1969:I, with confidence intervals stretching as wide as 1965:I to 1973:I. Furthermore, in a univariate analysis of labor productivity, Hansen (2001) finds evidence of a possible break in its growth rate in December 1963, and none in the early seventies. Finally, in a study based on the I(1) specification for the growth rate of productivity also adopted here, French (2001) picks up the first signs of a slowdown in the smoothed series for trend productivity sometime around 1966.

On the basis of this evidence, and of the other considerations mentioned above, this section further widens the reach of our simulations to include the period between 1965 and 1973. For reasons that will be shortly apparent, we do this in two steps, limiting our investigation at first to the period 1969 to 1983. The results of this experiment are illustrated in figure 11 for the case of Orphanides' Pre-Volcker rule with mismeasured output gap. First, note that the data series for inflation starts now at 2%, reflecting the fact that its value in 1969 was much closer to the pre-sample average than that of 1973. This, together with the model's prediction of an intial value of inflation around 1.5%, closes the gap that was evident in figure 8 almost completely. Moreover, with the exception of a counterfactual peak of around 5% in 1971, the model accounts rather precisely for the two major peaks in inflation in 1975 and 1981, in terms both of their magnitude and timing.

Finally, the model continues to reproduce the entire quantitative extent of the observed disinflation of the early eighties, in fact indicating an excessively low inflation rate, 2% below the pre-sample average by the end of 1983.

The next step is then to consider figure 12, in which we have extended the period under consideration all the way back to 1965. Interestingly, while the inflation data start now virtually at the average, the model predicts an initial value of inflation of more than 2%. More in general, the hypothesis that early symptoms of the productivity slowdown might help explain the steadily rising profile of inflation does not find much support in the simulation. On the contrary, the model's ability to reproduce the movements of inflation significantly deteriorates when compared to figure 11. In the first few years of the sample the model and the data are negatively correlated, with inflation reaching a trough of around -2% in the simulation, in correspondence with the slow but steady increase displayed by the data. Moreover, the predicted peaks in the mid-seventies and early eighties are now off by more than 1%, with a disinflation that reaches -4% at the end of the sample, against an observed level of approximately +1.5%.

Our conclusion from the analysis of this section is therefore twofold. First, we provided further evidence in favor of the idea that the Federal Reserve, having relied on an imperfect measure of the output gap as an indicator of inflationary pressures in its interest rate rule, was at least in part responsible for the disappointing inflationary outcomes of the seventies. Second, we have shown that the model based on this policy rule is able to provide a satisfactory account of the inflation and output evolution for the entire period between 1969 and 1983, including Volcker's disinflation, but that its performance significantly deteriorates if we try to extend the beginning of the sample to 1965. This suggests that, while the model is compatible with the possibility that the productivity slowdown started to manifest itself as early as 1969, it cannot explain the progressive buildup of inflation observed starting in 1965.

In the next section then we will start from the assumption that the model provides a good description of the inflationary experience of the period 1969 to 1983 and proceed to draw some normative indications from the simulated performance of alternative monetary policy rules.

6 Some Normative Conjectures

As documented in the previous sections, our model provides a remarkably close account of the joint evolution of inflation and real activity in the years between 1969 and 1983, taking as input the observed fluctuations in labor productivity during that period, and given a particular representation of the rule followed by the Federal Reserve to set its interest rate. Given the model's positive performance, it is pretty natural to ask what would have happened to inflation and real activity under different specifications of the interest rate rule. Although the nature of this exercise is clearly normative, its results should be considered very cautiously, as normative conjectures, rather than conclusions, for at least two important reasons. First, even if a particular policy rule can dramatically reduce the volatility of inflation in the simulation, without a significant impact on output fluctuations, this result is conditional on one particular realization of the shocks. The same policy could therefore produce undesirable outcomes under different circumstances. Second, it is outside the scope of the present analysis to derive a welfare criterion to quantitatively assess the trade-offs between inflation and output volatility associated with different rules. Our approach consists instead in providing a graphic illustration of the counterfactual paths of inflation and output growth that would have emerged under alternative policy rules, given the observed realization of the productivity shock. These caveats notwithstanding, the exercise provides some interesting comparisons of the effectiveness of different approaches to policy in preventing the kind of inflationary developments experienced in the seventies, as clearly illustrated in the rest of the section. To better organize the exposition, we follow the template of section 3, analyzing first the role of output gap mismeasurement and then that of the interest rate rule's coefficients.

An Alternative Measure of the Output Gap As argued above, one of the key ingredients for reproducing the broad evolution of observed inflation in the seventies and early eighties is a monetary policy authority ready to react fairly decisively to its real-time perceptions of the deviation of real activity from potential. In this paper's formalization, this real-time notion of "slack" has been identified with the deviation of measured output from a continuously updated linear trend. As shown for example in the bottom panel of figure 9, this measure persistently underestimates the model's "true" output gap, i.e. the actual determinant of agents' pricing decisions, for the entire second half of the seventies, leading the monetary authority to conduct a looser policy than that granted by the actual deflationary pressures present in the economy. In our simulations, such a policy translates in inflationary dynamics very close to those observed in the data, with inflation growing as much as 10% between 1973 and 1975, as illustrated in figure 8. However, figure 5 shows that this spike in inflation all but disappears if the true output gap is substituted for the deviation from trend in the monetary policy rule, leaving us with virtually no inflation throughout the seventies. Of course, since one of the premises of this work is that the monetary authority might lack the ability to monitor the true output gap in real time, therefore resorting to some ad hoc measure of inflationary pressure in its policy rule, it is interesting to study the performance of measures other than the deviation from trend.

One such measure that has received some attention in the literature is the deviation of the growth rate of GDP from its sample average, continuously updated to reflect the possibility of breaks, as detailed in section 2.2. We can think of this specification as the first differenced version of the level specification adopted so far. Its main advantage, at least in the context of the non stationary model for productivity considered here, is to impose a true unit root restriction on the representation of output underlying the policy rule. Note that, just like with the level specification, implementation of this rule only requires policy makers to observe the growth rate of output.

The striking result of adopting this alternative notion of "slack", coupled with Orphanides'

Pre-Volcker coefficients, is illustrated in figure 13. Inflation virtually disappears from the picture, with in fact a fair amount of deflation (at least with respect to the pre-1969 average of 2.3%) taking hold towards the end of the sample. Not surprisingly, this extremely benign outcome in terms of inflation comes at the cost of somewhat wider fluctuations in real activity. For example, the recession of 1974/75 reaches a trough of around -12% of GDP growth over a year earlier, compared with an observed value of about -7% and a simulated value under the level specification of a little more than -10% (see figure 11). However, the reduction in the volatility of inflation would seem to comfortably justify the increase in output fluctuations according to most reasonable weights of these two objectives, at least conditionally on the shocks' realizations of the period under consideration. In this particular sample in fact, the standard deviation of inflation drops from 2.45% to 1.52% when moving from the level to the first difference specification of the rule, while that of output growth only increases from 3.81% to 4.53%.

The key to this improvement can be immediately identified by turning to the second panel of figure 14, from which we see that the persistent underestimation of the actual output gap that characterized the level specification has now disappeared. Note however that even under the first difference specification the measurement of the output gap is far from perfect. In fact, true and estimated output gap are now quite clearly negatively correlated. This suggests that a more detailed investigation of the welfare implications of this particular rule are necessary before we could safely recommend it as a desirable guide for monetary policy. Nevertheless, it is clear from our simulations that the transitory nature of the measurement errors committed under the difference specification, when compared with the extremely persistent errors that emerged under the level specification, results in a significant improvement in the ability of the rule to prevent episodes of high inflation, even if at the cost of somewhat more pronounced fluctuations in real activity.

Finally, we should also remark that, in equilibrium, this improvement does not require the central bank to impose any significant increase in the nominal interest rate, which is in fact observed to decline steadily in the simulation. This apparently puzzling behavior of the policy interest is explained by two facts. The first is that, given the much lower level of inflation in the simulation than in the data, the stance of policy, as measured by the real interest rate, is in fact much closer in the two scenarios than the nominal interest rate suggests. The second explanation resides instead in the effect of the policy rule specification on agents's expectations. Under the level specification, the persistent discrepancy between actual and estimated output gap affects not only the current level of the interest rate, but also agents' assessments of its future value. After observing a persistent measurement error, agents expect a similar error to be committed also in the future, anticipating therefore that policy will be looser (or tighter) than necessary not only today, but also in future periods. Given the forward looking nature of their pricing decisions, this magnifies the impact of current mistakes on demand and prices. Under the difference specification on the contrary, measurement error, if not quite i.i.d., is still far from being predictable. This implies in turn that agents expect future policy to be "correct" on average, therefore significantly limiting the impact of current measurement errors on inflation. As a consequence, under the difference specification, the same results in terms of inflation can be achieved through more limited movements in the policy instrument, thanks to the relatively favorable shift in agents' expectations.

