

Optimal Monetary Policy and Productivity Growth*

Andrea Tambalotti[†]
Princeton University

February 11, 2003

Abstract

This paper studies the optimal monetary policy response to persistent changes in the growth rate of productivity. To this end, we formulate a sticky price model in which agents are uncertain about the persistence of observed productivity shocks. In this environment, trend shocks are initially perceived as transitory, causing wage demands to lag behind realized productivity. As a result, marginal costs and inflation fall after a positive shock, consistent with the data—and with conventional wisdom. From a positive standpoint, the model's dynamics in response to a productivity shock are in line with the main features of available VAR estimates. They also provide a remarkably accurate account of the observed evolution of output and inflation in the wake of the recent productivity pickup. From a normative perspective, we find that the adjustment in the real wage required by a positive productivity shock should be shared between an increase in the nominal wage and a reduction in the price level, in proportions that depend on the relative stickiness of the two prices. Finally, the model lends formal support to the popular idea that the Federal Reserve reacted appropriately to the productivity resurgence of the late 1990s.

JEL Classification: E31, E52, E58

Keywords: Growth rate shocks, sticky prices, partial information, stagflation, New Economy.

*I wish to thank the Board of Governors of the Federal Reserve System for its hospitality, Ben Bernanke, Alan Blinder, Gauti Eggertsson, Giovanni Favara, Mark French, Dale Henderson, Thomas Laubach, Guido Lorenzoni, Ann Morning, Fabrizio Perri, Bruce Preston, Giorgio Primiceri, John Roberts, Ernst Schaumburg, Chris Sims, Lars Svensson, Lawrence Uren, Mark Watson, seminar participants at Princeton University, the Federal Reserve Board, the Inter-University Student Conference at the University of Pennsylvania and the Society for Economic Dynamics 2002 Meetings for comments and suggestions, Marty Eichenbaum for providing the figure of the estimated impulse responses and especially Pierre-Olivier Gourinchas and Mike Woodford for numerous discussions and precious advice. This paper represents the views of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff.

[†]Department of Economics, Princeton University, Princeton, NJ 08544. Email: tamba@princeton.edu. Revisions available at: <http://www.princeton.edu/~tamba>.

“By now, the story of the boom in information technology is well known, and nearly everyone perceives that the resulting more rapid growth of labor productivity is at least partly enduring. (...) With output per hour having accelerated, cost pressures have been patently contained. For the most part, the Federal Reserve generally recognized these changing fundamentals and calibrated American monetary policy accordingly. (Remarks by Chairman Alan Greenspan, *Challenges for Monetary Policymakers*, at the 18th Annual Monetary Conference: *Monetary Policy in the New Economy*, Cato Institute, Washington, D.C. October 19, 2000).

1 Introduction

The investigation of the role of technology and monetary policy shocks in economic fluctuations has been at the center of macroeconomic research for the last two decades, starting from the prototypical real business cycle models of Kydland and Prescott (1982) and Long and Plosser (1983) and arriving at the latest estimated structural models with nominal rigidities of Altig, Christiano, Eichenbaum and Linde (2002) and Smets and Wouters (2002), with the monetary vector autoregression literature and the New Neoclassical/Keynesian Synthesis to represent the key empirical and theoretical stepping-stones along the way.¹ Remarkably, given the historical evolution of this broad research agenda, the question of the interaction between monetary policy and technological progress, and in particular of the role that monetary policy might play in dampening economic fluctuations originating from productivity shocks, has received far less attention.² This is all the more surprising if we recall that the first attempts at introducing money and prices in a dynamic stochastic general equilibrium model were conducted within an otherwise standard RBC model, with fluctuations driven exclusively by technology shocks.³

This paper attempts to fill this gap on three complementary levels. First, at the positive level, it proposes a sticky price model of the transmission of productivity shocks, whose specification is guided by the observed dynamic behavior of real and nominal variables in response to an identified technology shock. Besides reproducing the main features of the response of the U.S. economy to an “average” shock to productivity growth, the model provides a surprisingly accurate account of the puzzling evolution of inflation and real activity that accompanied the productivity revival of the second half of the 1990s (Ball and Moffitt, 2001; Staiger, Stock and Watson, 2001). It also provides some clues on the possible causes of the “great stagflation” of the 1970s.⁴ Second, shifting

¹ See also Prescott (1986), Plosser (1989) and King and Rebelo (1999); Rotemberg and Woodford (1997, 1999), Christiano, Eichenbaum and Evans (2001), Boivin and Giannoni (2001) and Amato and Laubach (2002); Bernanke and Blinder (1992), Leeper, Sims and Zha (1997), Bernanke and Mihov (1998) and Christiano, Eichenbaum and Evans (1999); King and Wolman (1996), Woodford (1996), Yun (1996), Goodfriend and King (1997), Clarida, Galí and Gertler (1999) and Woodford (2002).

² Ireland (1996), Galí (2001) and Galí, López-Salido and Vallés (2002) are among the most significant exceptions.

³ Among the early contributions were King and Plosser (1984) and Cooley and Hansen (1989). See also Cooley and Hansen (1995).

⁴ Blinder (1979) is a classic reference on the subject. De Long (1997) and Sargent (1999) are two modern and

onto normative grounds, we formulate a policy problem that is approximately equivalent to that of maximizing the utility of the representative agent in the model. By solving this problem, we are able to characterize the optimal path of the policy instrument following a growth rate shock, together with the associated equilibrium fluctuations of the endogenous variables. These fluctuations are then compared to those that would emerge from some simple policy prescriptions, like strict inflation targeting, as well as to those that were observed after the major productivity shocks of the early seventies and mid-nineties. Finally, from a methodological perspective, the paper introduces a variant of the imperfect information story of Lucas (1972, 1975) into a New Keynesian model with nominal rigidities, thus mixing important aspects of the two predominant modern theories about the importance of money for business cycle fluctuations (Cooley and Hansen, 1995).

The first step in our analysis is to ask whether the dynamics triggered by a shock to the growth rate of productivity in a “standard” sticky price model match, at least qualitatively, what is observed in the data. In conducting this comparison, we mainly rely on the VAR evidence presented by Altig *et al.* (2002) (henceforth ACEL), whose estimated impulse responses to a permanent technology shock are reported in figure 1.⁵ One of the most significant features of these impulse responses is the negative conditional correlation between output and inflation, which is in contrast with the more familiar positive relation associated with movements along the Phillips curve. In particular, after a positive shock to the growth rate of productivity, inflation declines significantly for at least two quarters, and possibly for as much as three years. In a standard sticky price model on the contrary, a surge in productivity reduces unit labor costs, but with perfectly competitive labor markets, the ensuing increase in real wages exactly offsets the initial impact of the shock. This general equilibrium effect then, by severing the link between real marginal costs and productivity, insulates prices and inflation from technology shocks.⁶ The counterfactual behavior of the model along this particular dimension is especially troublesome for our analysis, since one of its main objectives is to derive a set of constraints for the optimal policy problem that reflect the nature of the trade-offs actually faced by central banks as they strive to insulate the economy from the effects of productivity shocks.

Among the maintained assumptions of the standard model, one in particular results both instrumental to the counterfactual behavior just described, and especially unrealistic in the context of our study. This is the assumption of perfect information on the distribution of shocks, which implies that all agents in the economy can precisely forecast the entire future evolution of observed shocks. In reality, one of the consistent historical features of the public’s and policy makers’ response to persistent growth rate shocks has been an extremely gradual updating of their expected long run effect.⁷ Following this lead, we then formulate a model in which the response of wages

influential reconstructions of those same events.

⁵ ACEL identify this shock as the only source of the unit root in average labor productivity, as in Galí (1999). According to their estimates, the *growth rate* of the shock is distributed as an AR(1) process with an autoregressive coefficient of 0.8.

⁶ Section 1.1 illustrates this point more formally.

⁷ This is witnessed for example by the extreme caution with which the media routinely receive new evidence on

to a productivity disturbance is delayed by a particular form of imperfect information, whereby agents, upon observation of a shock, cannot conclusively infer its degree of persistence. As a result, whenever the shock is in fact persistent, workers come to this realization only gradually, causing wage demands to lag behind realized productivity. Decoupling the movements of real wages and productivity in the short run, this mechanism contributes to a temporary fall in marginal costs and inflation following a positive productivity shock. The extent to which this transmission channel can quantitatively reconcile the model's predictions with the data is at the center of our positive investigation.

The idea that wage contracts might incorporate adjustments to slowly evolving wage aspirations, or to filtered estimates of the growth rate of productivity, and that the resulting inertia might drive a wedge between wages and productivity following significant changes in its rate of growth dates back at least to Jackman *et al.* (1982) and Braun (1984), and has witnessed a recent revival as a possible explanation of the unusual joint behavior of inflation and real activity in the second half of the nineties (Blinder, 2000; De Long, 2000; Ball and Moffitt, 2001). Our contribution in this respect is to take this idea from the realm of simple hypothesis, or of loosely specified regression models, and to incorporate it into a fully specified model of price and wage setting. We do this by relying on the version of the story based on the slow adjustment of growth estimates (Blinder, 2000), first, because it does not require any departure from standard rational accounts of agents' behavior and, second, because uncertainty on the persistence of the recent productivity pickup has been one of the dominant notes in the debate spurred by this important event.⁸

As a possible alternative to the research strategy pursued in this paper, we could have indeed directly adopted the theoretical model proposed in ACEL, and shown there to provide a good match to the empirical impulse responses that we are also trying to replicate. Two main considerations kept us from following this approach. First, we wanted to limit the complexity of our model to concentrate on the transmission channels of productivity shocks that are likely to be more relevant for short term policy considerations. This meant in particular abstracting from capital accumulation and all the related real frictions. As an important by-product of this simplification, it is possible to derive a second order approximation to the utility function of the representative agent as a model-consistent welfare criterion. Second, even with its wide array of real frictions, ACEL's model is not particularly successful at replicating the negative conditional correlation between inflation and real activity, which we consider a fundamental stylized fact of the transmission of productivity shocks. In this respect, imperfect information can be thought of as a complementary mechanism

the productivity acceleration that has affected the U.S. economy since the mid-nineties. On November 8th 2002 for example, after the release of strong third quarter productivity data, the New York Times reported that "a significant share of the step-up in economic efficiency achieved during the boom of the late 1990's *could continue* over the long run" (emphasis added). See also the Economics focus in The Economist of November 2nd 2002 ("Productivity promises: How much of America's surge in productivity growth can be sustained?") and the evolution of the Federal Reserve's views on productivity in the period 1997-1999 as reconstructed in section 4.

⁸ Jorgenson (2001) and Gordon (2002) are two prominent academic contributions to this debate; Greenspan (2000) provides the perspective of an important policymaker, while Woodward (2000) contains a journalistic account and a popular interpretation of the events.

to the ones identified in ACEL to explain the observed behavior of aggregate variables in response to technology shocks.

The paper’s main results can be summarized as follows. First, following a persistent shock to the growth rate of productivity, our model can generate reductions in costs and inflation that are close to those observed in the data. However, other informational frictions besides partial information are quantitatively at least as important for this result as the latter. Second, monetary policy should maintain a procyclical stance in reaction to technology disturbances, decreasing the nominal interest rate below its flexible price counterpart following a positive growth rate shock. Such policy has the effect of inducing the desired long run adjustment of the real wage through increases in the nominal wage rate and reductions in the price level, in proportions that depend on the relative flexibility of the two prices. As a consequence, policies that strictly target the price or wage inflation rate are clearly undesirable. On the other hand, a policy that tries to stabilize the central bank’s forecast of the model consistent output gap is close to optimal. Finally, when perturbed by an appropriately scaled shock, the model’s optimal equilibrium displays a path of average productivity, interest rates, inflation and output growth, which is remarkably close to that observed in the second half of the 1990s. Hence the conclusion that in that period the Federal Reserve calibrated American monetary policy according to the model...

The rest of the paper is organized as follows. Section 1.1 offers some more detail on the “standard” sticky price model and on its behavior following a productivity shock. Section 2 presents the model economy, starting with assumptions on tastes, technology and information. In particular, section 2.6 reviews the main steps in the derivation of the welfare criterion, whose details are relegated to appendix A. Section 3 describes the dynamics of the model following a productivity shock for a given policy rule, while section 5 discusses the properties of the optimal equilibrium. Sections 4 and 6 compare the macroeconomic developments that accompanied the two major post-war productivity shocks with simulations from the model, from both a positive and a normative perspective. Section 7 concludes. Finally, appendices B and C contain some details on the solution and filtering of the system of linearized first order conditions and on the computation of the optimal equilibrium.

1.1 The Neutrality of Productivity Shocks in a Standard Model

This section provides a complete description of what we referred to above as the “standard” sticky price model. There is certainly a significant degree of arbitrariness in defining a “standard” model in the context of any field. Nevertheless, we found it useful to crystallize the existing literature in a paradigmatic model, not with the intent of denouncing its assumptions as flawed, but rather as a way of organizing our thoughts around a familiar and largely diffused benchmark. Three fundamental assumptions characterize in our judgment this prototypical model. First, output is produced with labor as the only input. Second, goods’ prices are sticky, and set in staggered fashion. Third, labor markets are perfectly competitive. For instance, these are the assumptions

that underlie several contributions in the influential Taylor (1999) volume, as well as the models in Woodford (1996), Rotemberg and Woodford (1997) and Clarida, Galí and Gertler (1999), just to name a few prominent examples.

A simple but fairly representative rendition of this model is that in Galí (2001), which we summarize here for expositional convenience.⁹ The representative household's utility is separable between consumption and hours

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{Y_t^{1-\sigma}}{1-\sigma} - \frac{H_t^{1+\varphi}}{1+\varphi} \right)$$

and a continuum of monopolistically competitive firms indexed by z produces differentiated goods according to the log-linear production function

$$y_t(z) = a_t + h_t(z)$$

where a_t denotes the logarithm of productivity and $H_t \equiv \int_0^1 \exp[h_t(z)] dz$. Given an economy-wide competitive labor market, firms face a common real marginal cost

$$s_t = w_t - p_t - a_t$$

while workers' optimization on the consumption-leisure margin implies

$$w_t - p_t = \sigma y_t + \varphi h_t$$

where $w_t - p_t$ is the logarithm of the real wage and lower case letters denote the logarithms of their upper case counterparts. Making use of the approximation $y_t = a_t + h_t$, we find that in equilibrium

$$s_t = (\sigma - 1) a_t + (\sigma + \varphi) h_t$$

The reason for focusing on this cost measure is that in this class of models real marginal costs are the only determinants of firms' optimal pricing behavior. In the presence of price setting frictions *à la* Calvo (1983) for example, this results in a Phillips curve of the form

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \hat{s}_t$$

where π_t is inflation, \hat{s}_t denotes the deviation of s_t from its steady state level and κ is a positive coefficient that depends on the frequency of price adjustment.

Two important considerations follow from this derivation. First, with the restriction $\sigma = 1$ required for the existence of a balanced growth path, positive productivity shocks are reflected one to one in the equilibrium real wage, offsetting entirely the partial equilibrium cost reduction normally associated with productivity improvements. Evidently, under these conditions technology shocks cannot have any direct impact on prices and inflation. Intuitively, the assumption of perfectly

⁹ We refer the reader to Galí (2001) for further details on the model and the derivation of its equilibrium.

competitive labor markets imposes the desired cointegration between income and real wages by locking them together in the short run. Introducing a temporary wedge between the comovements of these two variables, imperfect information contributes to generate an environment in which increases in output per hour help to contain cost pressures, as observed by Chairman Greenspan in our opening quote. Note also that wage stickiness alone would not necessarily produce the desired result. This is because in a forward looking model, wage inflation would in fact overshoot the currently observed surge in productivity if agents anticipated further productivity increases in the future, as would be the case for the kind of persistent growth rate shocks considered here.¹⁰

A second important point is that, even in the case $\sigma \neq 1$, a monetary policy that stabilized the real marginal cost around its steady state level, or equivalently the level of output around its flexible price equilibrium, would also result in zero inflation, again severing the link between productivity and prices. This is precisely the case emphasized by Galí (2001), one in which productivity shocks do not pose any stabilization trade-off to the monetary authority. Interestingly, Galí *et al.* (2002) present some VAR evidence that, during the Volcker-Greenspan era, inflation did not significantly respond to identified productivity shocks, while this response was negative and significant in the pre-Volker period. This evidence would then seem to suggest that the normative question posed in this study has a very simple answer: stabilize the output gap, and that this is exactly what the Federal Reserve has been doing for at least twenty years. Several considerations make us doubt this conclusion. First, according to most accounts, the Federal Reserve spent a considerable amount of effort to “calibrate” its monetary policy in response to the productivity acceleration of the mid-nineties.¹¹ Second, Galí *et al.* (2002) do not present any explicit empirical evidence on the statistical significance of the difference in the behavior of inflation across the two sub-periods in their analysis.¹² Third, as we will see in section 5, although in our model output gap stabilization closely approximates the optimal policy, it does not imply a stable inflation rate. Quite to the contrary, a policy that did stabilize inflation would result in very undesirable welfare consequences. We find that these exceptions provide enough ground to proceed with our analysis, which begins in the next section with the description of our structural model.

2 The Model

This section presents a dynamic stochastic general equilibrium model of the monetary transmission mechanism, whose microeconomic foundations derive from the work of Rotemberg and Woodford (1997) and Amato and Laubach (2002). As in Amato and Laubach (2002) and Erceg, Henderson and Levin (2000), we introduce wage setting frictions in the labor market, along with price stickiness in the market for consumption goods. As argued in the previous section, abandoning the assumption

¹⁰ This is a consequence of the well-known front loading effect described by Christiano *et al.* (2001). Mankiw and Reis (2002) present a “sticky information” model in which this effect is absent and that is very similar in spirit to the model developed here.

¹¹ See for example Woodward (2000).

¹² In fact, according to preliminary results in ACEL, the evidence of subsample instability is at best marginal.

of perfect competition in the labor market is a necessary condition for productivity shocks to play any role in this class of models. Assuming that labor and goods markets share the same kind of price frictions is a particularly convenient way of fulfilling this condition.

The section’s main contribution is twofold. First, it considers explicitly the possibility of secular growth. This requires to impose some restrictions on utility and production functions, of the kind that King, Plosser and Rebelo (1988, 2002) showed to be necessary to produce long run implications of a model economy in accordance with the basic growth facts. It also illustrates how these restrictions, together with a fairly natural assumption on wage indexation, are sufficient to maintain almost intact the basic structure of the standard models of pricing based on Calvo (1983). Second, it presents an environment in which productivity is the sum of two stochastic components with differing degrees of persistence, but only its overall level is observable. As a consequence, agents forecast its future evolution through a Kalman filter, which imparts significant inertia to their reaction to shocks, contributing to generate the negative conditional correlation between inflation and real activity that is the distinctive feature of the model.

The structural representation of the economy is completed with the derivation of a welfare criterion for policy evaluation, which is based on a second order expansion to the utility of the representative agent. Given appropriate assumptions on the magnitude of the distortions in the economy, this expansion, evaluated through a linear approximation to the structural equilibrium relations, provides an accurate gauge of the effect of alternative policies on the welfare of the representative agent, up to a residual that is of third order in the amplitude of the shocks.

2.1 Tastes, Technology and Information

We consider an economy populated by three classes of agents: a government, a continuum of households indexed by $j \in [0, 1]$ and a continuum of firms indexed by $z \in [0, 1]$, whose ownership is equally shared among the households.

Government The government is composed of two branches, a fiscal and a monetary authority. The fiscal authority levies a proportional tax τ on the sales of consumption goods and rebates the proceeds to the households through a lump-sum transfer. The monetary authority sets the level of the short term nominal interest rate, as further detailed in section 2.7 below.

Firms Firm z is a monopolist on the market for its output, that is produced using the composite labor input $h_t(z)$. As in Erceg *et al.* (2000), this is a Dixit-Stiglitz (1977) aggregate of a continuum of specialized “skills” distributed among the households

$$h_t(z) \equiv \left[\int_0^1 h_t^z(j)^{\frac{\theta_w-1}{\theta_w}} dj \right]^{\frac{\theta_w}{\theta_w-1}}$$

Each of these specialized “skills” trades at a wage $w_t(j)$ on an economy-wide market, from which all firms draw their supply. They generate a total demand for skill j

$$h_t(j) = H_t \left(\frac{w_t(j)}{W_t} \right)^{-\theta_w}$$

where $H_t \equiv \int_0^1 h_t(z) dz$ denotes aggregate hours and

$$W_t \equiv \left[\int_0^1 w_t(j)^{1-\theta_w} dj \right]^{\frac{1}{1-\theta_w}} \quad (1)$$

is the wage index associated with the minimum expenditure purchase of one unit of H_t .

Output $y_t(z)$ is produced according to the production function

$$y_t(z) = A_t f(h_t(z))$$

where f is increasing and concave, and A_t represents a technology factor.¹³ Following a long-standing tradition in the New Keynesian analysis of the monetary transmission mechanism, we abstract from capital accumulation.

Our main point of departure from this literature is instead in the distributional assumptions on A_t . In particular, we assume that the growth rate of productivity, $\gamma_t^a \equiv \ln A_t - \ln A_{t-1}$, follows a stationary process, which is the sum of a persistent (“trend”) and an i.i.d. component, as in

$$\gamma_t^a = \gamma_t + \varepsilon_t^a \quad (2a)$$

$$(1 - \rho_L)(\gamma_t - \gamma) = \varepsilon_t^\gamma \quad (2b)$$

where ε_t^a and ε_t^γ are orthogonal i.i.d. shocks with mean zero and variances σ_a^2 and σ_γ^2 respectively, $0 < \rho < 1$ and γ is the average growth rate.¹⁴ We furthermore assume that all agents in the economy, including the monetary authority, can observe the level of productivity A_t , along with the vector of endogenous variables X_t , but not the shocks ε_t^γ and ε_t^a separately. More formally, we denote agents’ information set by $I_t \equiv \{A_\tau; X_\tau\}_{\tau \leq t} \subset I_t^f \equiv \{\varepsilon_\tau^\gamma, \varepsilon_\tau^a; X_\tau\}_{\tau \leq t}$, where I_t^f represents full information. Note that this form of limited information is not necessarily a “binding” constraint on agents’ decision making, at least as long as their actions do not require them to forecast future economic conditions. This would for instance be the case for a firm that can reoptimize its price every period. It is only when decision making involves expectations about the future evolution of productivity, as for a firm that is allowed to reset its price only at certain intervals, that partial information becomes relevant. In this case, we assume that agents would make the best possible use of all the available information and use a Kalman filter to forecast the future evolution of productivity, as detailed in appendix B.

¹³ In what follows we refer to A_t indifferently as technology or (total factor) productivity. When the distinction is relevant, we will refer explicitly to (average) labor productivity.

¹⁴ The model can easily accommodate a much richer specification of the productivity process. See appendix B for more details.

Households Household j is endowed with one unit per period of the specialized labor $h_t(j)$, on whose market it is a monopolist. The household sets the wage rate $w_t(j)$ for period t and supplies the amount of hours demanded by firms at the posted price, as further discussed in section 2.4. It also chooses a sequence of consumption bundles $\left\{c_s^j(z), z \in [0, 1]\right\}_{s=t}^{\infty}$ to maximize a time separable expected utility function of the form

$$E_{t-1} \sum_{s=t}^{\infty} \beta^{s-t} [u(C_s(j)) - v(h_s(j))] \quad (3)$$

where $\beta \in (0, 1)$ is the subjective discount factor and $C_t(j)$ is an aggregator of the levels of consumption of different goods, with a constant elasticity of substitution $\theta_p > 1$

$$C_t(j) \equiv \left[\int_0^1 c_t^j(z)^{\frac{\theta_p-1}{\theta_p}} dz \right]^{\frac{\theta_p}{\theta_p-1}}$$

The expectation in equation (3) is based on the information publicly available as of time $t-1$, I_{t-1} . This captures an information processing (or implementation) delay, as for example in Christiano *et al.* (2001). In this context, besides preventing demand from jumping immediately in response to shocks, this assumption has the realistic feature of making the output gap unobservable, in a sense that will be made more precise in section 2.3. Note also that money does not enter the utility function, and in fact does not play any explicit role in our model. We are in other words assuming to be in a “cashless limiting economy”, in which the role of money balances in facilitating transactions is negligible (Woodford, 2002).

If we assume the existence of complete financial markets, and therefore of a unique kernel $Q_{t,s}$ for the pricing of stochastic flows of nominal income, given a sequence of goods’ prices $\{p_s(z), z \in [0, 1]\}_{s=t}^{\infty}$, we can write the household’s intertemporal budget constraint looking forward from time t as

$$\sum_{s=t}^{\infty} E_{t-1} \left[Q_{t-1,s} \left[\int_0^1 p_s(z) c_s^j(z) dz \right] \right] \leq E_{t-1} Q_{t-1,t} B_t(j) + \sum_{s=t}^{\infty} E_{t-1} Q_{t-1,s} [w_s(j) h_s(j) + \Pi_s(j)] \quad (4)$$

where $B_t(j)$ is initial wealth and $\Pi_t(j)$ are the profits accruing to the household, net of any lump sum taxes.¹⁵ In general, depending on the realized frequency with which they are offered the chance to reset their wage, different agents will experience very different employment histories, and thus very different levels of human wealth. Nevertheless, if at any point in time $\tau < t$ the distribution of wealth $\{B_\tau(j)\}_j$ is such that the right hand side of (4) is the same for all j , then, from then on, the optimal consumption profile will be identical for all households. In what follows we will assume that this is indeed the case, and consequently drop the j index that distinguishes individual consumption choices.

¹⁵ We refer the reader to Woodford (2002) for a detailed derivation of the intertemporal budget constraint in a closely related model and for a general discussion of models in which current consumption decisions are based on lagged information.

The solution to the household's consumption problem yields the following optimality conditions. First, optimal intratemporal allocation of a given amount of expenditure across the differentiated goods implies

$$c_t(z) = C_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta_p} \quad (5)$$

where $c_t(z)$ is total demand for good z , C_t is the index of aggregate consumption and P_t is defined analogously to (1). Second, the optimal choice of aggregate consumption satisfies

$$E_{t-1} [u_C(C_t)] = \Lambda_{t-1} E_{t-1} [Q_{t-1,t} P_t] \quad (6)$$

where Λ_{t-1} , the multiplier on the intertemporal budget constraint (4), measures the marginal utility of *nominal* income. Finally, the optimal allocation of resources across time and states achieved through complete financial markets implies

$$\Lambda_t Q_{t,s} = \beta^{s-t} \Lambda_s \quad (7)$$

and therefore

$$\Lambda_t = \beta E_t [R_t \Lambda_{t+1}] \quad (8)$$

where $R_t^{-1} \equiv E_t [Q_{t,t+1}]$ is the price of a riskless bond held between t and $t+1$, so that R_t is the gross short term nominal interest rate.

2.2 Flexible Price Equilibrium and the Balanced Growth Path

Having introduced the basic building blocks of the model, and before proceeding to the description of the price frictions that constitute the necessary premise to a well defined policy problem, it is useful to consider the model's behavior with flexible prices. This exercise becomes even more crucial in the context of our non-stationary economy, since it will help us to characterize the balanced growth path that the economy would follow in the absence of shocks. It will then be around this path that we approximate the model's first order conditions to provide a dynamic characterization of the equilibrium responses to shocks.

When prices are flexible, firms reset their price every period to maximize their instantaneous profits, taking as given the demand function for their output. From (5) and the fact that private consumption is the only source of demand, this function is simply

$$y_t(z) = Y_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta_p} \quad (9)$$

The result of profit maximization is an optimal price (net of the sales tax) that is set as a fixed markup over marginal cost. In real terms

$$\frac{p_t(z)}{P_t} (1 - \tau) = \mu_p s_t(z)$$

where $\mu_p \equiv \frac{\theta_p}{\theta_p - 1} > 1$ is the optimal (gross) markup for a monopolist facing a demand curve with constant elasticity $-\theta_p$, and $s_t(z)$ is the real cost function for firm z

$$s_t(z) = \frac{W_t}{P_t} \left[A_t f' \left(f^{-1} \left(\frac{y_t(z)}{A_t} \right) \right) \right]^{-1}$$

As for wages, they will similarly be set as a markup over the marginal rate of substitution between consumption and leisure

$$\frac{w_t(j)}{P_t} = \mu_w \frac{v_h(h_t(j))}{u_C(Y_t)}$$

where $\mu_w \equiv \frac{\theta_w}{\theta_w - 1} > 1$.

As a result, in a symmetric equilibrium with identical prices and wages across firms and households, the level of output is implicitly defined by the solution to

$$\frac{u_C(Y_t)}{v_h(f^{-1}(Y_{At}))} A_t f' (f^{-1}(Y_{At})) = \frac{\mu_p \mu_w}{1 - \tau}$$

where $Y_{At} \equiv Y_t/A_t = f(H_t)$ is the effective level of output. Note that, without some restrictions on the utility function u , we cannot guarantee the existence of a growth path with constant hours (and therefore a constant value of effective output), a very natural requirement to impose on any model economy. As shown by King *at al.* (1988), given separability between consumption and leisure in the period utility function, the restriction $u(C) = \ln C$ is a necessary and sufficient condition for the existence of a balanced growth path, which we define here as a stationary equilibrium in which hours worked and the shares of income accruing to wages and profits are bounded away from zero and one. With this restriction, and in the absence of shocks,

$$\frac{f'(f^{-1}(\bar{Y}_A))}{\bar{Y}_A v_h(f^{-1}(\bar{Y}_A))} = \frac{\mu_p \mu_w}{1 - \tau}$$

implicitly defines the constant level of effective output along the balanced growth path, \bar{Y}_A .

This expression also highlights how the presence of market power in the goods and labor markets, together with distortionary taxation, drives a wedge between the marginal rates of substitution and transformation of labor and consumption, whose equality characterizes the efficient equilibrium. To abstract from the role that monetary policy could (and optimally would) play in ameliorating this inefficiency, we then assume that the fiscal authority subsidizes the sales of consumption goods exactly enough to offset the monopoly distortions, setting $\tau = 1 - \mu_p \mu_w$ (Rotemberg and Woodford, 1999), from which we get

$$\bar{Y}_A^{-1} = \frac{v_h(f^{-1}(\bar{Y}_A))}{f'(f^{-1}(\bar{Y}_A))} \tag{10}$$

By the second welfare theorem, as well as by simple inspection of (10), $\bar{Y}_t \equiv A_t \bar{Y}_A$, the efficient level of output, is seen to maximize the utility of the representative agent, given the resource constraint. As a result, a linear approximation to the structural equilibrium relations will be all what is needed to compute a second order accurate approximation to the utility of the representative agent, as further illustrated in section 2.6.

To complete the characterization of the balanced growth path we need only to remark that, when output is equal to its efficient level and A_t grows at the deterministic rate γ , equation (6) implies that $\lambda_{At} \equiv \beta \Lambda_t P_t A_t = \bar{Y}_A^{-1}$ so that from (8) we obtain

$$\bar{i} \equiv \ln \bar{R} = \gamma - \ln \beta + \pi$$

the value of the (continuously compounded) nominal interest rate along a growth path with steady inflation at rate π .¹⁶

We now turn to the description of the equilibrium fluctuations induced by the productivity shocks, and of their interaction with the price setting frictions in the goods and labor markets. We begin with the analysis of the demand side of the economy.

2.3 Demand

The demand side of the economy is completely characterized by a dynamic IS equation, which links current expenditures to expectations of future real interest rates. This relation is obtained starting from a log-linear approximation to the Euler equation (8) around the balanced growth path with no inflation

$$\hat{\lambda}_{At} = E_t \left[\hat{i}_t - \pi_{t+1} + \hat{\lambda}_{At+1} - \hat{\gamma}_{t+1}^a \right] \quad (11)$$

where $\hat{\lambda}_{At} \equiv \ln (\Lambda_t P_t A_t / \bar{Y}_A^{-1})$, $\hat{i}_t \equiv \ln (R_t / \bar{R})$ is the deviation of the nominal interest rate from its steady state value, $\pi_t \equiv \ln (P_t / P_{t-1})$ is the inflation rate and $\hat{\gamma}_t^a \equiv \gamma_t^a - \gamma$. We can then translate this expression in terms of expenditures if we note that, with logarithmic utility, and defining the output gap variable $x_t \equiv \ln (Y_t / A_t \bar{Y}_A)$, (6) and (7) imply the approximate relation

$$x_t = -E_{t-1} \hat{\lambda}_{At} - (\ln A_t - E_{t-1} \ln A_t)$$

which together with (11) yields

$$x_t = E_{t-1} x_{t+1} - E_{t-1} [\hat{i}_t - \pi_{t+1} - \hat{r}_t^e] - (\hat{\gamma}_t^a - E_{t-1} \hat{\gamma}_t^a) \quad (12)$$

Here x_t denotes the deviation of output from its *efficient* level, our preferred notion of output gap, while $\hat{r}_t^e \equiv E_t \hat{\gamma}_{t+1}^a$ is the real interest rate that would prevail under flexible prices and *without* informational delays, expressed in deviation from its value along the balanced growth path, \bar{i} .