The Role of the Coefficients After having illustrated the excellent inflation fighting performance of the difference rule, we now turn to investigating if similar results can be achieved instead by a judicious choice of the policy rule's coefficients. To this end we will compare the inflation performance of the baseline Orphanides' Pre-Volcker rule, with that of both Orphanides' and CGG's Volcker-Greenspan rules, again under the assumption that the gap is measured as the deviation of output from its linear trend. Note from table 1 that, while both alternative rules respond fairly strongly to changes in inflation, CGG's rule displays a coefficient on the output gap more than three times as large as that estimated by Orphanides with real time data. As we will immediately discover, this difference has major implications for the ability of the two rules to prevent inflation.

We start our discussion from figure 15, in which monetary policy follows Orphanides' Volcker-Greenspan policy rule, but maintaining output deviations from trend as its measure of slack. Note that the main difference between this rule and its Pre-Volcker version is not in the strength of the response to inflation, which only slightly increases from 1.64 to 1.8, but rather in the significant reduction in the output gap coefficient, which is cut approximately in half, from 0.57 to 0.27. This results in a significant reduction in the amplitude of inflation fluctuations. Even if the simulations still display a distinctive peak in 1975, and less sharp ones in the early seventies and early eigthies, mimicking the broad qualitative movements of the data, these peaks are quantitatively far from the ones produced under the Pre-Volcker rule. As already stressed however, this improved performance is not due to the elimination of sunspot fluctuations, but rather to a much more cautious reliance on the estimated output gap as a measure of inflationary pressures.

The message of the previous experiment is further strengthened by the results in figure 16, in which we report simulations of the CGG Volcker-Greenspan rule, now under the level specification for the estimated output gap. Comparing the evolution of inflation in the upper panel with that in figure 11, based on Orphanides' Pre-Volcker rule, we note that the change in the coefficients has virtually no effect on the dynamics of inflation. Even if under CGG's rule the monetary authority responds to inflation with a coefficient of 2.15, as compared with the 1.64 of Orphanides' specification, the effect of this more forceful response is neutralized by the concomitant stronger reaction to the mismeasured output gap. On net, these two opposing forces cancel out almost exactly, leaving us with an evolution of inflation which is virtually indistinguishable between the two specifications. This result is important, because it illustrates how a configuration of the rule coefficients like that estimated by CGG on the Volcker-Greenspan period, which is usually considered to reflect "good" policy, is in fact compatible, at least in the context of our model, with an evolution of inflation very close to that observed in the seventies, and universally considered "bad". Once again, the key to this result lies in the measure adopted by the central bank to gauge the intensity of inflationary pressures, and in the strength with which the bank reacts to this measure. Under CGG's rule in particular, the relatively high weight on the estimated output gap in the reaction function results in a very bad inflation performance, the high coefficient on inflation notwithstanding.

In conclusion, our normative investigation leads us to advance two conjectures. First, violation of the Taylor principle is not a necessary condition to produce levels of inflation comparable to those observed in the seventies, so that compliance with the principle is not enough to guarantee that episodes of high inflation will be avoided in the future. Second, the way in which the central bank measures the output gap, and the weight given to this measurement in the policy rule, seem to hold the key to avoiding the recurrence of such episodes, at least in the presence of significant and persistent disturbances to the growth rate of productivity, like the ones that characterized the early seventies. In this respect then, an interesting extension of our analysis is to consider the more recent period of significant changes in the growth rate of productivity, the second half of the nineties. This task is taken on in the next section.

7 The New Economy

This section shifts the focus of the analysis from the great inflation of the 1970s to the remarkable macroeconomic performance of the late nighties, that was characterized by low inflation and extraordinarily high growth and low unemployment. Many observers have attributed this performance to the advent of the so-called "new economy", whose more fundamental and by now less controversial premise was the acceleration of productivity around the middle of the decade. Much has been made, especially by the popular media, of the skills with which Chairman Greenspan has been able to steer American monetary policy during this period of structural change, even if his reputation has more recently been tarnished in part by the prolonged slowdown started in March 2001.[21] Given the attention recently captured by this episode, which unfolded in many respects as a replay of the productivity slowdown, except with the "opposite sign", it is interesting to investigate if our model can provide a satisfactory account of it, and under what assumptions on the policy rule followed by the central bank.

We begin our analysis by considering Orphanides' Volcker-Greenspan rule, with the gap measured as the deviation of output from a linear trend. The results of this experiment are reported in figure 17. Looking first at inflation, we note that the model replicates its decline throughout the end of 1998, but departs quite significantly from the data in the last two years of the sample, indicating a further decline, while the data reverse course and return to their pre-sample average at the end of 2000.[22] As for output growth, similar remarks apply, with a fit that visibly deteriorates in the second half of the sample, where we can even clearly distinguish a negative correlation between the simulation and the data.

[21] A complete journalistic account of the events of the late nineties can be found for example in Woodward (xxxx).

[22] In the simulations referring to the 1990s, all variables are measured as deviations from their average in the year 1995. See Tambalotti (2002) for further details on the rationale behind this choice.

The question then arises of whether a better match could be obtained by simply changing the coefficients of the rule. The answer provided by figure 18, in which we consider CGG's Volcker-Greenspan rule, is clearly no. In fact, the model's ability to replicate the evolution of inflation significantly deteriorates. Nevertheless, this experiment highlights an important feature of the level specification of the monetary policy rule, namely the prediction of a steadily declining inflation throughout the sample. A similar tendency was also detectable in figure 17, but it was at least in part masked by the smaller scale of the decline, and by the fact that this decline did match what observed in the data for the first few years of the sample. What figure 18 makes clear, is the extent to which, even in the case of the nineties, a central bank that adopted output's deviations from trend as a measure of slack, would have been fooled into interpreting the observed surge in output as a source of inflationary pressures, when, with the benefit of hindsight, we now know that surge was at least in part due to a partial reversal of the productivity slowdown experienced in the seventies. This interpretation of the data would have in turn spurred the central bank to tigthten its policy, causing a decline in inflation that, under CGG's rule, the model quantifies at almost 6%. What is remarkable then is that, contrary to what happened in the seventies, the Federal Reserve clearly did not follow such a misleading notion of the output gap, or at least did not react to it to the extent indicated by CGG's rule.

Having established that Greenspan's Fed is unlikely to have followed the same policy rule that led to the mistakes of the seventies, the next question on the agenda is to discover what kind of rule might provide a satisfactory account of the observed evolution of inflation in the second half of the nineties. In this investigation, we will confine ourselves to considering only one set of coefficients for the rule, those estimated by Orphanides on the Volcker-Greenspan sample, varying instead the notion of output gap to which the policy authority is assumed to react. In particular, we considered a rule that reacts to the differenced version of the estimated output gap (in figure 19), together with one rule in which, besides responding to the latter, the central bank also centers its interest rate rule around a time varying notion of the equilibrium real interest rate, as defined in section 2.2 (see figure 20). Starting with the first figure, we can immediately observe that, once monetary policy abandons output deviations from trend as its notion of the gap, the pronounced downward trend in inflation found in figures 17 and especially 18 disappears, substituted however by an also counterfactual, if less pronounced, upward trend. This suggests then that a simple switch to a less biased estimation procedure of the output gap by the Fed, if closer to the truth, is not enough to provide a completely satisfactory account of the observed inflation path.

A much closer match to this path is provided instead by the rule in which, besides responding to inflation and the growth rate of output, the central bank also tries to target an estimate of the natural rate of interest, as illustrated in figure 20. This evidence, if hardly conclusive, is nevertheless very interesting in light of some recent efforts by the Federal Reserve and other central banks to devise efficient ways of measuring the natural rate of interest consistently with the New Keynesian microfoundations from which our model is also derived (see for example Laubach and

Williams, 2003; Nelson and Nikolov, 2002). What the simulations seem to suggest is that an attempt at tracking the fluctuations of the natural rate might have been already in place as a response to the shift in the growth rate of productivity associated with the new economy. What is left to investigate however is whether such an attempt has desirable welfare properties within a fully specified model of the central bank objectives.

Finally, we complete the picture on the positive performance of this rule turning to figure 21, in which we report actual and simulated interest rates and true and estimated output gaps. As we can see, the simulated Federal Funds rate matches the data very closely, with the exclusion of the temporary dip in 1999, which reflects the Federal Reserve's reaction to the Russian debt crises and the failure of LTCM, and therefore could not possibly be captured by the model. Moreover, from the lower panel of the figure we observe that, in the particular circumstances of the late nineties, the deviation of the growth rate from its average is a reasonably accurate indicator of the true output gap, devoid of the very persistent errors that were characteristic of the deviations from trend. This helps of course to explain the much better performance of monetary policy in the second half of the nineties than during the seventies. As already noted in the previous section however, the differenced series appears to be negatively correlated with the true output gap. This leaves open the question of whether this particular measure of slack might produce less desirable outcomes in periods in which other shocks than those to productivity are predominant.