Note that if we employed an alternative notion of the output gap, as the deviation of output from its equilibrium level with flexible prices, but *with* informational delays, $x_t^n \equiv \ln (Y_t / E_{t-1} A_t \bar{Y}_A)$, the IS equation could be written in the simpler form

$$x_t^n = E_{t-1} x_{t+1}^n - (E_{t-1} \hat{i}_t - E_{t-1} \pi_{t+1} - \hat{r}_t^n)$$

where now $\hat{r}_t^n \equiv E_{t-1} \hat{\gamma}_{t+1}^a$ is the familiar natural rate of interest. The reason for choosing the more cumbersome formulation in (12) is that x_t is the notion of output gap that is relevant for welfare comparisons, as shown in section 2.6.

¹⁶ In what follows we will restrict our attention to the case $\pi = 0$.

Another important characteristic of this demand formulation is that the output gap is not part of the information set I_{t-1} on which consumers base their decisions at time t . In fact, even if expenditure demand is predetermined one period in advance, it is only with the realization of the productivity shock at time t that the amount of hours needed to satisfy that level of demand, and therefore the efficient level of output, become known. As a result, the model exhibits the realistic feature of an output gap whose exact realization is uncertain at the time at which agents are required to take their decisions. Moreover, since only forecastable movements in the policy instrument \hat{i}_t have an effect on the output gap at time t , the goal of output stabilization, although desirable, cannot be perfectly achieved. The extent to which this influences the optimal policy is one of the objects of our analysis in section 5.

2.4 Wage Setting

The wage setting mechanism is based on the random staggering device of Calvo (1983), modified to accommodate indexation to productivity growth and decision lags. In particular, we assume that every period a fraction $1 - \alpha_w$ of outstanding labor contracts is randomly selected for renegotiation. A fraction $\tilde{\psi}_w$ of these new contracts is set on the basis of I_{t-1} information and carries an initial wage w_t^1 . The remaining contracts are set on the basis of I_{t-2} information, with wage w_t^2 . In general, this structure of delays has the objective of moderating the response of wages to shocks in the short run, to bring it more in line with the available empirical evidence.¹⁷ Here, it has the added benefit of attributing a consistent information set to members of the same household operating in the goods and labor markets.

We complete the characterization of the wage setting process describing the behavior of the fraction α_w of wages excluded from renegotiation. Similarly to Altig *et al.* (2002), we assume that these wages are indexed to past price inflation and to productivity growth, so that a wage contract j that is not renegotiated between $t - 1$ and t carries a wage $w_t(j)$ such that

$$\ln w_t(j) = \ln w_{t-1}(j) + \Omega_t$$

where Ω_t is defined as

$$\Omega_t \equiv \lambda_w \pi_{t-1} + \lambda_\gamma E_{t-1} \gamma_t^a + (1 - \lambda_\gamma) \gamma_{t-1}^a$$

and $\lambda_\gamma \in \{0, 1\}$. This means that, every quarter, besides being automatically increased by a fraction λ_w of past quarter's inflation, contractual wages are also allowed to incorporate an adjustment for either past quarter's observed productivity growth, or past quarter's forecast of its current value.

The reasons for assuming this particular form of indexation are twofold. First, in a stochastic equilibrium that exhibits stationary fluctuations around a balanced growth path, real wages and productivity need to be cointegrated. Given that under Calvo wage setting a positive fraction

¹⁷ The forward looking nature of standard optimizing models of pricing implies a front loaded response of prices to forecastable movements in the fundamentals, as observed for example by Christiano *et al.* (2001). Information delays help to moderate this effect.

of wages is not drawn for renegotiation for an arbitrarily large number of periods, contractual indexation guarantees that those wages do not fall excessively out of line with the general level of productivity. Second, consistency with the informational assumptions laid down so far requires that wages be indexed to a measure of productivity included in the information set I_{t-1} . Our indexation scheme considers two such measures, the expectation of the future value of productivity, $E_{t-1}\gamma_t^a$, and its latest available measurement, γ_{t-1}^a . As we will see, some of our results are sensitive to the prevalent form of indexation, which explains why we propose a specification that nests both measures.

The expression for the wage index

$$W_t = \left[(1 - \alpha_w) \left(\tilde{\psi}_w (w_t^1)^{1-\theta_w} + (1 - \tilde{\psi}_w) (w_t^2)^{1-\theta_w} \right) + \alpha_w (W_{t-1}\Omega_t)^{1-\theta_w} \right]^{\frac{1}{1-\theta_w}}$$

provides a useful summary of the assumptions introduced so far. To determine the equilibrium value of its components, we first note that wages reoptimized in period t on the basis of $t-1$ information, w_t^1 , are chosen to maximize the expected present discounted value of wage income, expressed in terms of utils and net of the disutility of labor

$$E_{t-1} \left\{ \sum_{s=t}^{\infty} (\alpha_w \beta)^{s-t} [\Lambda_s w_{t,s}^1 h_s^1 - v(h_s^1)] \right\}$$

where $w_{t,s}^1 \equiv w_t^1 \left(\prod_{k=t+1}^s \Omega_k \right)$ is the wage rate at time s , conditional on no renegotiation having occurred between t and s , and $h_s^1 \equiv H_s \left(\frac{w_{t,s}^1}{W_s} \right)^{-\theta_w}$ is the amount of labor demanded at that wage. The optimal choice of w_t^1 therefore satisfies

$$0 = E_{t-1} \left\{ \sum_{s=t}^{\infty} (\alpha_w \beta)^{s-t} h_s^1 [w_{t,s}^1 \Lambda_s - \mu_w v_h(h_s^1)] \right\}$$

which we log-linearize around the balanced growth path with no inflation to obtain

$$E_{t-1} \left\{ \sum_{s=t}^{\infty} (\alpha_w \beta)^{s-t} \left[(1 + \nu \theta_w) \hat{w}_t^1 - (1 + \nu \phi) x_s - (1 + \nu \theta_w) \sum_{k=t+1}^s (\hat{\pi}_k^w - \hat{\Omega}_k) + \hat{\omega}_{As} \right] \right\} = 0$$

where $\nu = \frac{v_{hh}\bar{h}}{v_h}$ measures the curvature of the disutility of labor, $\phi^{-1} \equiv \frac{f'_{\bar{h}}}{f}$ is the elasticity of the production function, $\hat{w}_t^1 \equiv \ln(w_t^1/W_t)$, $\hat{\pi}_t^w \equiv \ln(W_t/W_{t-1}) - \gamma$ is the deviation of nominal wage inflation from its steady state value, $\hat{\Omega}_t \equiv \ln \Omega_t - \gamma$ and $\hat{\omega}_{At} \equiv \ln(W_t/A_t P_t \bar{\omega}_A)$ is the percentage deviation of the effective real wage from its value on the balanced growth path

$$\bar{\omega}_A = \mu_w \bar{Y}_A v_h (f^{-1}(\bar{Y}_A))$$

The first order condition can then be solved for \hat{w}_t^1 and, together with the fact that $w_t^2 = E_{t-2} w_{t-1}^1$, and after some algebra, this is enough to characterize the evolution of the wage index. This is usefully summarized by a wage Phillips curve of the form

$$\tilde{\pi}_t^w = (1 - \psi_w) E_{t-2} \tilde{\pi}_t^w + \psi_w E_{t-1} [\kappa_w x_t - \xi_w \hat{\omega}_{At} + \beta \tilde{\pi}_{t+1}^w] + \zeta_w \Phi_t \quad (13)$$

where $\psi_w \equiv \frac{\alpha_w \tilde{\psi}_w}{1 - \psi_w (1 - \alpha_w)}$, κ_w , ξ_w and ζ_w are positive coefficients, $\tilde{\pi}_t^w \equiv \hat{\pi}_t^w - E_{t-2} \hat{\Omega}_t$ is the deviation of wage inflation from the forecastable component of its index and $\Phi_t \equiv \hat{\Omega}_t - E_{t-2} \hat{\Omega}_t$ is its forecast error.¹⁸ According to this equation, deviations of wage inflation from the value that would be predicted on the basis of indexation alone are a function of the expected future values of the effective wage level and of the output gap. The latter serves as a proxy for the amount of hours that workers expect to supply at the posted wage, and for the marginal disutility thereof.

The fact that in this class of models with indexation the resulting Phillips curve is conveniently expressed in terms of deviations of inflation from its index is familiar from Woodford (2002). In this particular example, we used instead the deviation $\tilde{\pi}_t^w$ from the forecast because this is the notion of wage inflation that is relevant for welfare comparisons, as shown in section 2.6. Note however that even if we adopted the more conventional deviation $\hat{\pi}_t^w - \hat{\Omega}_t$, we would still need to append a “shock” term similar to Φ_t to the resulting wage setting relation.

2.5 Price Setting

The price setting process in the goods’ market has the same basic structure described above, except that contractually fixed prices are now assumed to be partially indexed only to past observed inflation. Therefore, a price $p_t(z)$ that is not drawn for reoptimization at time t is updated as

$$\ln p_t(z) = \ln p_{t-1}(z) + \lambda_p \pi_{t-1}$$

Moreover, we maintain the assumption that a fraction $\tilde{\psi}_p$ of the prices to be reoptimized each period is chosen on the basis of time $t - 1$ information, while the remaining $1 - \tilde{\psi}_p$ prices are based on information dated $t - 2$. The price index is now

$$P_t = \left[(1 - \alpha_p) \left(\tilde{\psi}_p (p_t^1)^{1-\theta_p} + (1 - \tilde{\psi}_p) (p_t^2)^{1-\theta_p} \right) + \alpha_p \left(P_{t-1} \left(\frac{P_{t-1}}{P_{t-2}} \right)^{\lambda_p} \right)^{1-\theta_p} \right]^{\frac{1}{1-\theta_p}}$$

Prices reoptimized on the basis of I_{t-1} information, p_t^1 , are chosen to maximize the expected present discounted value of future profits

$$E_{t-1} \left\{ \sum_{s=t}^{\infty} \alpha_p^{s-t} Q_{t,s} \left[(1 - \tau) p_t^1 \left(\frac{P_{s-1}}{P_{t-1}} \right)^{\lambda_p} A_s y_{As}^1 - W_s f^{-1}(y_{As}^1) \right] \right\} \quad (14)$$

where

$$y_{As}^1 \equiv Y_{As} \left[\frac{p_t^1}{P_s} \left(\frac{P_{s-1}}{P_{t-1}} \right)^{\lambda_p} \right]^{-\theta_p}$$

is the level of demand for a good whose price has not been reoptimized between times t and s . The first order condition for the maximization of (14) is

$$0 = E_{t-1} \left\{ \sum_{s=t}^{\infty} (\alpha_p \beta)^{s-t} \frac{\Lambda_s A_s}{\Lambda_t} y_{As}^1 \left[\frac{p_t^1}{P_t} \left(\frac{P_{s-1}}{P_{t-1}} \right)^{\lambda_p} - \frac{\mu_p}{1-\tau} \frac{W_s}{P_s A_s} \frac{P_s}{P_t} f_1^{-1}(y_{As}^1) \right] \right\}$$

¹⁸ The details of the calculations, which follow very closely the steps illustrated in Rotemberg and Woodford (1998), are collected in an Appendix that is available from the author upon request. The values of the coefficients that appear in the log-linearized supply block, equations 13 and 15, are collected in table 4.

which we again log-linearize and solve to obtain an expression for $\hat{p}_t^1 \equiv \ln(p_t^1/P_t)$

$$\frac{1+\omega_p\theta_p}{1-\alpha_p\beta}\hat{p}_t^1 = E_{t-1} \left\{ \sum_{s=t}^{\infty} (\alpha_p\beta)^{s-t} \left[\omega_p x_s + \hat{\omega}_{As} + (1 + \omega_p\theta_p) \sum_{k=t+1}^s \tilde{\pi}_k \right] \right\}$$

where $\omega_p \equiv -\frac{f_p f''}{f'^2}$ is the elasticity of the real marginal cost function with respect to output, holding wages fixed, and $\tilde{\pi}_t \equiv \pi_t - \lambda_p \pi_{t-1}$ is the deviation of price inflation from its index.

As before, manipulation of the first order condition and of the expression for the price index results in a price Phillips curve of the form

$$\tilde{\pi}_t = (1 - \psi_p) E_{t-2} \tilde{\pi}_t + \psi_p E_{t-1} [\kappa_p x_t + \xi_p \hat{\omega}_{At} + \beta \tilde{\pi}_{t+1}] \quad (15)$$

where ψ_p is defined analogously to ψ_w and ξ_p and κ_p are positive coefficients. This equation relates the amount of inflation in excess of automatic indexation to the expected future values of the output gap and the effective real wage. These terms reflect in turn the two basic components of the marginal cost, one capturing the effect of higher output demand on the efficiency of the marginal input, the other measuring its unit cost. Differently from the wage equation 13, this relation is not perturbed by terms like Φ_t because the amount of the automatic indexation for prices is perfectly forecastable on the basis of information dated $t - 2$. This also accounts for the different definition of the inflation deviation $\tilde{\pi}_t$.

2.6 The Welfare Criterion

As already pointed out in section 2.1, the availability of a complete set of financial markets, providing households with perfect insulation from idiosyncratic income fluctuations, implies the existence of a representative consumer. Nevertheless, since the source of those fluctuations is the realized frequency with which wages can be reoptimized, different households still experience heterogenous employment histories, and therefore different levels of *ex-post* utility. It seems then natural to compare the performance of different monetary policies on the basis of their effect on the utility of the “average” household, defined as

$$E_{-1} \sum_{t=0}^{\infty} \beta^t U_t \equiv E_{-1} \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t - \int_0^1 v(h_t(j)) dj \right\}$$

Following Woodford (2002), this section sketches the main steps in the derivation of a second order expansion to this utility function. Given first order accuracy of the log-linearized constraints under which monetary policy is conducted, this approximation guarantees that, given a sequence of economies ordered by progressively tighter bounds $\|\varepsilon\|$ on the amplitude of the exogenous shocks, policies that appear to be better under the approximate criterion, indeed deliver higher levels of utility for the representative agent in economies far enough along the sequence.¹⁹

¹⁹ More precisely, $\|\varepsilon\|$ is a uniform bound on the elements of the vector of stationary shocks $\varepsilon_t \equiv (\varepsilon_t^c, \varepsilon_t^z)$.

Starting from the consumption term, we see that no approximation is actually needed, since

$$\ln Y_t = \ln (Y_t/A_t \bar{Y}_A) + \ln \bar{Y}_A A_t = \hat{Y}_{At} + t.i.p. \quad (16)$$

where we used the fact that consumption is the only source of demand in the model (i.e. $C_t = Y_t$) and that the exogenous term $\ln \bar{Y}_A A_t$ does not alter the preference ranking of alternative paths for the endogenous variables, hence the notation *t.i.p.* (term independent of policy). As for the disutility of labor, we have

$$\int_0^1 v(h_t(j)) dj = v_h \bar{h} \left\{ E_j \hat{h}_t(j) + \frac{1}{2} (1 + \nu) \left[\text{Var}_j \hat{h}_t(j) + \left(E_j \hat{h}_t(j) \right)^2 \right] \right\} + \mathcal{O}(\|\varepsilon\|^3)$$

where $\hat{h}_t(j) \equiv \ln(h_t(j)/\bar{h})$ is the percentage deviation of total demand for “skill” j from its steady state value. With a fair amount of algebra, this expression can be turned into

$$\int_0^1 v(h_t(j)) dj = v_h \frac{\bar{Y}_A}{f'} \left\{ \hat{Y}_{At} + \frac{1}{2} (1 + \omega) \hat{Y}_{At}^2 + \frac{1}{2} (1 + \omega_p \theta_p) \theta_p \text{Var}_z \hat{p}_t(z) + \frac{1}{2} (1 + \nu \theta_w) \theta_w \phi^{-1} \text{Var}_j \hat{w}_t(j) \right\}$$

where $\omega \equiv \omega_p + \nu \phi$ and the reminder is omitted for notational simplicity. Together with (16) this yields

$$U_t = \left(1 - v_h \frac{\bar{Y}_A}{f'} \right) \hat{Y}_{At} - v_h \frac{\bar{Y}_A}{f'} \left\{ \frac{1}{2} (1 + \omega) \hat{Y}_{At}^2 + \frac{1}{2} (1 + \omega_p \theta_p) \theta_p \text{Var}_z \hat{p}_t(z) + \frac{1}{2} (1 + \nu \theta_w) \theta_w \phi^{-1} \text{Var}_j \hat{w}_t(j) \right\}$$

Three important points emerge from this approximation. First, even though we wrote the first order term in the expansion of U_t as $\left(1 - v_h \frac{\bar{Y}_A}{f'} \right) \hat{Y}_{At}$, the efficiency of the balanced growth path maintained by fiscal policy implies the equality of the marginal rates of substitution and transformation, $v_h \bar{Y}_A = f'$, so that in fact this term is equal to zero. In other words, movements away from the optimum—which coincides with the efficient equilibrium by the second welfare theorem—do not change the level of utility of the representative agent, up to first order. This is the reason why a *first* order approximation to the equilibrium fluctuations in the endogenous variables is sufficient to compute a *second* order approximation to U_t . If the first order term in the expansion of U_t received a positive weight, second order terms in the approximation of the equilibrium dynamics would generate terms of the same order in the expansion, and could therefore not be omitted. Assuming that fiscal policy eliminates the first order inefficiency generated by monopolistic competition avoids this further complication.

Second, utility depends on the deviations of output Y_t from its *efficient* level $A_t \bar{Y}_A$, the notion of output gap that was proposed in section 2.3 and that appears in the pricing relations derived above. The reason for this is simple. Agent’s utility is decreasing in the amount of hours worked, which in turn depends on the level of effective demand Y_t/A_t . It is then fluctuations in this variable that are going to reduce households’ welfare.

Third, the other fundamental source of welfare losses is the distortion associated with price and wage dispersion. As shown in appendix A, this dispersion can be expressed as a function of inflation in the relevant price indexes, which results in

$$E_{-1} \sum_{t=0}^{\infty} \beta^t U_t = -2\Omega E_{-1} \sum_{t=0}^{\infty} \beta^t L_t + \mathcal{O}(\|\varepsilon\|^3) \quad (17)$$

where

$$L_t \equiv \Lambda_p \left\{ (\tilde{\pi}_t - E_{t-2}\tilde{\pi}_t)^2 + \psi_p (E_{t-2}\tilde{\pi}_t)^2 \right\} + \Lambda_w \left\{ (\tilde{\pi}_t^w - E_{t-2}\tilde{\pi}_t^w)^2 + \psi_w (E_{t-2}\tilde{\pi}_t^w)^2 \right\} + \Lambda_x x_t^2 \quad (18)$$

and the weights are defined in the appendix. Note that in the loss function L_t the level of price dispersion is related to the deviations of price and wage inflation from their indexes, as defined in sections 2.5 and 2.4. This should come as no surprise, since if indexation were the only source of price changes, no price dispersion would result.

2.7 Monetary Policy

This section completes the description of the macroeconomic environment with a discussion of the operations of the monetary authority. A first important point in this respect is that, given the negligible role of money balances in this economy, monetary policy is effective only through the impact of interest rate changes on agents' willingness to substitute their consumption over time. Moreover, because of the informational delays that characterize demand behavior, only the forecastable component of the nominal interest rate has an impact on their spending decisions, as it is evident from Euler equation (12). This implies that, even if the policy authority could collect information with no delay, incorporating this information in its decisions would simply increase the volatility of its instrument, with no other discernible effect on the economy. We will therefore assume that the central bank refrains from exploiting this information, or equivalently, and more realistically, that it shares the same information processing deficiencies as the private sector. As already pointed out, this includes in any case imperfect information on the state of productivity.

As for the authority's procedures to set its instrument, we consider three possible alternatives. The first is that the central bank targets one of three endogenous variables, price inflation, wage inflation or (its forecast of) the output gap. Rather than for their prescriptive content, we consider these policies because they represent useful benchmarks to which to compare the optimal equilibrium. Note also that strict output gap targeting is not within the central bank's menu of available choices due to the one period delay with which demand reacts to changes in the policy instrument. The second alternative is that interest rate policy is set according to a feedback rule of the form

$$\hat{i}_t = \phi_i \hat{i}_{t-1} + (1 - \phi_i) \left[\phi_\pi \pi_t + \frac{1}{4} \phi_x E_{t-1} x_t \right]$$

with parameters fixed at the values estimated by Clarida *et al.* (2000) in a closely related model and summarized in table 6. We take this rule to represent a reasonable approximation to actual U.S. monetary policy and use it to generate the simulated impulse responses that will be compared to those estimated from the data. Finally, we consider the state contingent policy that maximizes the unconditional expectation of the intertemporal loss function (17), under the constraint represented by the model's dynamics.²⁰

²⁰ The exact formulation of the problem and the fundamental steps in its solution can be found in appendix C. Giannoni and Woodford (2002) contains a detailed discussion of the reasons for and implications of considering the unconditional expectation in the optimal policy problem.

Note also that, in all our experiments, we are going to report the values of the losses associated with different policies and/or different models in terms of an “inflation equivalent” (Dennis and Söderström, 2002), defined as the amount of steady inflation that agents would be willing to accept in exchange for the elimination of all fluctuations. This is computed as

$$\pi^* \equiv \sqrt{\frac{(1-\beta)\text{Loss}}{\Lambda_p \psi_p}}$$

and is reported in our figures as an annualized percentage.

2.8 Calibration

The calibration of the model’s parameters, summarized in tables 2 to 5, is based on the estimates of Amato and Laubach (2002), with two important exceptions. The first regards the elasticity of intertemporal substitution in consumption. This coefficient, estimated at $(.26)^{-1}$ by Amato and Laubach (2002), is fixed here at one, the standard value in the real business cycle literature, because of the restriction on utility imposed by the presence of long run growth. This discrepancy is due to the fact that, in this class of models, aggregate consumption is usually interpreted as overall private expenditure, a composite with a much higher sensitivity to interest rate changes than simple non-durable consumption (Rotemberg and Woodford, 1997). Therefore, in our calibration a given adjustment of the policy instrument has a smaller impact on the level of demand than what is more commonly assumed, a difference that should be taken into consideration when interpreting the simulated impulse responses. The second exception regards the assumed degree of stickiness of prices and wages, which the benchmark calibration fixes at the values estimated by ACEL.²¹

To calibrate the productivity parameters, we start by noting that the Wold representation of the state space process in equation (2) can be written as

$$(1 - \rho L) \hat{\gamma}_t^a = (1 - \theta L) \varepsilon_t \tag{19}$$

with $\theta = \rho - K$ and $\text{Var}(\varepsilon_t) = \sigma_a^2 + P$, where K and P are the steady state gain and MSE of the Kalman filter for (2).²² The fundamental disturbance ε_t would then be the structural shock identified by ACEL’s long run restriction if (2) were the true data generating process, as this is the shock inducing the observed unit root in average productivity. Since one of the aims of our positive analysis is to compare our model’s impulse responses to those in ACEL, we first restrict the parameters of (19) to match the moments of the productivity process estimated there. Note though that since ACEL assume productivity growth to follow a simple AR(1) process, we are still missing one restriction, which we can think of as providing information on the “signal to noise ratio” $\eta \equiv \sigma_\gamma / \sigma_a$. To obtain this information we then turn to available empirical estimates of state space models of productivity, like those in Roberts (2001) and French (2001).

²¹ The main reason for favoring the estimates of ACEL is that Amato and Laubach (2002) restrict the two coefficients to be the same in their estimation procedure. This restriction is particularly questionable given the evidence in Christiano *et al.* (2001) that points to wage stickiness as the key nominal friction.

²² See Hamilton (1994, pg. ???) and appendix B for more details.

More specifically, we proceed as follows. First, we match ACEL’s estimated autocorrelation coefficient of 0.8 with the one that would be obtained by fitting an AR(1) process to data generated by (19). This produces a value for ρ of 0.93. Second, we choose the scale of ε_t so that the long run effect on output of a one standard deviation shock to ε_t matches the VAR estimate of 0.6% (see figure 1). This results in an overall volatility of the growth rate of productivity of 0.735%, extremely close to the 0.763% found by Prescott (1986) in estimates based on the Solow residual and that is often used as a benchmark in the real business cycle literature (see for example Cooley and Prescott, 1995). Finally, we calibrate η to reproduce the Kalman gain implicit in Roberts’ (2001) baseline estimates, which are based on a model of labor productivity of the same form as (2), but in which γ_t is assumed to follow a random walk. This results in a value of η of 14%. The reason for following this procedure, rather than for example directly matching the estimated signal to noise ratio, is that the gain provides more direct evidence on agents’ learning speed. In fact, in Roberts’ (2001) I(2) specification, K measures exactly the fraction of each observed change in productivity growth that the Kalman filter attributes to its trend, rather than to its transitory component. Moreover, everything else being equal, reducing the persistence of γ_t reduces the updating speed, while the opposite effect follows from an increase in the signal to noise ratio. With $\rho = 0.93$ then, our parametrization requires a value of η higher than Roberts’ (2001) 8% to reproduce the same gain. To the extent that a low signal to noise ratio helps the model to generate a decline in inflation then, this procedure yields a more conservative estimate of this parameter, in the sense of making it less likely to spuriously produce the desired results. In any event, given the importance of this parameter for the transmission of productivity shocks to costs and inflation, we also checked the robustness of our simulations across a fairly wide range of alternative values. According to these experiments, the paper’s main results are not affected by values of η of up to 50%, which we consider well above the range of reasonable values for this parameter.

To provide a visual illustration of the effect of this calibration on the persistence of the trend shocks and on the learning speed, figure 2 displays impulse responses of the actual and forecasted values of γ_t^a to the shocks, under full and partial information. As we can see, the relatively low value of η results in a very gradual movement of the forecast and in a fairly persistent forecast error, which only disappears after approximately five years.

3 The Transmission of Productivity Shocks

Having introduced the central features of the model, we now turn to the analysis of its equilibrium behavior. In this section, we begin by considering the case of a monetary authority that sets its instrument according to a deterministic policy rule. This allows us to study the transmission mechanism of the productivity shocks built into the behavior of the private sector, and to highlight the role of imperfect information in this mechanism. Then, in section 5, we analyze the optimal policy problem and compare its solution to the prescriptions of some simple targeting rules. We

will show that, even if in the optimum the output gap is nearly stabilized, this does not result in a stable inflation rate. Quite to the contrary, policies that stabilize inflation produce undesirable welfare outcomes. Finally, in sections 4 and 6, we compare the model’s predictions to the observed evolution of some key macroeconomic variables during two much discussed episodes of the recent past, the seventies’ productivity slowdown and the speedup of the last decade.

In what follows, we will mainly focus our attention on the model’s dynamic responses to persistent shifts in the growth rate of productivity. We do this for two reasons. First, from a theoretical perspective, these changes generate the kind of misperception of actual production possibilities that, according to our reasoning, should moderate wage demands and cause them to trail realized productivity. We argued that it is through this channel that a surge in productivity growth might result in lower inflation. Second, from the empirical point of view, these shocks closely fit the description of the productivity disturbances that hit the U.S. economy in the 1970s and 1990s. However, when comparing the simulated impulse responses to those identified through VAR-based long run restrictions, we need to take into account that such restrictions can only recover a linear combination of the model’s structural shocks, the “unconditional” disturbance ε_t in the Wold representation of equation (19). Therefore, for the purpose of that comparison, the model will be perturbed with this particular disturbance, rather than with the theoretically more interesting, but unobservable, trend shock ε_t^γ .

Ideally, to isolate the role of private behavior in the transmission of shocks, we would like to describe their dynamic effects controlling for any systematic response of policy to observed economic conditions. Given the general equilibrium nature of our model though, in which today’s choices depend not only on past conditions, but also on expectations, this is clearly impossible. In other words, there is no obvious “neutral” policy stance whereby the monetary authority can be usefully thought of as “doing nothing”, since defining a rational expectation equilibrium requires to specify some (possibly state contingent) rule of conduct for the central bank. This task is further complicated by the well-known indeterminacy issues that are typical of this class of models (Woodford, 2002). In particular, a peg of the interest rate to any exogenous quantity, like its steady state level, or the exogenously fluctuating natural rate of interest, would result in indeterminacy of the rational expectations equilibrium, therefore ruling out these policies as candidate implementations of a neutral policy stance. As a second best alternative then, we consider a simple class of targeting rules, whose objective is simply to stabilize one of the endogenous variables in the system. As such, these policies do not involve explicit feedback from the economy, which might in turn obscure the direct impact of the productivity shock on private behavior.

Equilibrium dynamics in the benchmark model To introduce the basic features of the model’s equilibrium, figure 3 displays impulse responses of the key endogenous variables to a positive shock to the trend growth rate of productivity, given a policy that stabilizes the one period ahead forecast of the output gap, i.e. $E_{t-1}x_t = 0$. The shock is realized at time $t = 1$ and is assumed

to increase the growth rate of productivity by one percentage point. The impulse responses of the actual and forecasted productivity growth to this particular shock are depicted in figure 2, under full and partial information.²³ Starting from the upper right corner of the figure, we observe that the response of the output gap is a properly scaled mirror image of the persistent forecast error incurred by agents trying to predict the future growth rate, the difference between the continuous and dashed lines in the right panel of figure 2. This follows immediately from equation (12), which implies $x_t - E_{t-1}x_t = -(\hat{\gamma}_t^a - E_{t-1}\hat{\gamma}_t^a)$ and therefore $x_t = -(\hat{\gamma}_t^a - E_{t-1}\hat{\gamma}_t^a)$, given that policy maintains $E_{t-1}x_t = 0$. In economic terms, with demand predetermined at time $t - 1$, an unexpected increase in productivity decreases the amount of labor that firms need to hire to satisfy that demand, thus also decreasing the output gap.²⁴ Lower labor demand and wage setters' conservative estimates of the persistence of the shock restrain the response of wage inflation, which translates in turn into a fairly pronounced deflation in the price of consumption. Finally, this equilibrium is accompanied by a decline in the gap between the nominal interest rate and the natural rate of interest, which is our measure of the nominal stance of policy. Note however from the approximate Euler equation that stabilization of the expected output gap implies that the forecastable component of the real interest rate, $E_{t-1}[i_t - E_t\pi_{t+1}]$, is equal to the natural rate, implying a zero (expected) real interest rate gap. This also confirms that output gap targeting can in some sense be considered a reasonable approximation to a "neutral" policy.

The role of informational frictions As a further illustration of the role of informational frictions in the model's equilibrium dynamics, figure 4 displays impulse responses to the trend shock across several different models, which incorporate various combinations of those frictions. More specifically, the thick dashed line (denoted by ALp-F-s in the legend) represents an extremely simplified version of our model, with no decision delays, full information and indexation of contractual wages to the observed *contemporaneous* growth rate of productivity. This model is included in the simulations as a normative benchmark, since under this specification a policy that targets the (now observable) output gap can attain the unconstrained optimum; it also represents the most direct extension of the standard model of section 1.1 to our non stationary environment. The dashed-dotted line (ALp-F) depicts the responses of a model that reintroduces the decision delays, and indexation to past productivity, but maintains the observability of all productivity shocks. This last assumption is then dropped in the model represented by the continuous thin line (ALp-P). Finally, the thick continuous line represents our baseline specification, with partial information and indexation of wages to forecasted productivity growth. Note also that the simulations assume that policy follows the Volcker-Greenspan interest rate rule, which we argued is an empirically

²³ At this stage, the choice of the shock's size is only a normalization, since we are not yet in the position of quantitatively comparing simulated and empirical impulse responses. This comparison is undertaken in figure 5, in which the model is perturbed by the unconditional shock ε_t .

²⁴ Recall that in this simple model, with labor as the only input and no other shocks besides those to productivity, the output gap is a monotonic transformation of hours. In the log-linear approximation considered here this relation is $\hat{H}_t = \phi x_t$.

satisfactory representation of actual policy in the last twenty years.

As it is evident from the figure, the combination of our assumptions produces a model in which the response of nominal wage inflation to a surge in productivity is significantly dampened with respect to all the other specifications. This implies in turn that, while in those models the response of inflation is either close to zero or significantly positive, in our benchmark specification inflation takes a fairly pronounced and persistent dip. Note in particular that the standard model predicts a sizable positive jump in inflation, which is caused by a loose monetary policy and accompanied by an initial response of nominal wage inflation that amplifies the 1% impulse originated by the increase in productivity growth. Moreover, the figure illustrates that the assumption that informational delays also affect the wage indexation process, once interacted with partial information, provides the key added value in explaining the decline in inflation. It is in fact only moving from the thin to the thick continuous line that inflation is seen to turn negative after the positive shock. This is because indexing wages to past observed productivity, as in the thin line, would undo most of the sluggishness in newly set wages generated by workers' errors in forecasting future productivity growth. This consideration is especially relevant in quantitative terms under our calibration, in which 78% of wages are subject to indexation every period. Ultimately, what assumption represents a more accurate depiction of reality remains an empirical question, which is not the objective of this paper to investigate. Nevertheless, at least when judged in terms of its macroeconomic implications, indexation to forecasted productivity produces results that are clearly more in line with the evidence, explaining our preference for this specification.