8 Robustness

As pointed out in section 2.3, one of the key parameters in our simulations is the constant gain k used by the central bank to compute its estimates of potential output. This is also one of the hardest parameters to calibrate, since the only direct evidence on its value would come from real time estimates of the trend computed by the Federal Reserve, which to our knowledge are not readily available. In our benchmark calibration then, we more or less arbitrarily chose the value $k = 0.1$, as the "round" number closest to the Kalman gain $K = 0.08$ derived from Roberts's (2001) estimates.[23] The effect of this choice on the updating speed of the estimates of the trend slope following 1973 are reported in figure 22, where we compare four different values of k between 0.05 and 0.2. The figure shows that for all values of k the updating is fairly slow for the first two years of the sample. The following acceleration though, brought about by the accumulation of several quarters of low growth, is still fairly smooth for $k \leq 0.1$, but significantly more pronounced for $k \geq 0.15$. By the end of the period under consideration however, estimates of the slope are still around 2.65% for $k = .05$, and in a more reasonable 1.8% to 2.2% range for the other values of k considered. As a possible reference point, consider that Blinder (1979), writing some time in 1978/79, while expressing considerable uncertainty about the extent of the productivity slowdown,

[23] Recursive estimation of a constant growth rate (see equation (7) and filtering of the trend component in an I(2) model (see equation (5) are formally equivalent problems. Hence our treatment of the respective gains k and K as similar objects.

proposes the average growth rate between 1967 and 1973 as a reasonable measure of the slope of the new productivity trend.[24] In our data, this average is equal to 2.66%. In comparison, the model's estimates of the trend slope in 1979:I are around 2.8% with $k = 0.1$ and 2.6% with $k = 0.15$. All in all then, our benchmark value for k appears to be in line with the (scant) available evidence, but some check on the effect of alternative values for this parameter on our main conclusions is certainly in order.

Just as an example, figure 23 reports the alternative predicted paths for inflation and GDP growth under the four values of k considered above, in the benchmark case of Orphanides' pre-Volcker specification of the interest rate rule. With the exception of the $k = 0.05$ case, which predicts values of inflation around 20% in the early eighties, matching the extremely slow updating of trend growth observed above, all the other paths of inflation are reasonably close to each other, and even more to the data. In particular, a value of $k = 0.15$, which was shown to be compatible with our evidence on real time estimates of productivity trends, would if anything improve the model's ability to replicate the inflation profile, especially in the central years of the sample. Another interesting observation is that the differences in GDP growth associated with different values of k are almost nonexistent. Given the same profile for productivity, these differences feed directly into differences in hours, and therefore in the true output gap. This suggests in turn that the observed differences in the predicted path of inflation are primarily due to differences in agents' expectations about future values of the central bank's measure of the output gap, rather than to genuine differences in inflationary pressures. A more systematic investigation of this issue awaits further research.

9 Conclusions

This paper reexamined some recent accounts of the unusual behavior of the US macroeconomy in the seventies and late nineties, in light of a structural model of the relationship between inflation, productivity and monetary policy developed in Tambalotti (2002). The central feature of the model is the recognition that all economic agents, including the policy authority, do not know the exact distribution of the productivity shocks buffeting the economy. In particular, they are uncertain about their persistence. This implies that, following a persistent productivity shock, agents' optimal pricing and demand decisions give rise to transitional dynamics of the kind observed during the "great stagflation" and the "new economy", with inflation rising (falling) and output falling (rising) in response to a negative (positive) shock. While in Tambalotti (2002) we focused on the optimal policy response to such fluctuations, this paper has investigated what kind of monetary policy, if any, could account quantitatively for these historical episodes, taking as given the observed sequence of productivity growth rates.

[24] In line with our proposed reconstruction of real time estimates of trends, Blinder (1979) assumes the new productivity trend to be linear (see for example his figure 4.6).

Our analysis followed Taylor (1993) and the recent work of Clarida, Galí, and Gertler (2000) and Orphanides (2000b; 2000a; 2001; 2002) in assuming that monetary policy can be effectively summarized by an interest rate feedback rule. Within this class of policies, we compared specifications in which the central bank fully relies on an imperfect measure of the output gap as its gauge of economic "slack", as suggested by Orphanides, with specifications in which policy leans only lightly against the inflationary wind, but appropriately reacts to a model-consistent measure of the output gap, as in CGG. We did not consider however policy rules that do not conform with the "Taylor principle".

From a positive standpoint, we found that the latter policies, once coupled with our model of private behavior, have serious difficulties in accounting for the high inflation of the seventies. The former, on the contrary, successfully reproduce the two inflationary spikes of the mid-seventies and early eighties, and can also account for the Volcker disinflation that followed this second episode, irrespective of the exact configuration of the rule's parameters. The evidence is more mixed for the "new economy", even if it clearly points towards a change in the policy rule with respect to the earlier period. Our best guess is that this new rule involves targeting the growth rate of output (rather than deviations from potential), and possibly a time-varying notion of the "natural" rate of interest.

Finally, from a normative perspective, we showed that, even in the presence of uncertainty on the "true" output gap, the Federal Reserve could have avoided the inflationary outcome of the seventies following a policy rule that targeted the deviations of output growth from a continuously updated sample average. Moreover, this improvement would have come at the cost of only negligible increases in the volatility of output.

In conclusion, our analysis provides further support for Orphanides' reconstruction of the American policy history of the last thirty years, and for his advocacy of "difference rules" for monetary policy. However, much work remains to be done to strengthen these results. First, to provide a formal criterion to compare the positive performance of alternative models of monetary policy. Second, to analyze more generally the welfare properties of difference rules. Third, to verify empirically if in fact difference rules of the kind considered here provide a better characterization of policy in the nineties than more traditional "level" specifications.

References

Amato, J. D., and T. Laubach (2003): "Estimation and Control of an Optimization-Based Model with Sticky Prices and Wages," *Journal of Economic Dynamics and Control*, 27(7), 1181–1215.

Bai, J., R. L. Lumsdaine, and J. H. Stock (1998): "Testing for and Dating Common Breaks in Multivariate Time Series," *Review of Economic Studies*, 65(3), 395–432.

Barro, R. J., and D. B. Gordon (1983): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91(4), 589–610.

Barsky, R. B., and B. J. De Long (1993): "Why Does the Stock Market Fluctuate?," *Quarterly Journal of Economics*, CVIII (2), 291–311.

Blinder, A. S. (1979): *Economic Policy and the Great Stagflation.* Academic Press, New York.

Boivin, J., and M. P. Giannoni (2003): "Has Monetary Policy Become More Powerful?," NBER Working Paper No. 9459.

Bullard, J., and S. Eusepi (2003): "Did the Grat Inflation Occur Despite Policymaker Committment to a Taylor Rule?," mimeo, Federal Reserve Bank of St. Louis.

Calvo, G. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12(3), 383–98.

Christiano, L. J., and T. J. Fitzgerald (2003): "Inflation and Monetary Policy in the Twentieth Century," *Federal Reserve Bank of Chicago Economic Perspectives*, First Quarter, 22–45.

Christiano, L. J., and C. Gust (1999): "The Grat Inflation of the 1970s," mimeo, Northwestern University.

Clarida, R., J. Galí, and M. Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and some Theory," *Quarterly Journal of Economics*, 115(1), 147–180.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46(2), 281–313.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics.* Princeton University Press, Princeton, NJ.

French, M. W. (2001): "Estimating Changes in Trend Growth of Total Factor Productivity: Kalman and H-P Filters versus a Markov-Switching Framework," Federal Reserve Board Finance and Economics Discussion Series 2001-44.

Hansen, B. E. (2001): "The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity," *Journal of Economic Perspectives*, 15(4), 117–128.

Judd, J. P., and G. D. Rudebusch (1998): "Taylor's Rule and the Fed: 1970-1997," *Federal Reserve Bank of San Francisco Economic Review*, 3, 3–16.

Kydland, F. E., and E. C. Prescott (1977): "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85(3), 473–492.

Lansing, K. J. (2000): "Learning About a Shift in Trend Output: Implications for Monetary Policy and Inflation," Federal Reserve Bank of San Francisco Working Paper 00-16.

Laubach, T., and J. C. Williams (2003): "Measuring the Natural Rate of Interest," *Review of economics ans Statistics*, forthcoming.

Lubik, T. A., and F. Schorfheide (2002): "Testing for Indeterminacy: An Application to U.S. Monetary Policy," mimeo, University of Pennsylvania.

Nelson, C. R., and C. R. Plosser (1982): "Trends and Random Walks in Macroeconmic Time Series. Some Evidence and Implications," *Journal of Monetary Economics*, 10(2), 139–162.

Nelson, E., and K. Nikolov (2002): "Monetary Policy and Stagflation in the UK," CEPR Discussion Paper No. 3458.