Further clues on the effect of the model's informational frictions on the transmission of productivity shocks come from the evolution of the variables depicted in the last line of figure 4. These variables represent the fundamental determinants of price and wage inflation identified by the Phillips curves (13) and (15), namely the "long horizon forecasts" of unit labor costs (ω_{At}) and the output gap (x_t). They are formally defined as

$$E_{t-1} \sum_{s=0}^{\infty} \beta^s z_{t+s} \tag{20}$$

$\forall t$ along the horizon of the impulse response simulation, with $z_t = \ln \frac{W_t}{A_t P_t}$ or x_t . Note that these forecasts are scaled by different coefficients in the two Phillips curves, but their dynamic evolution is captured entirely by expression (20). The responses of these two variables, which reflect the unfolding of agents' uncertainty about the productivity shock, provide a more complete picture of the role of learning in the equilibrium dynamics than the *ex-post* real wage and output gap evolution described above. First of all, we observe that partial information has a significant effect on the forecasts of future labor costs, as suggested by our intuition. This effect does not translate into a more pronounced deflation because in the benchmark model agents also expect a sizable and persistent surge in the output gap. The expectation of positive gaps compensates at least in part the expected decline in labor costs, contributing to restrain the deflationary effects of the productivity shock. This suggests that a richer demand specification, that included for example

habits in consumption (Fuhrer, 2000; Christiano *et al.*, 2001), by further smoothing the response of the output gap, might induce a more pronounced dip in inflation. This would in turn contribute to bringing the model’s dynamics even closer to the data, in which the decline in inflation appears quantitatively more pronounced, as we will see in the empirical comparison below. We leave the quantitative exploration of this effect for future research.

Empirical performance Our analysis so far has been limited to considering the model’s impulse responses to a trend productivity shock, since we argued that this is the theoretically more interesting case. As we already had the chance to remark though, these responses cannot be compared with those estimated by ACEL and reported in figure 1, because their VAR identification procedure, under the null represented by our model, would only isolate a particular linear combination of the model’s two structural shocks, namely the unconditional shock ε_t in the Wold representation of equation (19). Nevertheless, given the distribution of those shocks, we can still simulate the effect of an impulse to ε_t , making the exercise consistent with ACEL’s empirical experiment. Note also that, for the purpose of the quantitative comparison between the model and the data, the scale of the shocks is of course important. As indicated in section 2.8, this is chosen so that the unconditional shock’s long run effect on the level of output is equal to the 0.6% estimated in the VAR (see figure 1), a choice that implies a standard deviation of the growth rate of productivity very closely in line with standard values in the literature.

The findings of this experiment are illustrated in figure 5, which compares the same four models presented above, now under the empirically realistic assumption that monetary policy follows the Volcker-Greenspan interest rate rule. Starting with the upper row of the table, we immediately note that our preferred specification (the thick continuous line) is the only one to produce the negative correlation between output and inflation responses that is a clear feature of the VAR estimates. Moreover, the profile of the two responses is very similar to that depicted in figure 1. Output jumps by approximately 0.4% in the period immediately following the shock (recall that output is predetermined in the model), with only a small adjustment from there to its new long run level. Inflation reaches its trough response of slightly more than -0.2% two quarters after the shock, and adjusts back from there in less than two years. In the VAR point estimates, the maximum response happens on impact and is closer to -0.6% , but the subsequent adjustment back to steady state is similarly sharp. Moreover, the simulated impulse responses fit easily within the VAR 95% confidence bands represented by the shaded area, with the only exception of the impact response.

As for the rest of the variables, the model is fairly successful in replicating the behavior of real wages, that exhibit a smoother profile than the other real variables, and of consumption, that in our model is simply equal to output. The federal funds rate on the contrary is predicted to fall in the model, while the VAR point estimates show it rising, but the model’s prediction is comfortably inside the estimated confidence intervals. Finally, looking at the response of average hours in the last panel, we see that in the model they decline on impact, but then jump immediately to a positive

level, from which they adjust (almost) monotonically towards their steady state. Note that this is broadly in line with the estimated profile, except that there the peak response is only reached after one year, and hours actually never turn negative, at least according to the point estimates. Interestingly, this discrepancy with ACEL's estimates brings our model more in line with several other VAR studies, in which hours are found to decline in response to an identified technology shock (see for example Galí, 1999 and Francis and Ramey, 2001), and whose results have raised some doubts on the empirical relevance of standard RBC models.

As already pointed out, the initial dip in hours depends on the unexpected increase in labor productivity, coupled with a predetermined level of demand. This points in turn to what is perhaps the main weakness of the model, at least when observed through the lens of its responses to productivity shocks. One notable characteristic of the empirical impulse responses is in fact that most variables jump immediately in response to the identified technology shock. Some of them, like inflation, return monotonically to their steady state value, while others, like output, consumption and the real wage are either constant after an initial jump, or slowly increase towards their new steady state level. In the model on the other hand, even if the overall profile of the responses is very close to the VAR's, the initial response is constrained to be zero. This depends in turn on the fact that this class of models was originally devised to match estimated responses to monetary policy shocks. In particular, the informational delays that prevent the endogenous variables from jumping on impact originate from the identifying assumption, common to recursive monetary VARs, that decision variables do not react to changes in policy within the quarter. As it is clear from figure 1, the estimated responses to technology shocks show that this assumption is not quite compatible with the behavior of those variables under all circumstances, and that more research is needed to reconcile these observations.²⁵

4 The New Economy and The Productivity Slowdown

The extraordinary performance of the U.S. economy in the second half of the 1990s, with its unprecedented combination of rapid GDP growth, low unemployment and moderate inflation, and the contemporaneous surge in labor productivity to levels not seen for more than two decades, have attracted a great deal of attention among scholars, policy makers and the public.²⁶ In the midst of the excitement about these developments, some observers hailed the arrival of a “New Economy”, whose most thaumaturgic virtues were promptly disposed of by the recession of 2001. In the meantime however, the acceleration of productivity, the fundamental underpinning of most moderate views of the New Economy, has become a well documented fact, whose quantitative scope has survived the recent recession virtually untouched.²⁷ As extensively documented above,

²⁵ Woodford (2002a) presents a model that can in principle accommodate these different patterns of reaction to different shocks.

²⁶ See for example Gordon (2002), Jorgenson (2001), Oliner and Sichel (2000) and the references therein.

²⁷ The average growth rate of labor productivity from 1948 to the second quarter of 2002 was 2.25%. It was 2.8% from 1948 to 1973, 1.45% from 1973 to 1995 and 2.4% from 1995 to 2002:II (data are from the latest available release

the model proposed in this paper has the potential to generate falling inflation and high growth rates against the backdrop of a persistent increase in the growth rate of labor productivity, precisely the combination of macroeconomic conditions registered in the U.S. in the nineties. It is still an open question though if the model can quantitatively account for the evolution of these variables during that episode.

Answering this question requires first to take a stance on a reasonable chronology for the unfolding of the shock in the data. This task is considerably simplified for the period under study by the fact that 1995 is generally recognized as a year in which the US economy was close to a steady state (see for example Meyer, 2001), while it was in 1996 that the first signs of a possible productivity revival started to emerge. We therefore chose the twenty quarters from 1996:I to 2000:IV as the relevant period for our “case study”, under the assumption that the “New Economy shock” hit the U.S. in the first quarter of 1996 and its effects unfolded over the subsequent five years.

This reconstruction is corroborated for instance by the sequence of testimonies by Chairman Alan Greenspan accompanying the Federal Reserve Board’s semiannual monetary policy report to the Congress.²⁸ The first cautious references to a surge in productivity can be found in the testimony of February 1997, in which we read that “faster productivity growth last year meant that rising compensation gains did not cause labor costs per unit of output to increase any more rapidly”, but it is not until July 1997 that the Fed starts referring to this faster growth as potentially persistent: “although the anecdotal evidence is ample and manufacturing productivity has picked up, a change in the underlying trend is not yet reflected in our conventional data for the whole economy.” A further step in the updating of the Fed’s trend estimates can be detected in July 1998, when Greenspan reports that “evidence continues to mount that the trend of productivity has accelerated, even if the extent of that pickup is as yet unclear.” Finally, in July 1999 the Chairman declares that “to date, 1999 has been an exceptional year for the American economy (...). At the root of this impressive expansion of economic activity has been a marked acceleration in the productivity of our nation’s workforce.”

Note that this sequence of progressively more sanguine statements about the nature of the measured changes in productivity is in accordance with the informational assumptions of our model, according to which policy makers, as well as the public, update their estimates of the persistent component of productivity only gradually over time. Whether the assumed distribution of the productivity factor in the model is also a reasonable characterization of the stochastic properties of the “true” New Economy shock is indeed impossible to ascertain, either from these statements, or from the data, if we assume, as it seems realistic, that the econometrician shares the same information of the agents in the model. The main source of concern in this respect is indeed the degree of persistence of the growth rate shock. The fact that the surge in labor productivity growth that we traced to the beginning of 1996 has not yet significantly subsided suggests that the initial

(Septembr 5th, 2002) by the Bureau of Labor Statistics).

²⁸ Testimonies and Reports from 1996 to 2002 can be found at <http://www.federalreserve.gov/boarddocs/hh/>.

shock might have been even more persistent than what implied by our calibration. But then again, another persistent shock, or maybe a sequence of transitory shocks, might have hit productivity in the meantime. Trying to discriminate among these alternative scenarios is beyond the scope of the present case study, which therefore simply retains the model’s distributional assumptions.

Another key ingredient of this study is the magnitude of the productivity shock to be simulated. To calibrate this quantity, we followed the same procedure used to calibrate the “average” shock considered in section 3. First, we computed the long run impact of the New Economy shock in the data. The average growth rate of output between 1948:I and 1995:IV was 3.39%, but increased to 3.95% between 1996:I and 2000:IV. Compounded over these five years, the extra 56 basis points of annual growth result in a level of output at the end of the period that is 3.36% higher than what it would have been under average growth.²⁹ Assuming that this extra income can be entirely attributed to the acceleration in productivity, this implies an initial shock of approximately 1.2% at an annual rate. Note that even if this number is approximately three times the estimated standard deviation of the permanent shock, making this event highly unlikely, at least on a normal distribution, the growth rate of labor productivity in the U.S. did jump from an average of 1.5% in the period 1973-1995 to an average of 2.5% between 1996 and 2000.

The dynamics of labor productivity, output growth and inflation predicted by the model in response to the technology shock described above are reported in figure 6, along with their observed counterparts, as dashed and continuous lines respectively, under the assumption that monetary policy follows the Volcker-Greenspan estimated interest rate rule. Considering that the continuous lines simply represent the raw data, and that no special effort has been directed at specifying and calibrating the model to match those data, we think that the model captures the broad features of this episode remarkably well. In particular, the model closely matches the dynamics of average labor productivity, with the exception of the last few quarters, in which the monotonic decay of the dashed line falls short of what looks very much like a further resurgence of productivity, similar to that observed at the beginning of the sample. This discrepancy is reflected in turn in the behavior of output growth, that is more persistent in the data than what is predicted by the AR(1) decay built into the model. Similar considerations also apply to inflation, which declines by a full percentage point in the data, but by only half that much according to the simulation.

We should probably remark at this point that it is far from the spirit of the exercise to assess the “fit” of the model according to any formal criterion. It is therefore quite useful, also to provide some perspective on the previous results, to ask how the model compares with observations on the period following the seventies’ productivity slowdown, which we date to the third quarter of 1973.³⁰ Following the same procedure described above, the simulated paths for the variables are generated by a shock whose long run effect on output matches that measured in the data.³¹ As we

²⁹ Recall that the shock considered in the simulations in figure 5 resulted in a long run effect on the level of output of 0.6%.

³⁰ French (2001) conducts Andrews tests on the growth rate of TFP and finds a significant break in 1973:III. However, our results are not overly sensitive to the choice of another break date around this period.

³¹ The average growth rate of GDP between 1948:II and 1973:II was 3.89%, and slowed to 2.73% in the period

can see from figure 7, although the model predicts inflation and real activity to move in the right direction, the amplitude of their observed fluctuations is orders of magnitude wider than in the model. Interestingly however, the simulation correctly identifies the dates of the peak responses in both inflation and output growth, and is fairly close to the observations in the last part of the sample.

We interpret these results as suggesting that it is reasonable to ascribe the macroeconomic events of the second half of the nineties almost exclusively to the effects of the productivity speedup. On the other hand, the seventies' productivity slowdown emerges as a plausible, but quantitatively incomplete explanation for the period's stagflation. The extent to which these results are influenced by our representation of monetary policy, and in particular how the model's optimal equilibrium compares to the observed fluctuations are questions that we defer until section 6.

5 Optimal Policy

Having established that our model provides a reasonably accurate depiction of the transmission of productivity shocks, we are now ready to tackle the normative question posed by the paper's title. In particular, we wish to investigate the nature of the equilibrium fluctuations that minimize the welfare loss of the representative agent and the path of the policy instrument that implements them. We are also going to compare this optimal equilibrium with the ones that would emerge under some of the simple policy rules introduced in section 2.7, to gain further insights into the propagation of productivity shocks and the properties of the optimal policy. We leave instead for further research the question of implementation of the optimal policy and the systematic analysis of simple rules that might approximate this policy.

Optimal policy response to unconditional shocks We begin the normative analysis from figure 8, which considers impulse responses of the endogenous variables to the unconditional shock in the Wold representation of equation (19). We can thus compare the optimal responses both to those obtained under the Volcker-Greenspan interest rate rule, and to the VAR estimates, which are reported in figures 5 and 1 respectively. In our preferred model (the thick line), a drop in the nominal interest rate that more than compensates for the increase in the natural rate of interest is accompanied by a surge in demand, deflation in the price of consumption goods and virtually no change in the nominal wage. The comparison with the other specifications illustrates how the deflationary effect of the productivity shock, which is one of the defining characteristics of our model, and of the data, aides the monetary authority's attempt at balancing the fluctuations in price and wage inflation. With the exception of the simplest model (the dashed line), in which optimal policy actually achieves the unconstrained optimum, the monetary authority responds to the shock by turning distinctly contractionary. This results in a drop in demand, which together with the

1973:III to 1978:II, resulting in a long run effect on output of -6.84%.

higher labor productivity produces a steep and persistent reduction in hours. These adjustments in the real variables are indeed just an instrument through which the monetary authority reaches its objective of minimizing a combination of price and wage dispersion. All models display in fact similar degrees of price deflation and a reduction in wage inflation with respect for example to the corresponding equilibria illustrated in figure 5. In conclusion, figure 8 clearly illustrates how the transmission mechanism embedded in the benchmark model, which we argued is a better representation of reality than more standard sticky price models, has very distinctive implications for the desirable monetary policy response to productivity shocks and for the resulting equilibrium fluctuations in the endogenous variables.

Optimal policy response to trend shocks Turning now our attention to the optimal policy responses to the trend shock ε_t^γ , figure 9 presents the simulated dynamics of the endogenous variables, together with those of the target variables identified by the approximate welfare criterion (18), namely the forecastable and unforecastable components of wage and price inflation, expressed in deviation from their respective indexes.³² We can note immediately that the main source of inefficiency in the models in which wages are assumed to be indexed to past productivity growth (the thin continuous and dashed-dotted lines) comes from the behavior of wage inflation, which compounds high volatility of both its forecastable and unforecastable components. In the attempt to curb this volatility, optimal policy maintains a contractionary stance, inducing a recession that has the effect of moderating workers' wage demands. As a by-product, the recession also produces a drop in price inflation. This explains why it is not desirable to reach for the monetary breaks even more decisively. Such a policy would in fact reduce the welfare losses stemming from wage fluctuations, but cause an even deeper drop in price inflation, increasing the losses through that channel. The contrast with the dynamics just described helps to highlight the role of imperfect information in shaping the optimal equilibrium responses in the benchmark model. In this model, policy responds to the productivity shock by lowering the nominal interest rate below the natural rate, to limit the contractionary impact of the surge in productivity on labor demand and the output gap. Differently from what would happen under the other specifications though, this milder contraction in hours worked does not cause a jump in wage demands because of the slow adjustment of workers' estimates of future productivity. The net result is a price deflation of very similar proportions to that observed under the alternative specifications, but with very different welfare consequences.