Noah Williams, I.-K. C., and T. J. Sargent (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

Orphanides, A. (2000a): "Activist Stabilization Policy and Inflation: The Taylor Rule in the 1970s," Federal Reserve Board Finance and Economics Discussion Series 2000-13.

——— (2000b): "The Quest for Prosperity Without Inflation," European Central Bank Working Paper No. 15.

——— (2001): "Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches," Federal Reserve Board Finance and Economics Discussion Series 2001-62.

——— (2002): "Monetary Policy Rules and the Great Inflation," *American Economic Review*, 92(2), 115–120.

——— (2003): "Historical Monetary Policy Analysis and the Taylor Rule," Federal Reserve Board Finance and Economics Discussion Series 2003-36.

Orphanides, A., and J. C. Williams (2002): "Imperfect Knowledge, Inflation Expectations and Monetary Policy," Federal Reserve Bank of San Francisco Working Paper 02-04.

Primiceri, G. (2003): "Time Varying Structural Vector Autoregressions and Monetary Policy," mimeo, Princeton University.

Roberts, J. M. (2001): "Estimates of the Productivity Trend Using Time-Varying Parameter Techniques," Federal Reserve Board Finance and Economics Discussion Series 2001-8.

Rotemberg, J. J., and M. Woodford (1997): "An Optimization-Based Econometric Model for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

——— (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton.

Sims, C. A. (2001): "Stability and Instability in US Monetary Policy Behavior," mimeo, Princeton University.

Sims, C. A., and T. Zha (2002): "Macroeconomic Switching," mimeo, Princeton University.

Stock, J. H., and M. W. Watson (1998): "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model," *Journal of the American Statistical Association*, 93(441), 349–358.

Tambalotti, A. (2002): "Optimal Monetary Policy and Productivity Growth," mimeo, Princeton University.

Taylor, J. B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

——— (1999): "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*, ed. by J. B. Taylor, pp. 319–341, Chicago. University of Chicago Press.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton.

Tables and Figures

Interest Rate Rule	Pre-Volcker		Volcker-Greenspan	
Coefficients	Orphanides	CGG	Orphanides	CGG
ϕ_i	.70	.68	.79	.79
ϕ_π	1.64	.83 (1.1)	1.8	2.15
ϕ_x	.57	.27	.27	.93

Table 1: Interest rate rule coefficients. The Pre-Volcker column refers to estimates for the period 1960:I to 1979:II for CGG and 1966:I to 1979:II for Orphanides (2001). The value in brackets for CGG's Pre-Volcker ϕ_π coefficient is the one actually used in the simulations. The Volcker-Greenspan column refers to estimates for the period 1979:III to 1996:IV for CGG and 1979:III to 1995:IV for Orphanides (2001).

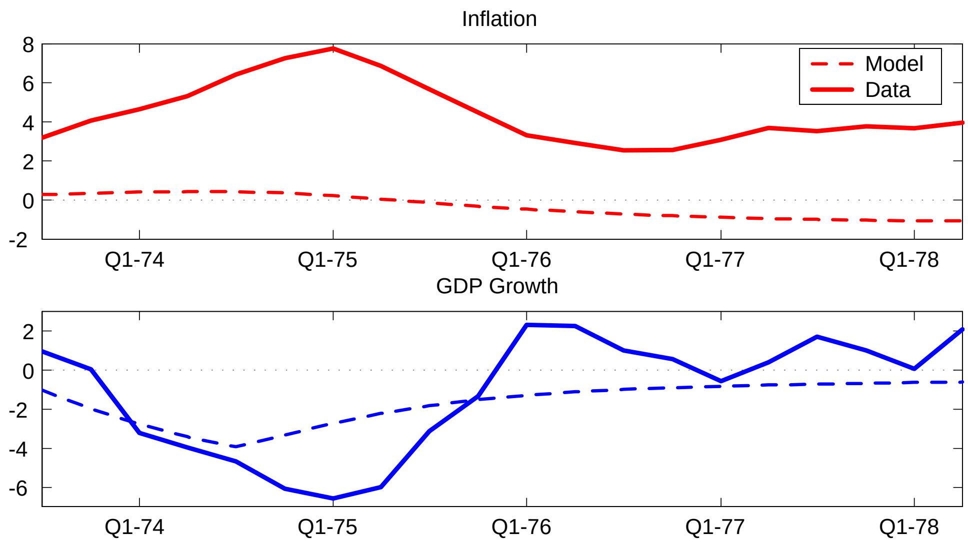

Figure 1: Impulse responses of inflation and GDP growth to a shock to "trend" productivity assumed to decay as an AR(1) process with autocorrelation coefficient 0.93. Monetary policy follows CGG's Volcker-Greenspan rule with $x_t^{CB} = x_t$.

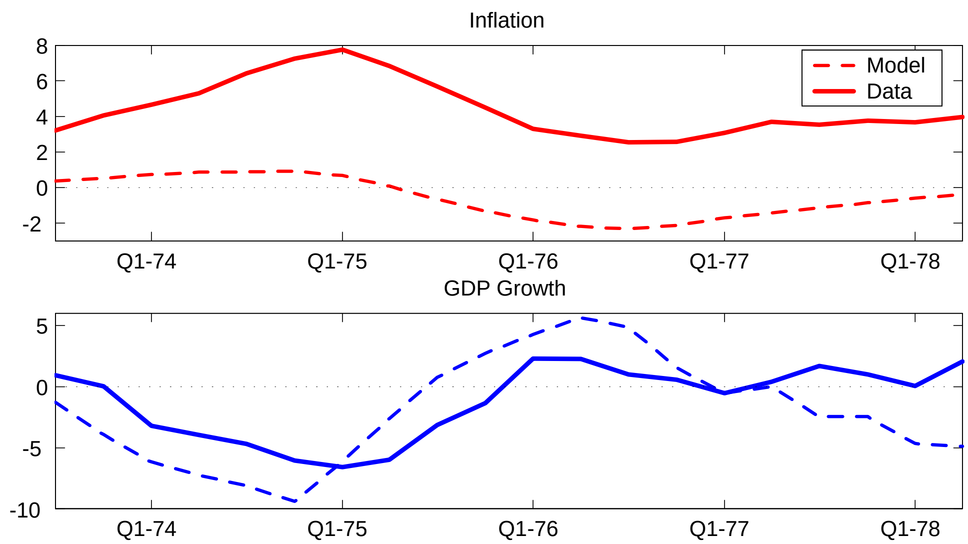

Figure 2: Simulations of inflation and GDP growth. Monetary policy follows CGG's Volcker-Greenspan rule with $x_t^{CB} = x_t$.

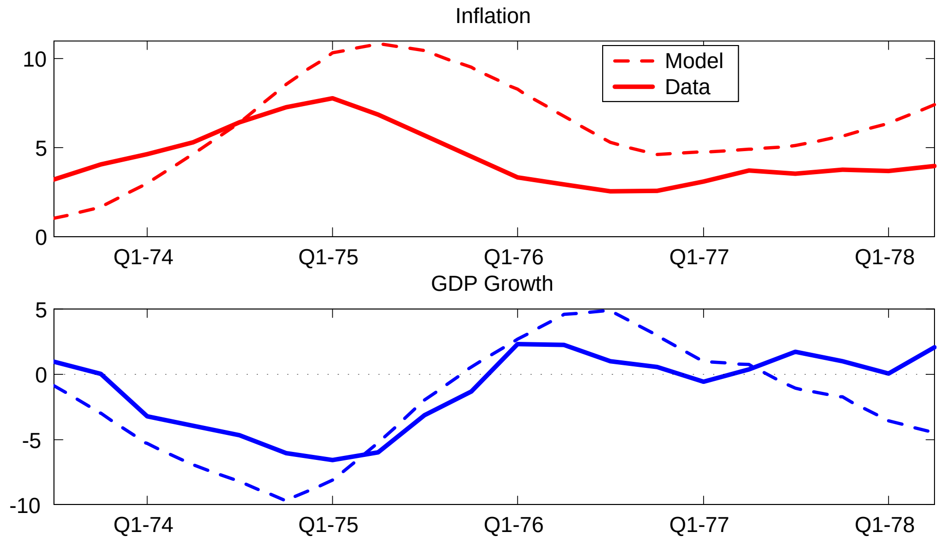

Figure 3: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

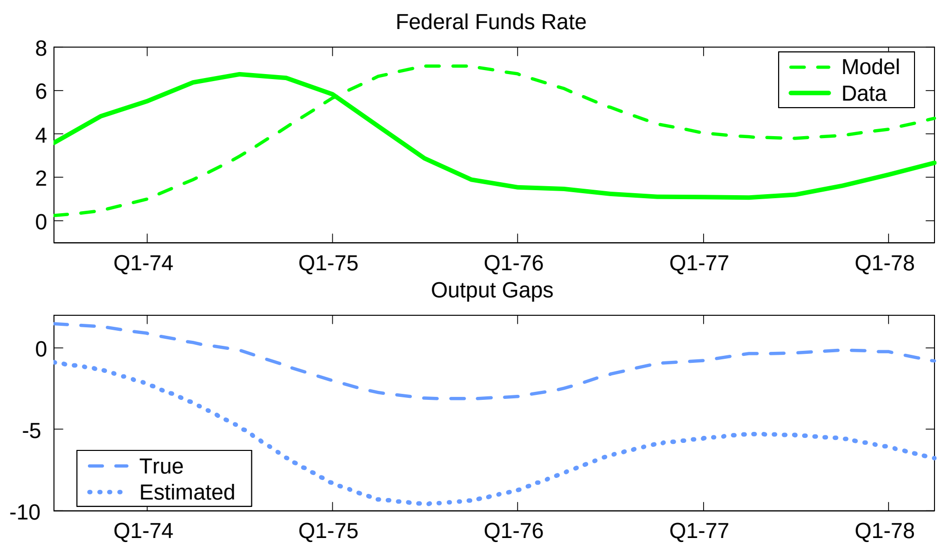

Figure 4: Simulations of federal funds rate and output gaps. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