One notable implication of the welfare ranking of the optimal policies across the models illustrated in figure 9 is that indexation of wages to observed productivity growth results in higher losses than indexation to its forecast. This might seem surprising, since in a perfectly competitive labor market real wages should efficiently incorporate all shocks to productivity, which is just the

³² The second line of the figure reports the forecastable components of $\tilde{\pi}_t$ and $\tilde{\pi}_t^{wv}$, and not of the measured inflation rates π_t and π_t^w , whose values are instead reported in the first line. This explains why the sum of the responses in the second and third line is not equal to those in the first.

opposite of what happens in our model. The solution to this apparent paradox is in the nature of the approximation to the “true” loss function adopted here, and in particular in the assumption that the economy is close to being efficient. This assumption implies that, up to first order, real wages are in fact indexed to productivity growth, or more precisely that fluctuations in productivity are small enough that the fact that wages are not indexed to its growth rate (but only to its steady state value) does not have first order welfare consequences. As a result, the only sources of welfare losses left in our approximation are the second order distortions emerging from price dispersion. But a high degree of price dispersion, and high losses, are exactly what will result from a model in which newly set wages are based on filtered estimates of future productivity, while contractual wages are indexed to its much more volatile realized value.

5.1 Simple Policy Rules

Targeting rules Our last look at the issue of optimal monetary policy design in response to growth rate shocks is through the window of the targeting rules introduced in section 2.7. Figure 10 compares equilibrium fluctuations and level of welfare, as measured by the “inflation equivalent”, under the optimal policy and three simple policies that target price inflation, wage inflation and the forecast of the output gap.

The first notable aspect of this comparison is that strict inflation targeting is a particularly undesirable policy to insulate the economy from the effect of trend productivity shocks.³³ Differently from what is commonly argued though, this is not because, in the attempt of stabilizing inflation, the central bank might end up “choking” real growth. On the contrary, according to the model, to prevent inflation from falling in response to a positive shock, monetary policy needs to maintain an expansionary stance. This expansion will then result in a boom in demand, which leads workers to set higher wages, increasing their volatility. Of course, from a welfare perspective, the higher volatility in wage inflation is compensated by the complete stabilization of price inflation. Quantitatively, these opposing effects appear roughly equivalent, if we compare the fluctuations under price inflation stabilization (the thick dashed line) to those under wage stabilization (the dashed-dotted line), leaving us with a significant unexplained welfare gap between the two policies. We can reconcile these apparently contradictory observations if we note that, under the benchmark calibration, the variability of wages receives nine times as much weight in the loss function as inflation fluctuations, due to the lower frequency with which wages are expected to adjust (see table 5). This asymmetry reflects in turn the intuitive observation that, for a given amount of volatility in the growth rate of a price index, higher stickiness gives rise to more dispersion in the underlying price distribution (Aoki, 2001; Benigno, 2000). If we are willing to accept this simple intuition as the basis of our welfare criterion, and we trust the empirical evidence of more stickiness in the labor rather than in the goods’ market, the result follows immediately. We find that this example offers

³³ Blanchard (1997) first pointed out the undesirable consequences of strict inflation targeting in a stylized model with wage stickiness. This point was further elaborated by Erceg *et al.* (2000).

an especially compelling illustration of the misleading welfare implications that might be drawn from a welfare criterion not carefully grounded in the microeconomic details of the model economy.

In light of the preceding observations, it is then not surprising that a policy that stabilizes nominal wage inflation around its steady state growth rate γ , would produce better outcomes than those obtained under inflation stabilization. Note also that wage inflation targeting is close to a mirror image of strict inflation targeting in that it requires a contraction in monetary policy to reduce the demand for goods and labor services and workers' wage demands. This avoids the surge in wage inflation that would otherwise follow the positive productivity shock, but also causes a fairly pronounced disinflation. As already pointed out though, this leaves the representative consumer better off because the undesirable fluctuations in inflation can be more easily absorbed, thanks to the relative flexibility of goods' prices.

Role of relative stickiness The role of relative price stickiness is confirmed by figure 11, in which we report impulse responses for the Amato and Laubach (2002) calibration, $\alpha_p = \alpha_w = .66$, which also implies almost equal weights for price and wage inflation in the loss function (see table 5). In this case, optimal policy tries to balance the opposite fluctuations in price and wage inflation, reflecting their similar weights in the welfare function. Note however that, due to wages' higher flexibility, wage inflation grows higher than in the benchmark case, even if the policy contraction is accompanied by a slightly deeper and more prolonged recession. Similarly, inflation fluctuations are at least in part dampened by the higher price stickiness, resulting in higher inflation (i.e. less deflation) regardless of the more pronounced recession.

Output gap forecast targeting Returning now to figure 10, we observe that the policy that stabilizes the forecast of the output gap produces intermediate results between the inflation targeting policies considered above. Under output gap targeting, a modest drop in the interest rate limits the negative impact of the surge in productivity on the output gap, preventing the boom that characterizes inflation targeting as well as the recession that accompanies the stabilization of wage inflation. What is more remarkable though is the fact that output gap forecast targeting closely approximates the optimal policy, almost exactly reproducing its equilibrium fluctuations. This does not at all depend on the direct welfare effect of a stable output gap on the volatility of hours, since this component of the policy objective receives a negligible share of the weight in the utility approximation (see table 5). This result stems rather from the fact that stabilizing the forecast of the output gap achieves (almost) the right balance between fluctuations in the two price indexes that dominate the welfare criterion (Woodford, 2002). Furthermore, this conclusion is robust to the alternative calibration considered in figure 11, even though in this case a slightly more pronounced contraction than that produced by the targeting policy would be desirable. Finally, we should point out that the desirability of output gap targeting, already noted for example by Erceg *et al.* (2000) in their related model with price and wage rigidities, carries over to our framework, in which the output gap is in fact unobservable, under the form of output gap *forecast* targeting.

In other words, even in an economy like ours, in which the monetary authority lacks an accurate measure of the output gap within the quarter, almost optimal results can be obtained by simply targeting the available estimates of the output gap.³⁴

In conclusion, we can summarize the findings in this section by noting that a positive shock to the growth rate of productivity, though temporary, has a permanent effect on the level of productivity, therefore requiring an equivalent increase in the real wage in the long run. According to our analysis, this should be achieved by a combination of positive nominal wage inflation and deflation in the price of consumption. The optimal combination of these two changes depends on the relative stickiness of the two prices, with the more flexible price carrying the bulk of the adjustment.

6 Optimal Policy, Stagflation and the New Economy

The attempt to bring our model to bear on some specific historical episodes, what we referred to as case studies, has been so far only positive in nature. In section 4, we simulated the model's responses to a technology shock assuming that monetary policy followed the interest rate rule estimated by Clarida *et al.* (2000) for the Volcker-Greenspan era, which we argued is a reasonable characterization of actual policy for the post-1979 period. This approach to the case studies left at least two important questions unanswered. First, how close is the Volcker-Greenspan interest rate rule to the optimal policy prescribed by the model for the second half of the 1990s? and, second, does the failure of the model to reproduce the dynamic properties of output and inflation in the 1970s depend on the assumed specification of the interest rate rule? We take up these questions in turn in the rest of the section.

Figure 12 provides a clear answer to the first question: actual policy in the 1990s was very close to optimal according to our model. In fact, optimal fluctuations in output, and especially in inflation, are significantly closer to the data than those implied by the interest rate rule underlying figure 6, and the same is true for the Federal Funds rate, especially in the first part of the sample.³⁵ Interestingly, the policy loosening of 1998:IV corresponds to the first significant departure of measured growth from the model's prediction, while the tightening of the last three quarters of the sample, which exceeds the model's prescription, is associated with a pronounced decline in the growth rate of real activity. In conclusion, our analysis corroborates the widespread opinion that the Federal Reserve skillfully steered the U.S. economy of the second half of the nineties in the face of the non trivial challenges posed by a significant shift in the growth rate of productivity.

Turning now to the period following the 1973 productivity slowdown, we recall from figure 7 that the model falls significantly short of replicating the extreme volatility found in the data, especially

³⁴ Interestingly, Amato and Laubach (2002), in a model much closer to ours than that of Erceg *et al.* (2000), find that "a strong interest rate response to the output gap can lead to severely sub-optimal outcomes." They however do not explicitly consider output gap targeting rules.

³⁵ The dip in the interest rate in 1998:IV corresponds to the loosening of policy that followed the Russian debt crisis and the default of LTCM in the summer of 1998.

in inflation. This is indeed not too surprising, especially if we contrast the model’s extremely restricted set of shocks with the turbulence in the fundamentals that is one of the period’s most salient features (Primiceri, 2002; Stock and Watson, 2002). One popular explanation for this turbulence has been advanced by Clarida *et al.* (2000), who argue that monetary policy in the pre-Volcker period did not fulfill the “Taylor principle”, resulting in indeterminacy of the rational expectations equilibrium and sunspot fluctuations. Since our solution algorithm cannot simulate an indeterminate system, we cannot directly test this theory within our model. Nevertheless, we can investigate the role of different parametrizations of the interest rate rule on the model’s equilibrium, at least as long as they guarantee its existence and uniqueness. Figure 13 then displays simulated and actual data on GDP growth and inflation under the pre-Volcker rule in table 6. This is the rule estimated by Clarida *et al.* (2000) for the period 1960:I to 1979:II, except for the coefficient on inflation, which we assume takes the value closest to the estimate, but that is still compatible with equilibrium determinacy.³⁶ Surprisingly, this rule dampens the predicted fluctuations in inflation when compared to the Volcker-Greenspan rule (figure 7), even leading to a fall in inflation in the second part of the sample. This evidence cannot of course rule out the interpretation of the data proposed by Clarida *et al.* (2000), since our model is silent on what would happen if the coefficient on inflation in the policy reaction function were in fact lower than one. Nevertheless, this exercise suggests that the failure of the model to account for the quantitative extent of the observed fluctuations does not depend on the kind of interest rate rule assumed in figure 7, at least as long as we restrict our analysis to the class of determinate rules.

There is however something more that we can learn from the model about the role of monetary policy in the economic landscape of the seventies. To this effect, figure 14 contains a fairly detailed comparison between the data and the model’s optimal equilibrium, which includes inflation, GDP growth, real wage inflation and the ex-post real interest rate. This picture confirms that the observed fluctuations in real activity and inflation were indeed much wider than those produced by the model, when perturbed by only one technology shock. Interestingly though, if we concentrate only on the first few quarters of the period, up to the beginning of 1975, we observe a combination of tight monetary policy, depressed output, a real wage inflation that is broadly in line with the model’s prediction, but still high inflation. This particular macroeconomic configuration is at least in principle compatible with the traditional idea that the run up in inflation was the product of the oil price shock of 1973, to which monetary policy clearly appears to have responded fairly vigorously at the outset. High real interest rates in the face of a “cost-push” shock could in turn explain the precipitous drop in demand and the high level of inflation, in the absence of significant cost pressures originating from the labor markets. On the other hand, the subsidence of the commodity price shock in 1975, and the concomitant loosening of policy, might explain the decline in inflation, the steep recovery of output and the subsequent build up of cost pressures in the labor market. Note also that policy did not turn loose in comparison to the model’s prescription until the end of 1976, at

³⁶ Clarida *et al.* (2000) estimate a value for the reaction coefficient to inflation of 0.83, with a standard deviation of 0.07.

which point inflation started inching up again. We hasten to emphasize that this reconstruction of the events is highly speculative, even though roughly in line with the chronology of the commodity price shocks emphasized by Blinder (1979), and with some of the conclusions of De Long’s (1997) influential work on the subject.³⁷ Whether augmenting the model with an exogenous source of cost pressure, like a commodity price shock, could help to reconcile quantitatively its predictions with the data is a question that we leave for future research.

7 Conclusions

In this paper, we investigated the relationship between monetary policy and productivity growth, from both a positive and a normative perspective. We argued that standard New Keynesian models, built around the assumption of a perfectly competitive labor market, do not capture a key feature of the data, the negative correlation between inflation and real activity conditional on productivity shocks. To solve this problem, we proposed a model in which wages respond to persistent growth rate shocks with a delay. This is due to a particular form of imperfect information, whereby agents, after observing a shock, cannot conclusively establish its degree of persistence. As a result, when the shock is in fact persistent, workers come to this realization only gradually, so that wage demands lag behind realized productivity, moderating the reaction of marginal costs and inflation.

An extensive investigation of the model’s equilibrium behavior documented that it represents a significant improvement over the existing literature. The model’s dynamics following a productivity shock, calibrated to match that identified by Altig *et al.* (2002) within a structural VAR, are qualitatively and quantitatively in line with the estimated responses. In particular, the model is able to reproduce the conditionally negative correlation between inflation and output that is one of the salient features of the data. Moreover, it provides a remarkably close account of the evolution of labor productivity, inflation, output and interest rates that followed the “New Economy shock” of the mid-nineties. All these considerations lead us to conclude that the model safely qualifies as an adequate laboratory to study the optimal policy response to technology shocks.

In the normative part of the analysis then, we found that this response is characterized by deflation in the goods’ market and by a moderate amount of nominal wage inflation, at least when compared with the predictions of the standard model. The economic nature of this response can be summarized by noting that cointegration of income and real wages implies that a permanent increase in productivity must be followed by an equivalent movement in real wages. In the optimum, this is achieved by an increase in the nominal wage, accompanied by a fall in the price of consumption, with the less sticky of the two price indexes carrying the bulk of the adjustment. In our benchmark model, this adjustment is engineered by the monetary authority through a reduction in the nominal interest rate. In other words, the popular argument that monetary policy should “accommodate” a

³⁷ DeLong (1997) writes for example: “(...) I would tentatively conclude that the supply shocks of the 1970s were in large part sound and fury. They came. They caused headlines, jumps in the price level, and recessions. And when they had passed the inflation situation was much the same as before their impact.”

surge in the growth rate of productivity survives the scrutiny of our formal analysis. Nevertheless, a reduction in nominal rates is optimal here not because it allows higher productivity to be turned into higher output, but rather because it contributes to balance wage and price adjustments, minimizing the distortions associated with price dispersion. Not surprisingly then, a policy of strict inflation targeting, forcing the adjustment of real wages to fall entirely on nominal wage changes, results in undesirable welfare consequences. Finally, we argued that the pickup in productivity was the predominant influence on the observed fluctuations of output and inflation in the second half of the 1990s, and that the stance of American monetary policy in that same period was very close to that prescribed by the model.

A Derivation of the Welfare Criterion

This appendix derives expressions for the price and wage dispersion measures $\text{Var}_z \hat{p}_t(z)$ and $\text{Var}_j \hat{w}_t(j)$ as a function of inflation in the associated price indexes, following the steps described in Rotemberg and Woodford (1998).