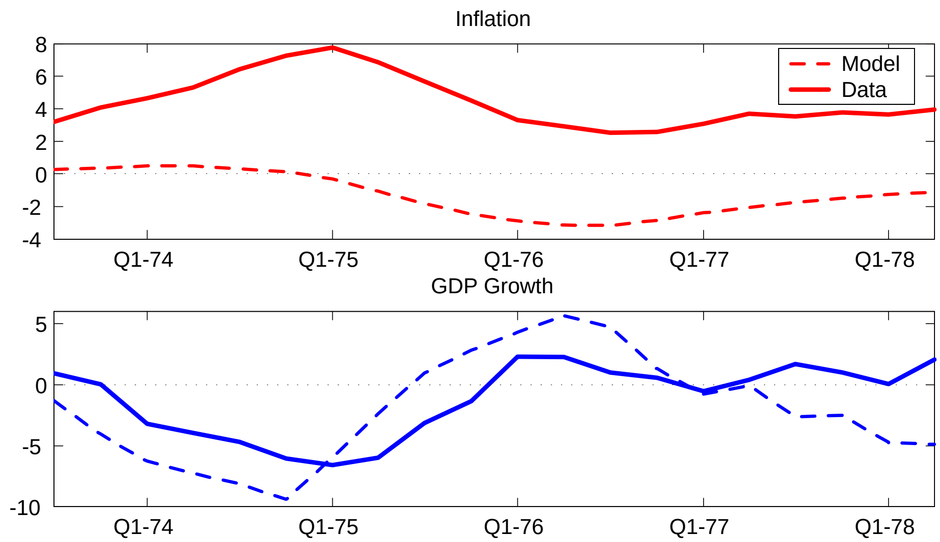

Figure 5: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = x_t$.

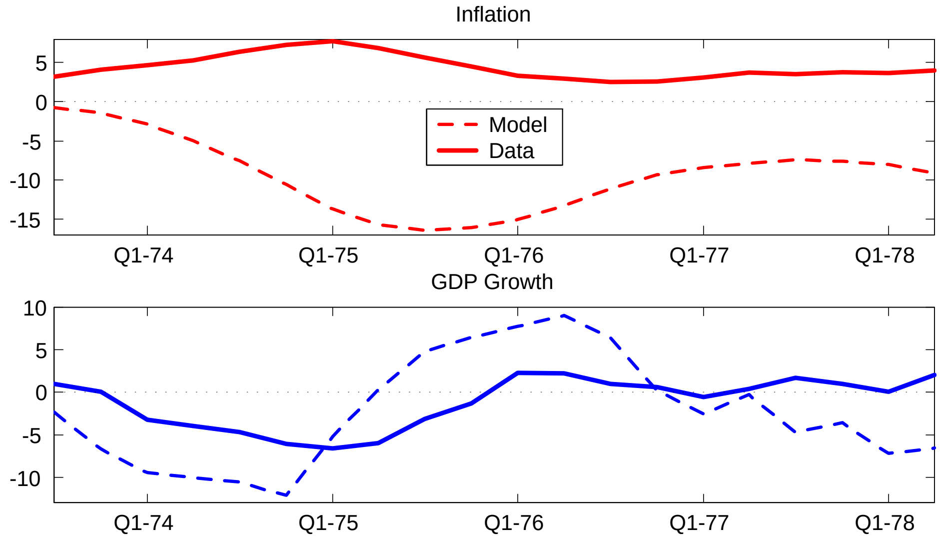

Figure 6: Simulations of inflation and GDP growth. Monetary policy follows CGG's Pre-Volcker rule with $x_t^{CB} = x_t$.

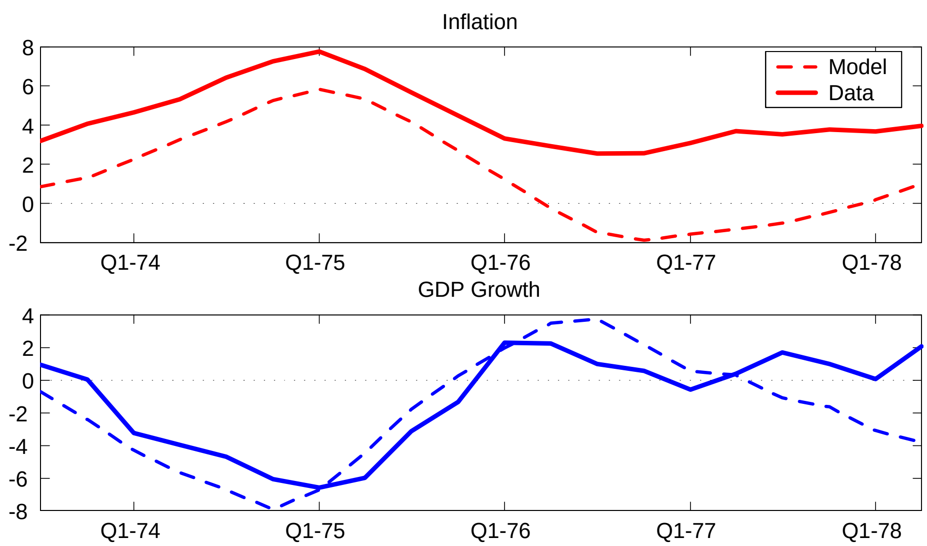

Figure 7: Simulations of inflation and GDP growth. Monetary policy follows CGG's Pre-Volcker rule with $x_t^{CB} = x_t$ and $\phi_n = 1$.

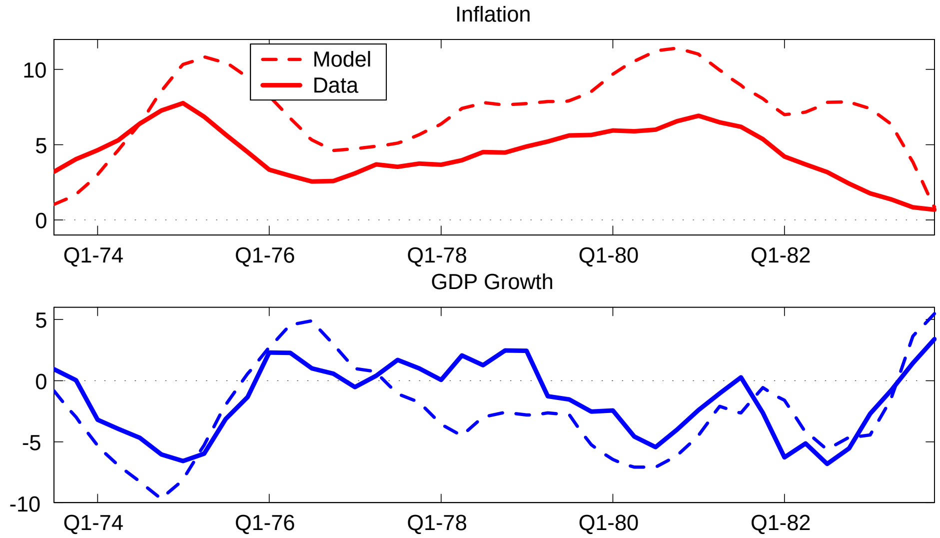

Figure 8: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

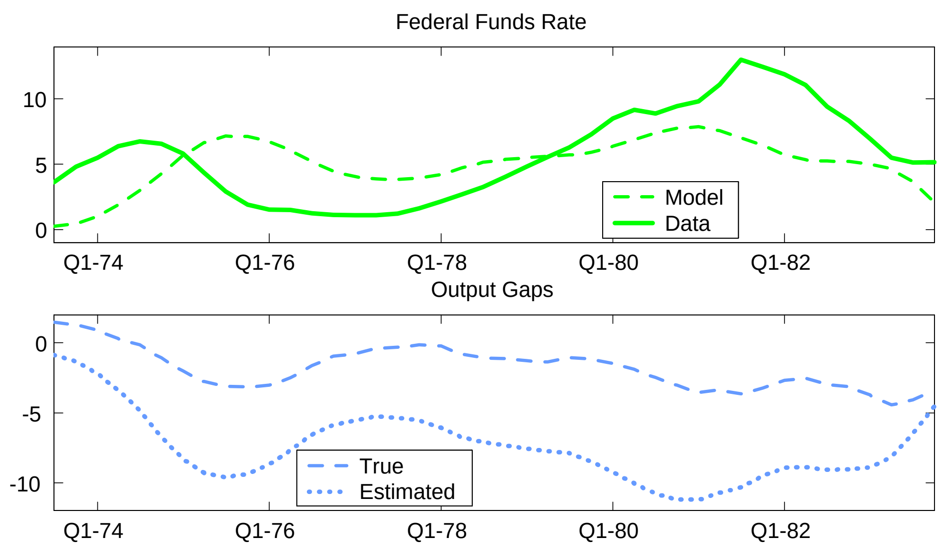

Figure 9: Simulations of federal funds rate and output gaps. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