Starting from the price dispersion measure, we define $\bar{p}_t \equiv E_z \ln p_t(z)$, $\bar{p}_{t-1}^i \equiv \bar{p}_{t-1} + \lambda_p \pi_{t-1}$ and $\Delta_t^p \equiv \text{Var}_z \ln p_t(z) = \text{Var}_z \hat{p}_t(z)$ and write

$$\begin{aligned} \Delta_t^p &= \text{Var}_z [\ln p_t(z) - \bar{p}_{t-1}^i] = E_z \left[(\ln p_t(z) - \bar{p}_{t-1}^i)^2 \right] - (E_z \ln p_t(z) - \bar{p}_{t-1}^i)^2 \\ &= \alpha_p E_z \left[(\ln p_{t-1}(z) + \lambda_p \pi_{t-1} - \bar{p}_{t-1}^i)^2 \right] + (1 - \alpha_p) \tilde{\psi}_p [\ln p_t^1 - \bar{p}_{t-1}^i]^2 \\ &\quad + (1 - \alpha_p) \left(1 - \tilde{\psi}_p \right) [\ln p_t^2 - \bar{p}_{t-1}^i]^2 - (\Delta \bar{p}_t - \lambda_p \pi_{t-1})^2 \\ &= \alpha_p \Delta_{t-1}^p + (1 - \alpha_p) \left\{ \tilde{\psi}_p [\ln p_t^1 - \bar{p}_{t-1}^i]^2 + \left(1 - \tilde{\psi}_p \right) [\ln p_t^2 - \bar{p}_{t-1}^i]^2 \right\} - (\Delta \bar{p}_t - \lambda_p \pi_{t-1})^2 \end{aligned} \quad (21)$$

and

$$\begin{aligned} \bar{p}_t - \bar{p}_{t-1}^i &= E_z [\ln p_t(z) - \bar{p}_{t-1}^i] \\ &= \alpha_p E_z [\ln p_{t-1}(z) - \bar{p}_{t-1}] + (1 - \alpha_p) \tilde{\psi}_p [\ln p_t^1 - \bar{p}_{t-1}^i] \\ &\quad + (1 - \alpha_p) \left(1 - \tilde{\psi}_p \right) [\ln p_t^2 - \bar{p}_{t-1}^i] \\ &= (1 - \alpha_p) \left\{ \tilde{\psi}_p [\ln p_t^1 - \bar{p}_{t-1}^i] + \left(1 - \tilde{\psi}_p \right) [\ln p_t^2 - \bar{p}_{t-1}^i] \right\} \end{aligned} \quad (22)$$

Taking expectations at $t-2$, and remembering that (up to our log-linear approximation), $\ln p_t^2 = E_{t-2} \ln p_t^1$ this becomes

$$E_{t-2} [\bar{p}_t - \bar{p}_{t-1}^i] = (1 - \alpha_p) [\ln p_t^2 - \bar{p}_{t-1}^i] + \mathcal{O}(\|\varepsilon\|^2)$$

which substituted back into (22) yields

$$\bar{p}_t - \bar{p}_{t-1}^i - \left(1 - \tilde{\psi}_p \right) E_{t-2} [\bar{p}_t - \bar{p}_{t-1}^i] = (1 - \alpha_p) \tilde{\psi}_p [\ln p_t^1 - \bar{p}_{t-1}^i] + \mathcal{O}(\|\varepsilon\|^2)$$

Squaring these last two expression and using $\Delta \bar{p}_t = \pi_t + \mathcal{O}(\|\varepsilon\|^2)$ we then obtain

$$[\ln p_t^2 - \bar{p}_{t-1}^i]^2 = (1 - \alpha_p)^{-2} (E_{t-2} \tilde{\pi}_t)^2$$

and

$$[\ln p_t^1 - \bar{p}_{t-1}^i]^2 = (1 - \alpha_p)^{-2} \tilde{\psi}_p^{-2} \left(\tilde{\pi}_t - E_{t-2} \tilde{\pi}_t + \tilde{\psi}_p E_{t-2} \tilde{\pi}_t \right)^2$$

where we are suppressing terms of third order or higher. Substituting this expression into (21) yields an AR(1) representation for the dispersion measure

$$\begin{aligned} \Delta_t^p - \alpha_p \Delta_{t-1}^p &= (1 - \alpha_p)^{-1} \tilde{\psi}_p^{-1} \left(\tilde{\pi}_t - E_{t-2} \tilde{\pi}_t + \tilde{\psi}_p E_{t-2} \tilde{\pi}_t \right)^2 + \frac{1 - \tilde{\psi}_p}{1 - \alpha_p} (E_{t-2} \tilde{\pi}_t)^2 - \tilde{\pi}_t^2 \\ &= \left(\frac{1 - \tilde{\psi}_p (1 - \alpha_p)}{\tilde{\psi}_p (1 - \alpha_p)} + 1 \right) (\tilde{\pi}_t - E_{t-2} \tilde{\pi}_t)^2 + \frac{1}{(1 - \alpha_p)} (E_{t-2} \tilde{\pi}_t)^2 - \tilde{\pi}_t^2 + unf \\ &= \frac{\alpha_p}{(1 - \alpha_p) \tilde{\psi}_p} \left[\tilde{\pi}_t^2 - (1 - \tilde{\psi}_p) (E_{t-2} \tilde{\pi}_t)^2 \right] + unf \end{aligned}$$

where unf collects terms that cannot be forecasted on the basis of information available at time -1 .

As for wage dispersion, we proceed analogously and after defining $\bar{w}_t \equiv E_j \ln w_t(j)$, $\bar{w}_{t-1}^i \equiv \bar{w}_{t-1} + E_{t-2} \ln \Omega_t$ and $\Delta_t^w \equiv \text{Var}_j \ln w_t(j) = \text{Var}_j \hat{w}_t(j)$ we write

$$\begin{aligned} \Delta_t^w &= \text{Var}_z [\ln w_t(z) - \bar{w}_{t-1}^i] \\ &= \alpha_w \Delta_{t-1}^w + (1 - \alpha_w) \left\{ \tilde{\psi}_w [\ln w_t^1 - \bar{w}_{t-1}^i]^2 + (1 - \tilde{\psi}_w) [\ln w_t^2 - \bar{w}_{t-1}^i]^2 \right\} \\ &\quad - (\Delta \bar{w}_t - E_{t-2} \ln \Omega_{t-1})^2 \end{aligned} \quad (23)$$

and

$$\bar{w}_t - \bar{w}_{t-1}^i = (1 - \alpha_w) \left\{ \tilde{\psi}_w [\ln w_t^1 - \bar{w}_{t-1}^i] + (1 - \tilde{\psi}_w) [\ln w_t^2 - \bar{w}_{t-1}^i] \right\}$$

Taking expectations at $t - 2$

$$E_{t-2} \tilde{\pi}_t^w = (1 - \alpha_w) [\ln w_t^2 - \bar{w}_{t-1}^i] + \mathcal{O}(\|\varepsilon\|^2)$$

which we again substitute back to obtain

$$\tilde{\pi}_t^w - (1 - \tilde{\psi}_w) E_{t-2} \tilde{\pi}_t^w = (1 - \alpha_w) \tilde{\psi}_w [\ln w_t^1 - \bar{w}_{t-1}^i] + \mathcal{O}(\|\varepsilon\|^2)$$

where we now used $\ln w_t^2 = E_{t-2} \ln w_t^1 + \mathcal{O}(\|\varepsilon\|^2)$ and $\Delta \bar{w}_t = \pi_t^w + \mathcal{O}(\|\varepsilon\|^2)$. Squaring these last two expressions and plugging the results into (23) we then find

$$\frac{1 - \alpha_w}{\alpha_w} \psi_w (\Delta_t^w - \alpha_w \Delta_{t-1}^w) = (\tilde{\pi}_t^w)^2 - (1 - \psi_w) (E_{t-2} \tilde{\pi}_t^w)^2 + unf$$

which, following Woodford (2002, chapter 6), can be solved backwards to yield

$$\frac{1 - \alpha_w}{\alpha_w} \psi_w (1 - \alpha_w \beta) E_{-1} \sum_{t=0}^{\infty} \beta^t \Delta_t^w = E_{-1} \sum_{t=0}^{\infty} \beta^t \left\{ (\tilde{\pi}_t^w)^2 - (1 - \psi_w) (E_{t-2} \tilde{\pi}_t^w)^2 \right\}$$

with an analogous expression for Δ_t^p

$$\frac{1 - \alpha_p}{\alpha_p} \psi_p (1 - \alpha_p \beta) E_{-1} \sum_{t=0}^{\infty} \beta^t \Delta_t^p = E_{-1} \sum_{t=0}^{\infty} \beta^t \left\{ (\tilde{\pi}_t^p)^2 - (1 - \psi_p) (E_{t-2} \tilde{\pi}_t^p)^2 \right\}$$

Putting everything together produces the period loss function

$$L_t \equiv \Lambda_p \left\{ \tilde{\pi}_t^2 - (1 - \psi_p) (E_{t-2} \tilde{\pi}_t)^2 \right\} + \Lambda_w \left\{ (\tilde{\pi}_t^w)^2 - (1 - \psi_w) (E_{t-2} \tilde{\pi}_t^w)^2 \right\} + \Lambda_x x_t^2$$

where the weights, normalized so that $\Lambda_p + \Lambda_w = 1$, are

$$\Lambda_p \equiv \frac{\theta_p (\xi_p \psi_p)^{-1}}{\Omega} ; \quad \Lambda_w \equiv \frac{\theta_w (\xi_w \psi_w \phi)^{-1}}{\Omega} ; \quad \Lambda_x \equiv \frac{1 + \omega}{\Omega}$$

B Solution and Filtering

Even though the paper reports results only for the parametrization of the productivity process discussed in section 2.8, our solution procedure can accommodate a much more general class of processes. This section discusses the outline of that procedure under these more general assumptions.

We consider a stochastic process for the growth rate of productivity of the form

$$\begin{aligned}\hat{\gamma}_t^a &= \hat{\gamma}_t + \mu^a(L) \varepsilon_t^a \\ \hat{\gamma}_{t+1} &= \rho(L) \hat{\gamma}_t + \mu^\gamma(L) \varepsilon_{t+1}^\gamma\end{aligned}$$

where ε_t^a and ε_t^γ are i.i.d. shocks with $E[\varepsilon_t^a] = E[\varepsilon_t^\gamma] = E[\varepsilon_t^a \varepsilon_t^\gamma] = 0$, $\text{Var}[\varepsilon_t^a] = \sigma_c^2$ and $\text{Var}[\varepsilon_t^\gamma] = \sigma_\gamma^2$, while $\rho(L)$, $\mu^a(L)$ and $\mu^\gamma(L)$ are polynomials in the lag operator. Following Hamilton (1994, section 13.1), this process can be represented in state space form as

$$\begin{aligned}\hat{\gamma}_t^a &= \mu' \xi_t \\ \xi_{t+1} &= \Xi \xi_t + \Theta \epsilon_{t+1}\end{aligned}\tag{24}$$

where μ , Ξ and Θ are appropriately defined matrices of parameters and $\epsilon_t' \equiv \begin{bmatrix} \varepsilon_t^a & \varepsilon_t^\gamma \end{bmatrix}$.

Given this representation, we can run a standard Kalman filter to forecast the growth rate of productivity as

$$\begin{aligned}\hat{\gamma}_{t|t-1}^a &= \mu' \xi_{t|t-1} \\ \xi_{t|t} &= \xi_{t|t-1} + K \mu' (\xi_t - \xi_{t|t-1}) \\ \xi_{t+1|t} &= \Xi \xi_{t|t}\end{aligned}$$

where the Kalman gain is

$$K = P \mu (\mu' P \mu)^{-1}$$

and $P \equiv E \left[\left(\xi_{t+1} - \xi_{t+1|t} \right) \left(\xi_{t+1} - \xi_{t+1|t} \right)' \right]$ solves the usual Riccati equation

$$P = \Xi \left[P - P \mu (\mu' P \mu)^{-1} \mu' P \right] \Xi' + \Theta \Sigma \Theta'$$

with $\Sigma \equiv E[\epsilon_t \epsilon_t']$. Note that in the full information case the state vector would simply be forecasted as

$$\xi_{t+1|t} = \Xi \xi_t$$

Defining now an appropriate partition of the vector of endogenous variables y_t between jumps (x_t) and states (X_t), we express the system of linearized first order conditions in the canonical form

$$\begin{bmatrix} A^0 E_t y_{t+1} \\ \xi_{t+1} \end{bmatrix} = \begin{bmatrix} B^1 & B^2 \\ 0_{n_\xi \times n_y} & \Xi \end{bmatrix} \begin{bmatrix} y_t \\ \xi_t \end{bmatrix} + \begin{bmatrix} B^3 \\ 0_{n_\xi \times n_\xi} \end{bmatrix} \xi_{t|t} + \begin{bmatrix} 0_{n_y \times n_\epsilon} \\ \Theta \end{bmatrix} \epsilon_{t+1}$$

Taking expectations at time t , this becomes

$$AE_t \begin{bmatrix} y_{t+1} \\ \xi_{t+1|t} \end{bmatrix} = B \begin{bmatrix} y_t \\ \xi_{t|t} \end{bmatrix}$$

This system can now be solved with standard methods to produce a state space representation

$$\begin{bmatrix} y_t \\ \xi_{t|t} \end{bmatrix} = \begin{bmatrix} D^1 & D^2 \\ 0_{n_\xi \times n_X} & I_{n_\xi} \end{bmatrix} \begin{bmatrix} X_t \\ \xi_{t|t} \end{bmatrix}$$

$$\begin{bmatrix} X_{t+1} \\ \xi_{t+1|t} \end{bmatrix} = \begin{bmatrix} G^1 & G^2 \\ 0_{n_\xi \times n_X} & \Xi \end{bmatrix} \begin{bmatrix} X_t \\ \xi_{t|t} \end{bmatrix}$$

which, together with

$$\xi_{t|t} = \xi_{t|t-1} + K\mu' (\xi_t - \xi_{t|t-1}) = K\mu' \xi_t + (I - K\mu') \xi_{t|t-1}$$

and (24), we write recursively as

$$y_t = \begin{bmatrix} D^1 & D^2(I - K\mu') & D^2K\mu' \end{bmatrix} \begin{bmatrix} X_t \\ \xi_{t|t-1} \\ \xi_t \end{bmatrix}$$

$$\begin{bmatrix} X_{t+1} \\ \xi_{t+1|t} \\ \xi_{t+1} \end{bmatrix} = \begin{bmatrix} G^1 & G^2(I_{n_\xi} - K\mu') & G^2K\mu' \\ 0_{n_\xi \times n_X} & \Xi(I_{n_\xi} - K\mu') & \Xi K\mu' \\ 0_{n_\xi \times n_X} & 0_{n_\xi \times n_\xi} & \Xi \end{bmatrix} \begin{bmatrix} X_t \\ \xi_{t|t-1} \\ \xi_t \end{bmatrix} + \begin{bmatrix} 0_{n_X \times n_\epsilon} \\ 0_{n_\xi \times n_\epsilon} \\ \Theta \end{bmatrix} \epsilon_{t+1}$$

or, more compactly

$$y_t = DX_t^\xi$$

$$X_{t+1}^\xi = GX_t^\xi + H\epsilon_{t+1}$$

with the obvious definition of the extended state X_t^ξ .

Note also that under full information we simply have $K\mu' = I_{n_\xi}$, which in turn simplifies the matrices above to

$$D \equiv \begin{bmatrix} D^1 & 0_{n_y \times n_\xi} & D^2 \end{bmatrix}; \quad G \equiv \begin{bmatrix} G^1 & 0_{n_X \times n_\xi} & G^2 \\ 0_{n_\xi \times n_X} & 0_{n_\xi \times n_\xi} & \Xi \\ 0_{n_\xi \times n_X} & 0_{n_\xi \times n_\xi} & \Xi \end{bmatrix}$$

C Computing the Optimal Policy

Given the particularly simple form of the objective function L_t , the optimal policy problem can be formulated as

$$\begin{aligned} \min_{\{i_{t,1}\}_{t=0}^{\infty}} \quad & E \sum_{t=0}^{\infty} \beta^t \{x_t' W_{xx} x_t + X_t' W_{XX} X_t\} \\ \text{s.t.} \quad & A_{xx}^0 E_t x_{t+1} + A_{xX}^0 X_{t+1} = B_{xx}^1 x_t + B_{xX}^1 X_t + B_{xi} \hat{i}_{t,1} \\ & A_{XX}^0 X_{t+1} = B_{Xx}^1 x_t + B_{XX}^1 X_t + B_{X\xi}^2 \xi_t + B_{X\xi}^3 \xi_{t|t} \\ & \xi_{t+1} = \Xi \xi_t + \Theta \epsilon_{t+1} \end{aligned}$$

The Lagrangian for this problem is therefore

$$\begin{aligned} \mathcal{L} = \quad & E \sum_{t=0}^{\infty} \beta^t \{x_t' W_{xx} x_t + X_t' W_{XX} X_t \\ & + 2\varphi_{t+1}^{x'} (A_{xx}^0 x_{t+1} + A_{xX}^0 X_{t+1} - B_{xx}^1 x_t - B_{xX}^1 X_t - B_{xi} \hat{i}_{t,1}) \\ & + 2\varphi_{t+1}^{X'} (A_{XX}^0 X_{t+1} - B_{Xx}^1 x_t - B_{XX}^1 X_t - B_{X\xi}^2 \xi_t - B_{X\xi}^3 \xi_{t|t}) \\ & + 2\varphi_{t+1}^{\xi'} (\xi_{t+1} - \Xi \xi_t - \Theta \epsilon_{t+1}) \end{aligned}$$

where it should be noted that (differently from the more general case treated in Svensson and Woodford, 2002), $\varphi_{t+1}^{x'}$ and $\varphi_{t+1}^{X'}$ are both measurable with respect to I_t (the information set including the endogenous variables and $\{A_{\tau}\}_{\tau \leq t}$), but not I_t^f (the full information set also including $\{\varepsilon_{\tau}^a, \varepsilon_{\tau}^{\gamma}\}_{\tau \leq t}$), while $\varphi_{t+1}^{\xi'}$ is measurable with respect to I_{t+1}^f information.