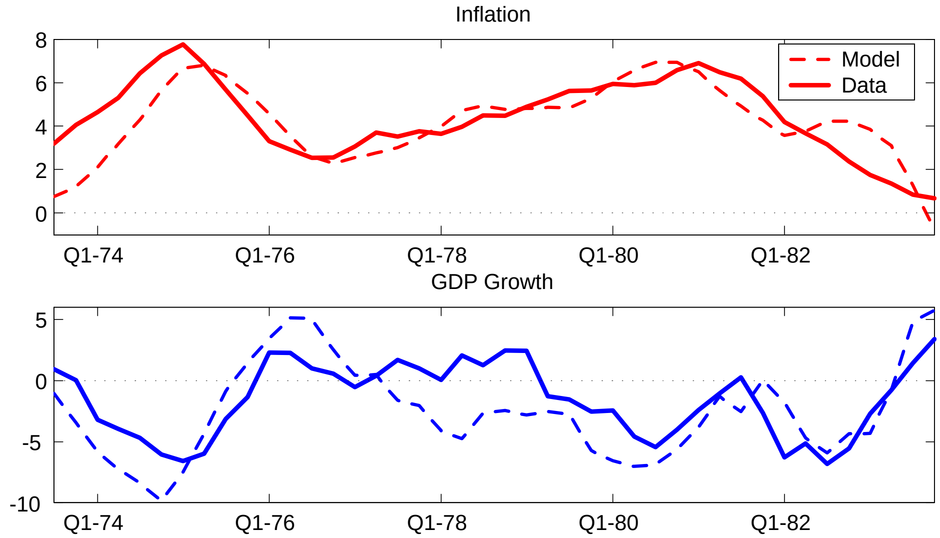

Figure 10: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $\phi_x = 0.4$ and $x_t^{CB} = \bar{x}_t$.

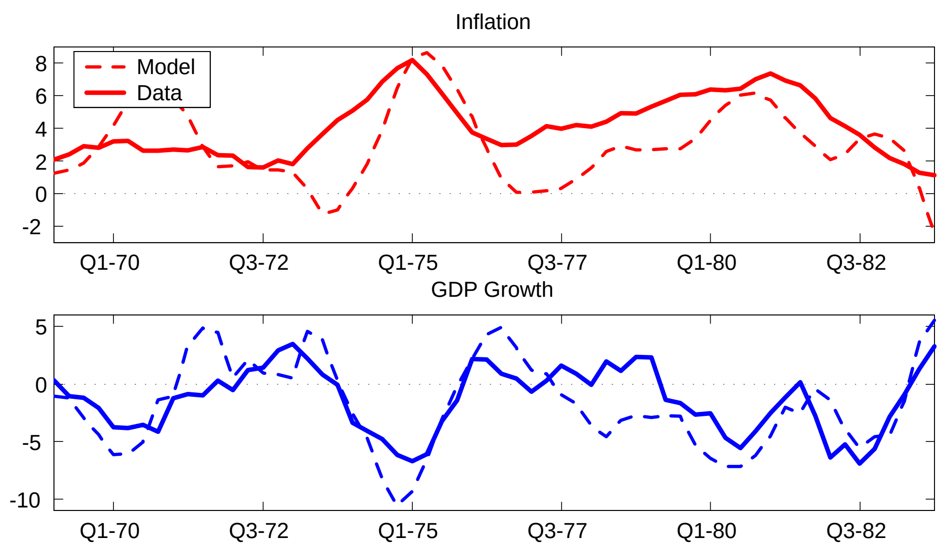

Figure 11: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

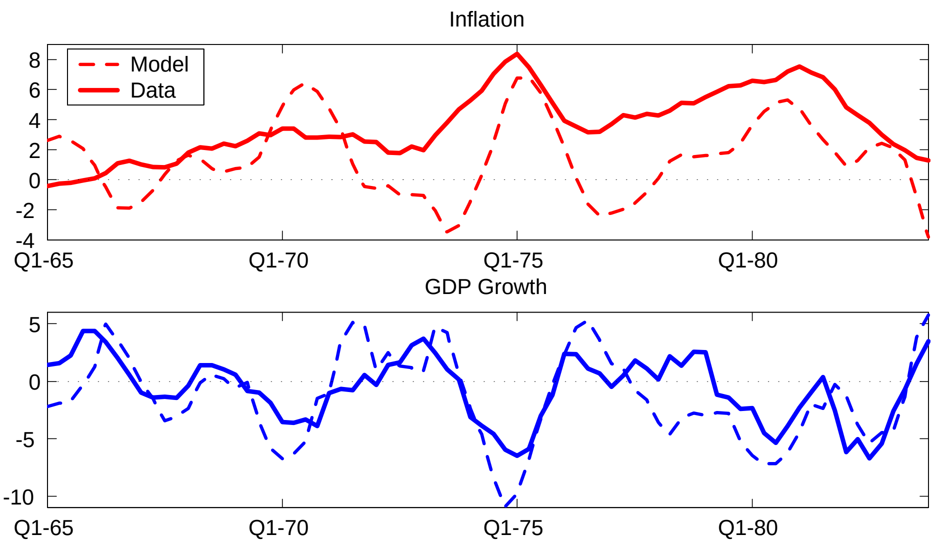

Figure 12: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

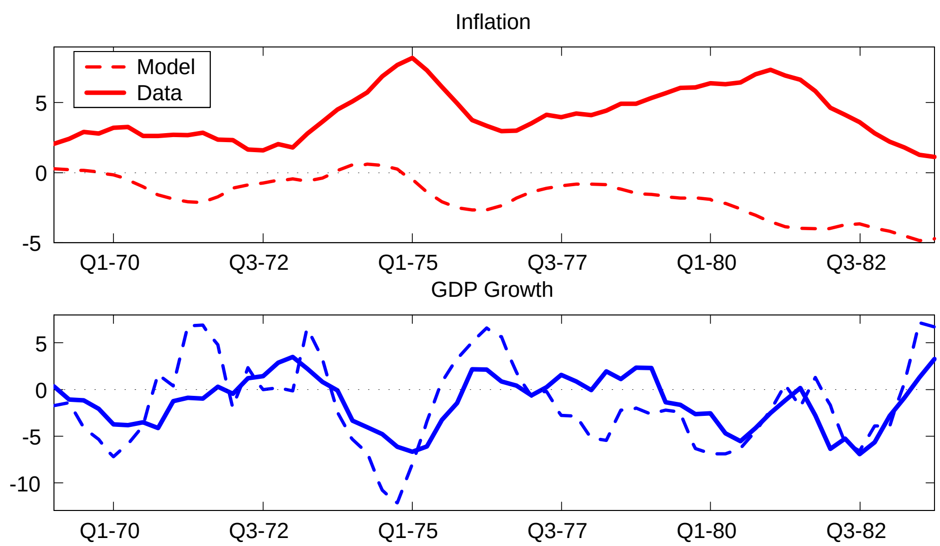

Figure 13: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \overline{\Delta x}_t$.

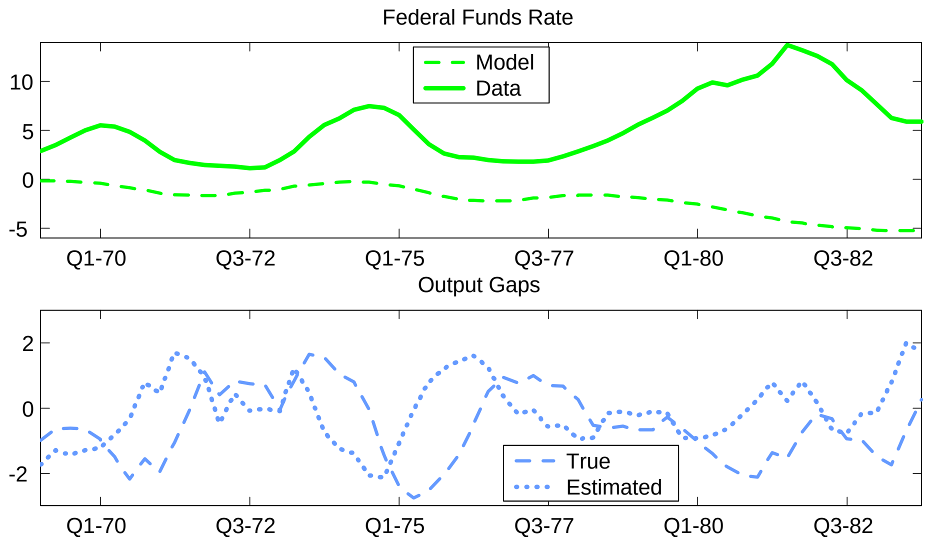

Figure 14: Simulations of federal funds rate and output gaps. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \overline{\Delta x_t}$.

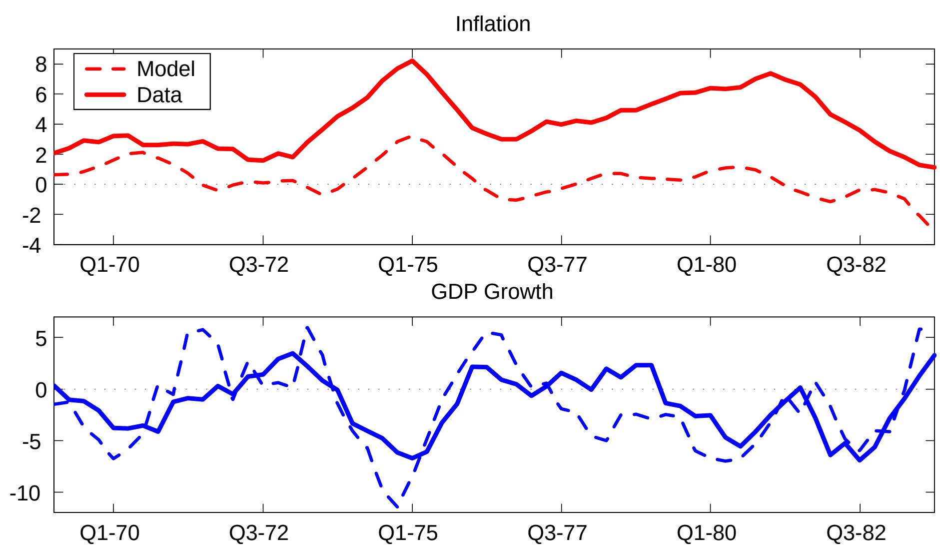

Figure 15: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Volcker-Greenspan rule with $x_t^{CB} = \bar{x}_t$.

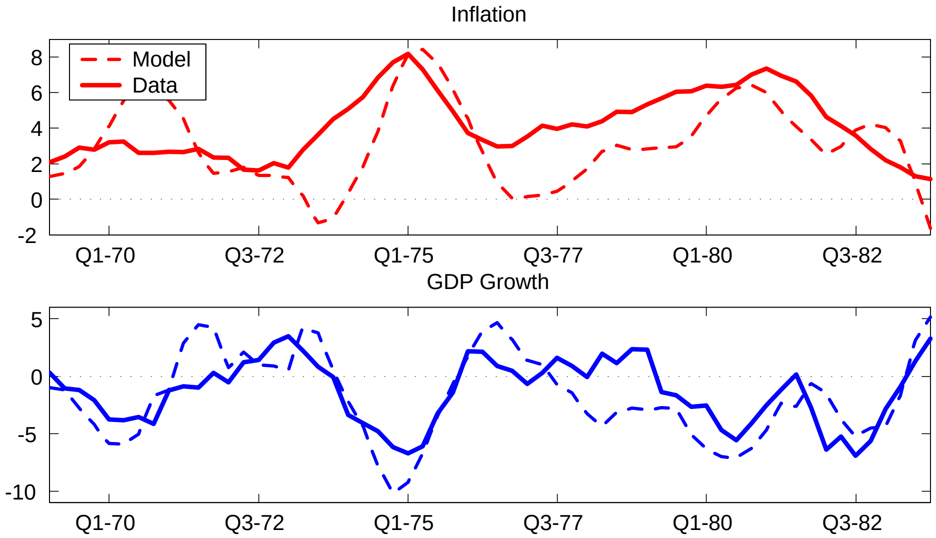

Figure 16: Simulations of inflation and GDP growth. Monetary policy follows CGG's Volcker-Greenspan rule with $x_t^{CB} = \bar{x}_t$.

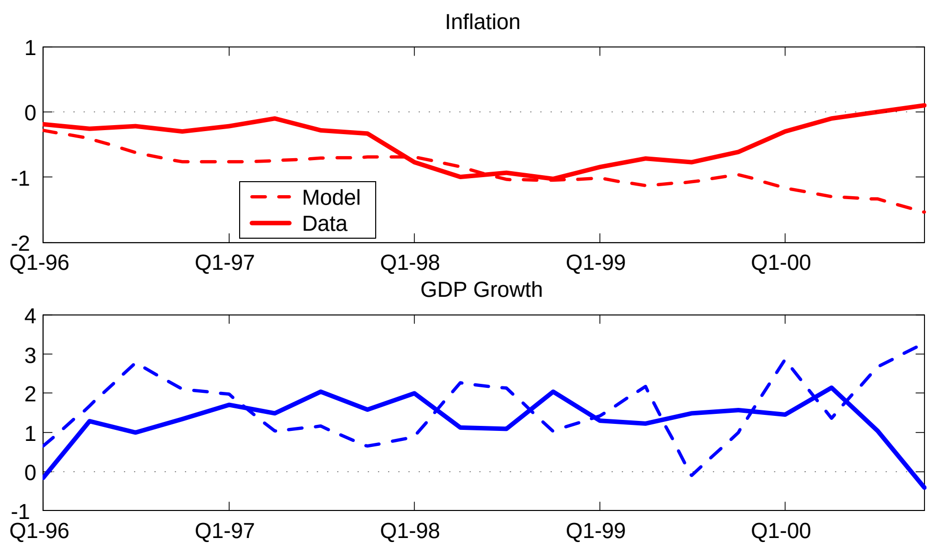

Figure 17: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Volcker-Greenspan rule with $x_t^{CB} = \bar{x}_t$.

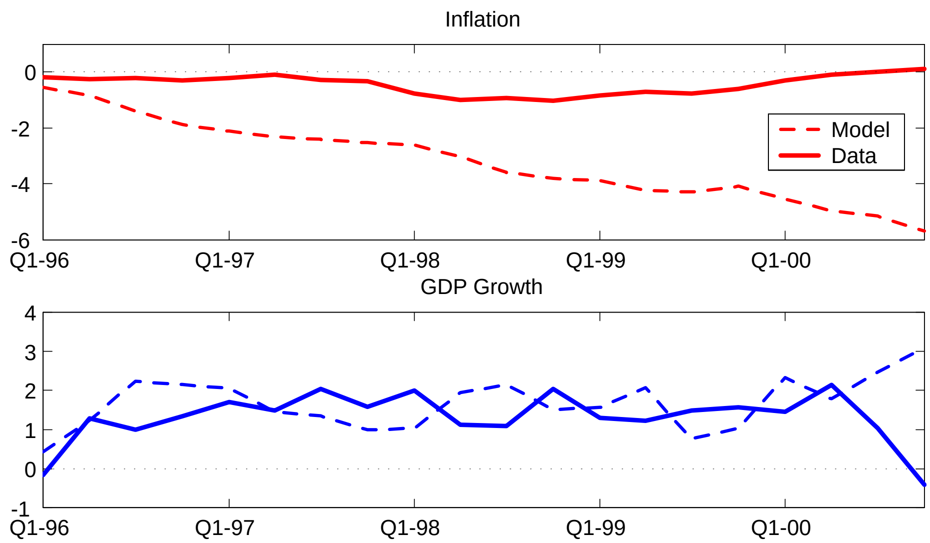

Figure 18: Simulations of inflation and GDP growth. Monetary policy follows CGG's Volcker-Greenspan rule with $x_t^{CB} = \bar{x}_t$.

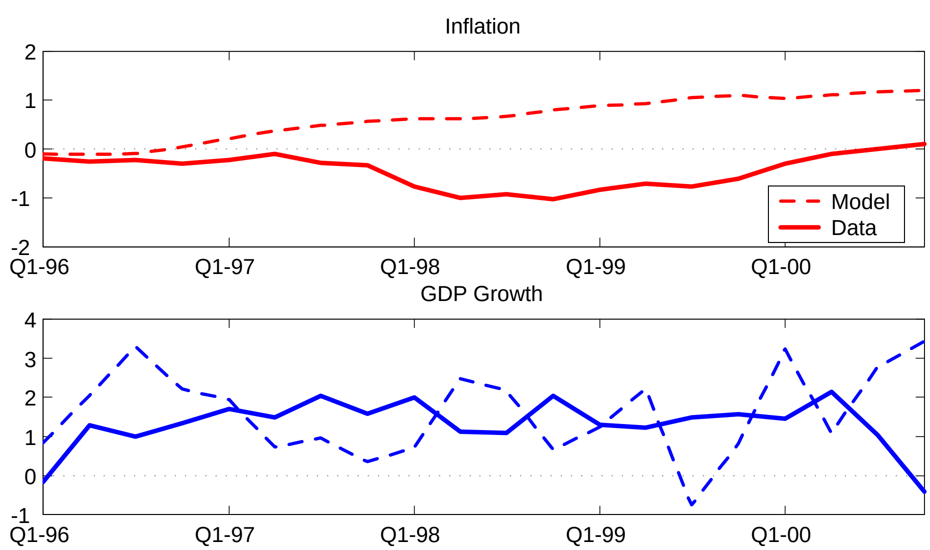

Figure 19: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Volcker-Greenspan rule with $x_t^{CB} = \overline{\Delta x}_t$.

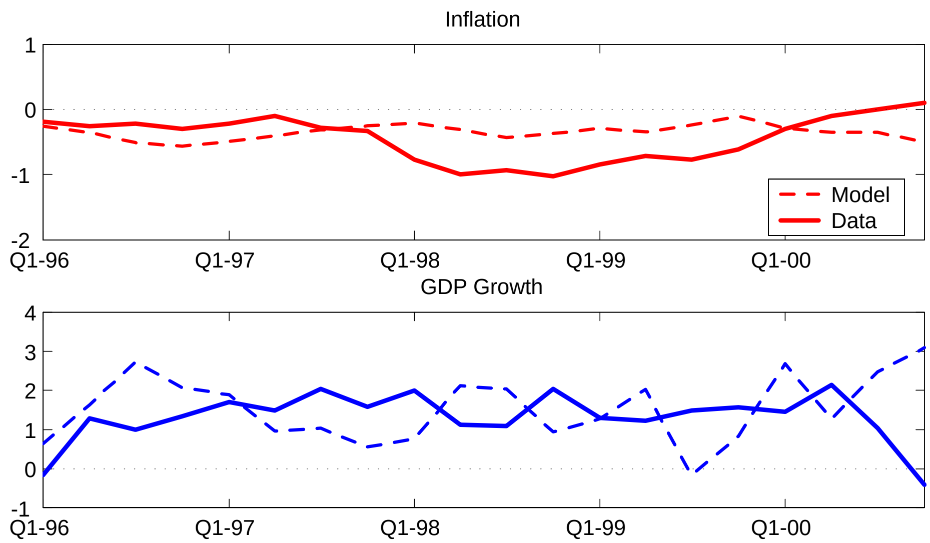

Figure 20: Simulations of inflation and GDP growth. Monetary policy follows Orphanides' Volcker-Greenspan rule with $x_t^{CB} = \overline{\Delta x}_t$ and $\phi_n = 1$.

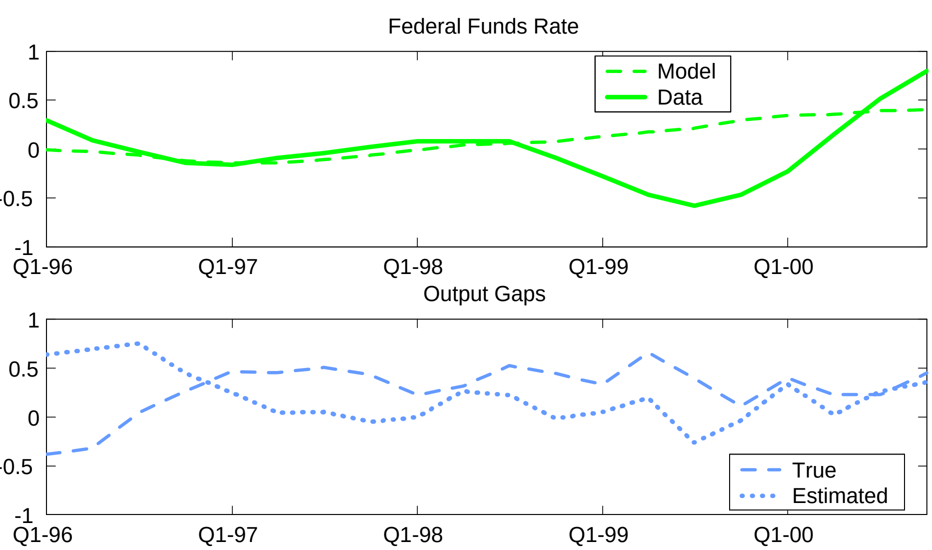

Figure 21: Simulations of federal funds rate and output gaps. Monetary policy follows Orphanides' Volcker-Greenspan rule with $x_t^{CB} = \overline{\Delta x}_t$ and $\phi_n = 1$.

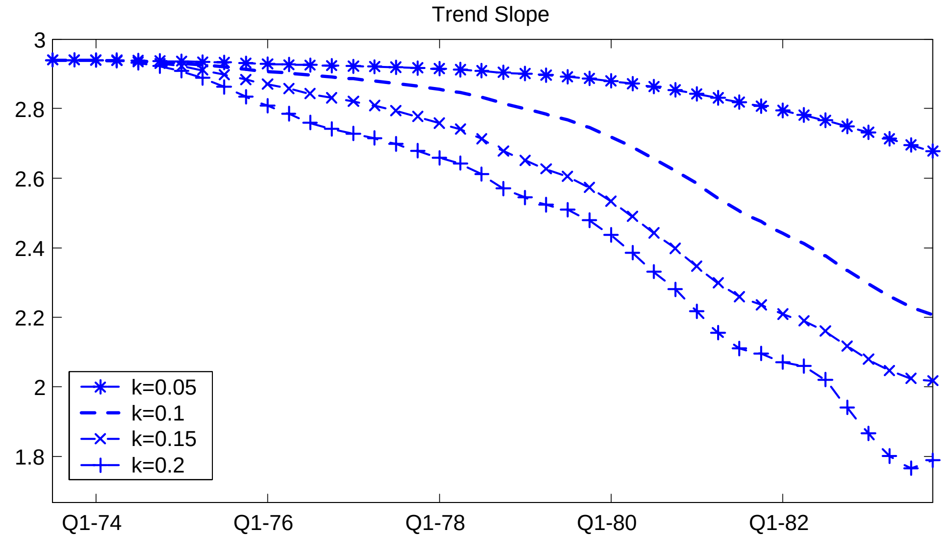

Figure 22: Evolution of "real-time" estimates of the slope of the output trend for different values of the constant gain parameter k. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.

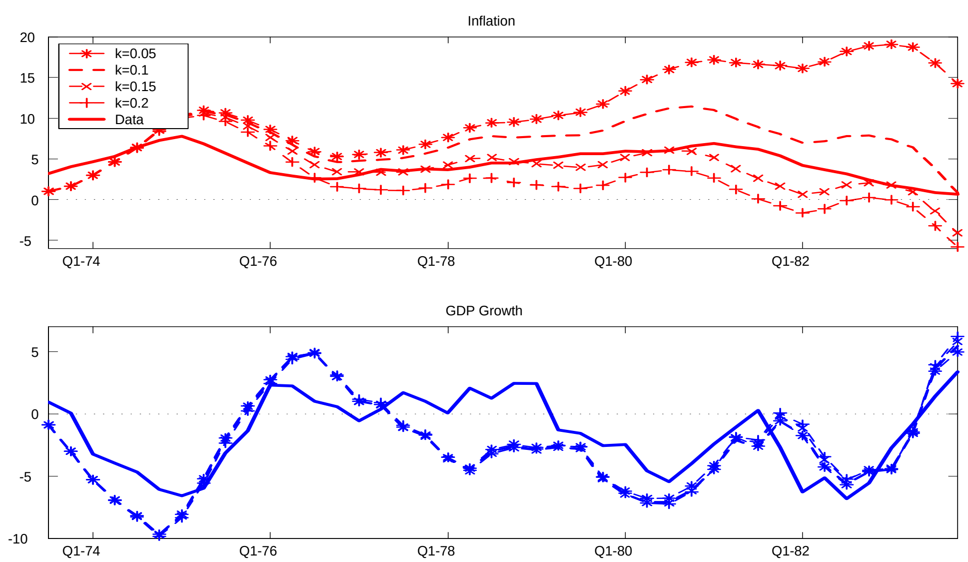

Figure 23: Simulations of inflation and GDP growth for different values of the constant gain paramater k. Monetary policy follows Orphanides' Pre-Volcker rule with $x_t^{CB} = \bar{x}_t$.