The first order conditions for this problem are

$$\begin{aligned} \hat{i}_{t,1} \quad & : \quad 0 = B_{xi}' \varphi_{t+1}^x \\ x_t' \quad & : \quad 0 = W_{xx} x_t + \beta^{-1} A_{xx}' \varphi_t^x - B_{xx}' \varphi_{t+1}^x - B_{xX}' E_t \varphi_{t+1}^X \\ X_t' \quad & : \quad 0 = W_{XX} X_t + \beta^{-1} A_{xX}' \varphi_t^x - B_{xX}' \varphi_{t+1}^x + \beta^{-1} A_{XX}' \varphi_t^X - B_{XX}' E_t \varphi_{t+1}^X \\ \xi_t' \quad & : \quad 0 = \beta^{-1} \varphi_t^{\xi} - (B_{X\xi}^2 + B_{X\xi}^3)' E_t \varphi_{t+1}^X - \Xi' E_t^f \varphi_{t+1}^{\xi} \\ \varphi_{t+1}^{x'} \quad & : \quad 0 = A_{xx} E_t x_{t+1} + A_{xX} X_{t+1} - B_{xx} x_t - B_{xX} X_t - B_{xi} \hat{i}_{t,1} \\ \varphi_{t+1}^{X'} \quad & : \quad 0 = A_{XX} X_{t+1} - B_{Xx} x_t - B_{XX} X_t - B_{X\xi}^2 \xi_t - B_{X\xi}^3 \xi_{t|t} \\ \varphi_{t+1}^{\xi'} \quad & : \quad 0 = \xi_{t+1} - \Xi \xi_t - \Theta \epsilon_{t+1} \end{aligned}$$

which, after taking expectations with respect to I_t information, can be cast in first order form as

$$\bar{A} E_t \begin{bmatrix} \bar{y}_{t+1} \\ \xi_{t+1|t} \end{bmatrix} = \bar{B} \begin{bmatrix} \bar{y}_{t|t} \\ \xi_{t|t} \end{bmatrix}$$

where the vector \bar{y}_t now also contains the multipliers. The solution to this system is again of the form

$$\begin{aligned} y_t &= \bar{D} \bar{X}_t^{\xi} \\ \bar{X}_{t+1}^{\xi} &= \bar{G} \bar{X}_t^{\xi} + \bar{H} \epsilon_{t+1} \end{aligned}$$

where $\bar{X}_t^\xi \equiv \begin{bmatrix} X_t & \varphi_t^x & \xi_{t|t-1} & \xi_t \end{bmatrix}$.

From this state space representation we can then derive

$$\bar{X}_t^\xi = \sum_{\tau=0}^{t-1} \bar{G}^\tau \bar{H} \epsilon_{t-\tau} + \bar{G}^t \bar{X}_0^\xi$$

and

$$\Sigma_{\bar{X}^\xi} \equiv E \left[\bar{X}_t^\xi \bar{X}_t^{\xi'} \right] = \bar{G} \Sigma_{\bar{X}^\xi} \bar{G}' + \bar{H} \Sigma \bar{H}'$$

which we substitute into the loss function to obtain the value of the optimal plan

$$\begin{aligned} \text{Loss} &= E \sum_{t=0}^{\infty} \beta^t y_t' W y_t \\ &= E \sum_{t=0}^{\infty} \beta^t \left[\bar{D} \left(\sum_{\tau=0}^{t-1} \bar{G}^\tau \bar{H} \epsilon_{t-\tau} + \bar{G}^t \bar{X}_0^\xi \right) \right]' \bar{W} \left[\bar{D} \left(\sum_{\tau=0}^{t-1} \bar{G}^\tau \bar{H} \epsilon_{t-\tau} + \bar{G}^t \bar{X}_0^\xi \right) \right] \\ &= \sum_{t=0}^{\infty} \beta^t \left\{ \sum_{\tau=0}^{t-1} \text{tr} \left[\bar{G}^{\tau'} \bar{D}' W \bar{D} \bar{G}^\tau \bar{H} E \left[\epsilon_{t-\tau} \epsilon_{t-\tau}' \right] \bar{H}' \right] + \text{tr} \left[\bar{G}^{t'} \bar{D}' W \bar{D} \bar{G}^t E \left[\bar{X}_0^\xi \bar{X}_0^{\xi'} \right] \right] \right\} \\ &= \frac{1}{1-\beta} \sum_{t=0}^{\infty} \text{tr} \left[\beta^t \bar{G}^{t'} \bar{D}' W \bar{D} \bar{G}^t \bar{H} \Sigma \bar{H}' \right] + \sum_{t=0}^{\infty} \text{tr} \left[\beta^t \bar{G}^{t'} \bar{D}' W \bar{D} \bar{G}^t \Sigma_{\bar{X}^\xi} \right] \\ &= \frac{1}{1-\beta} \text{tr} \left[\beta \bar{L} \bar{H} \Sigma \bar{H}' \right] + \text{tr} \left[\bar{L} \Sigma_{\bar{X}^\xi} \right] \end{aligned}$$

where \bar{L} solves the Sylvester equation

$$\bar{L} = \beta \bar{G}' \bar{L} \bar{G} + \bar{D}' W \bar{D}$$

References

- Altig, David, Lawrence J. Christiano, Martin Eichenbaum and Jesper Linde (2002), “Technology Shocks and Aggregate Fluctuations,” *manuscript*
- Amato, Jeffrey D. and Thomas Laubach (2002), “Estimation and Control of an Optimization-Based Model with Sticky Prices and Wages,” *Journal of Economic Dynamics and Control*, forthcoming
- Aoki, Kosuke (2001), “Optimal Monetary Policy Responses to Relative-Price Changes” *Journal of Monetary Economics* 48(1): 55-80.
- Ball, Laurence and Robert Moffitt (2001), “Productivity Growth and the Phillips Curve,” in A.B. Krueger and R. Solow, eds. *The Roaring Nineties: Can Full Employment be Sustained?* New York: The Russell Sage Foundation and The Century Foundation Press.
- Benigno, Pierpaolo (2000), “Optimal Monetary Policy in a Currency Area” CEPR Discussion Paper No. 2755.
- Bernanke, Ben S. and Alan Blinder (1992), “The Federal Funds Rate and the Channels of Monetary Transmission,” *American Economic Review* 82(4): 901-21.
- Bernanke, Ben S. and Ilian Mihov (1998), “Measuring Monetary Policy,” *Quarterly Journal of Economics* 113(3): 869-902.
- Blanchard, Olivier Jean (1997), “Comment”, *NBER Macroeconomics Annual* 12: 289-293.
- Blinder, Alan (1979), *Economic Policy and the Great Stagflation*, New York: Academic Press.
- Blinder, Alan (2000), “The Internet and the New Economy,” Brookings Institution Policy Brief 60.
- Boivin, Jean and Marc Giannoni (2002), “Has Monetary Policy Become Less Powerful?” Federal Reserve Bank of New York, Staff Report no. 144.
- Braun, Steven N. (1984), “Productivity and the NIIRU (and Other Phillips Curve Issues),” Working Paper 34, Federal Reserve Board Economic Activity Section.
- Calvo, Guillermo A. (1983), “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics* 12(3): 383-398.
- Christiano, Lawrence J., Martin Eichenbaum and Charles L. Evans (1999), “Monetary Policy Shocks: What Have We Learned and to What End?” in J.B. Taylor and M. Woodford, eds. *Handbook of Macroeconomics*, vol. 1A, Amsterdam: North Holland.

- Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (2001), “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *NBER Working Paper* 8403.
- Clarida, Richard, Jordi Galí and Mark Gertler (1999), “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature* 37(4): 1661-1707.
- Clarida, Richard, Jordi Galí and Mark Gertler (2000), “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics* 115(1): 147-180.
- Cooley, Thomas F. and Gary D. Hansen (1989), “The Inflation Tax in a Real Business Cycle Model,” *The American Economic Review* 79(4): 733-748.
- Cooley, Thomas F. and Gary D. Hansen (1995), “Money and the Business Cycle,” in T.F. Cooley, ed. *Frontiers of Business Cycle Research*, Princeton: Princeton University Press.
- Cooley, Thomas F. and Edward C. Prescott (1995), “Economic Growth and Business Cycles,” in T.F. Cooley, ed. *Frontiers of Business Cycle Research*, Princeton: Princeton University Press.
- DeLong, Bradford J. (1997), “America’s Peacetime Inflation: The 1970s,” in C. Romer and D. Romer, eds. *Reducing Inflation: Motivation and Strategy*, Chicago: University of Chicago Press.
- DeLong, Bradford J. (2000), “What Went Right in the 1990s? Source of American and Prospects for World Economic Growth,” in D. Gruen and S. Shresha, eds. *The Australian Economy in the 1990s*, Reserve Bank of Australia.
- Dennis, Richard and Ulf Söderström (2002), “How important is precommitment for monetary policy?” Working Paper No. 139, Sveriges Riksbank.
- Dixit, Avinash K and Joseph E. Stiglitz (1977), “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review* 67(3): 297-308.
- Erceg, Christopher J., Dale W. Henderson and Andrew T. Levin (2000), “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics*, 46(2): 281-313.
- Francis, Neville and Valerie A. Ramey (2002), “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *NBER Working Paper* 8726..
- French, Mark W. (2001), “Estimating Changes in Trend Growth of Total Factor Productivity: Kalman and H-P Filters versus a Markov-Switching Framework,” Board of Governors of the Federal Reserve System, Finance and Economics Discussion Paper Series: 2001-44.

- Galí, Jordi (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89(1): 249-71.
- Galí, Jordi (2001), "The Conduct of Monetary Policy in the Face of Technological Change: Theory and Postwar U.S. Evidence," in *Stabilization and Monetary Policy: the International Experience*, Banco de México.
- Galí, Jordi, J. David López-Salido and Javier Vallés (2002), "Technology Shocks and Monetary Policy: Assessing the Fed's Performance," forthcoming in the *Journal of Monetary Economics*.
- Giannoni, Marc and Michael Woodford (2002), "Optimal Interest-Rate Rules: I. General Theory," *mimeo*, Princeton University.
- Goodfriend, Marvin and Robert G. King (1997), "The New Neoclassical Synthesis and the Role of Monetary Policy," in *NBER Macroeconomics Annual* 12: 231-283.
- Gordon, Robert J. (2002), "Technology and Economic Performance in the American Economy," *NBER Working Paper* 8771.
- Greenspan, Alan (2000), "Challenges for Monetary Policymakers," Remarks at the 18th Annual Monetary Conference: *Monetary Policy in the New Economy*, Cato Institute, Washington, D.C., October 19.
- Grubb, Dennis, Richard Jackman and Richard Layard (1982), "Causes of the Current Stagflation," *Review of Economic Studies*, Special Issue 49(5): 707-730.
- Hamilton, James D. (1994), *Time Series Analysis*, Princeton: Princeton University Press.
- Ireland, Peter N. (1996), "The Role of Countercyclical Monetary Policy," *Journal of Political Economy* 104(4): 704-723.
- Jorgenson, Dale W. (2001), "Information Technology and the U.S. Economy," *American Economic Review* 91(1): 1-32.
- Kydland, Finn and Edward C. Prescott (1982), "Time to Build and Aggregate Fluctuations," *Econometrica* 50(6): 1345-70.
- King, Robert G. and Charles I. Plosser (1984), "Money, Credit, and Prices in a Real Business Cycle," *The American Economic Review* 74(3): 363-380.
- King, Robert G., Charles I. Plosser and Sergio Rebelo (1988), "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics* 21(2/3): 195-232.
- King, Robert G., Charles I. Plosser and Sergio Rebelo (2002), "Production, Growth and Business Cycles: Technical Appendix" forthcoming in *Computational Economics*.

- King, Robert G. and Sergio Rebelo (1999). "Resuscitating Real Business Cycles," in J.B. Taylor and M. Woodford, eds. *Handbook of Macroeconomics*, vol. 1B, Amsterdam: North Holland.
- King, Robert G. and Alexander L. Wolman (1996), "Inflation Targeting in a St. Louis Model of the 21st Century," *Federal Reserve Bank of St. Louis Review* 78: 83-107.
- Leeper, Eric M., Christopher Sims and Tao Zha (1996), "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity* 0(2): 1-63.
- Long, John B. and Charles I. Plosser (1983), "Real Business Cycles," *Journal of Political Economy* 91(1): 39-69.
- Lucas, Robert E. Jr., (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 31(1): 197-99.
- Lucas, Robert E. Jr., (1975), "An Equilibrium Model of the Business Cycle," *Journal of Political Economy* 83(6): 1113-44.
- Mankiw, Gregory N. and Ricardo Reis (2002), "Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps," in P. Aghion., R. Frydman, J.E. Stiglitz and M.D. Woodford, eds. *Knowledge, Information, and Expectations in Modern Macroeconomics: Essays in Honor of Edmund S. Phelps*, Princeton: Princeton University Press.
- Meyer, Lawrence H. (2001), "What Happened to the New Economy?" Remarks Before the New York Association for Business Economics and The Downtown Economists, New York, June 6.
- Oliner, Stephen D. and Daniel E. Sichel (2000), "The Resurgence of Growth in the Late 1990's: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(1): 3-22.
- Plosser, Charles I. (1989), "Understanding Real Business Cycles," *Journal of Economic Perspectives*, 3(1): 51-77.
- Prescott, Edward (1986), "Theory Ahead of Business Cycle Measurement," *Carnegie-Rochester Conference Series on Public Policy* 25: 11-44.
- Primiceri, Giorgio (2002), "Time varying structural vector autoregressions and monetary policy," *mimeo*, Princeton University.
- Roberts, John M. (2001), "Estimates of the Productivity Trend Using Time-Varying Parameter Techniques," *Contributions to Macroeconomics* 1: <http://www.bepress.com/bejm/contributions>.
- Rotemberg, Julio J. and Michael D. Woodford (1997), "An Optimization-Based Econometric Framework for the evaluation of Monetary Policy," *NBER Macroeconomics Annual* 12: 297-346.

- Rotemberg, Julio J. and Michael D. Woodford (1998), "An Optimization-Based Econometric Framework for the evaluation of Monetary Policy: Expanded Version," *NBER Technical Working Paper 233*.
- Rotemberg, Julio J. and Michael D. Woodford (1999), "Interest-Rate Rules in an Estimated Sticky-Price Model," in J.B. Taylor, ed. *Monetary Policy Rules*, Chicago: University of Chicago Press.
- Sargent, Thomas J. (1999), *The Conquest of American Inflation*, Princeton: Princeton University Press.
- Staiger, Douglas, James H. Stock and Mark Watson (2001), "Prices, Wages and the U.S. NAIRU in the 1990s," *mimeo*, Princeton University.
- Svensson, Lars E.O. and Michael D. Woodford (2002) "Indicator Variables for Optimal Policy under Asymmetric Information," *Journal of Economic Dynamics and Control*, forthcoming.
- Smets, Frank and Raf Wouters (2002), "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *mimeo*, European Central Bank.
- Stock, James H. and Mark. W. Watson (2002), "Has the Business Cycle Changed and Why?," *mimeo*, Princeton University.
- Taylor, John B. (1993), "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214.
- Taylor, John B. (1999), *Monetary Policy Rules*, Chicago: Chicago University Press.
- Woodford, Michael D. (1996), "Control of the Public Debt: A Requirement for Price Stability?" *NBER Working Paper 5684*.
- Woodford, Michael D. (2002), *Interest and Prices*, Book Manuscript, Princeton University.
- Woodford, Michael D. (2002a), "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion., R. Frydman, J.E. Stiglitz and M.D. Woodford, eds. *Knowledge, Information, and Expectations in Modern Macroeconomics: Essays in Honor of Edmund S. Phelps*, Princeton: Princeton University Press.
- Woodward, Bob (2000), *Maestro: Greenspan's Fed and the American Boom*, New York: Simon & Schuster.
- Yun, Tack (1996), "Monetary Policy, Nominal Price Rigidity, and Business Cycles," *Journal of Monetary Economics* 37(2): 345-70.

Figures and Tables

Productivity	
ρ	.93
σ_a	.687%
σ_γ	.096%

Table 1: Calibrated parameter values: productivity process.

Tastes		Technology	
β	.99	ϕ^{-1}	.75
ν	.2	ω_p	.33
$(\theta_p - 1)^{-1}$.19	$(\theta_w - 1)^{-1}$.13

Table 2: Calibrated parameter values: tastes and technology.

Stickiness	CE	AL	Indexation	ALf	ALp
α_p	.42	.66	λ_γ	1	0
α_w	.78	.66	λ_p	0	0
$\tilde{\psi}_p$.56	.56	λ_w	0	0
$\tilde{\psi}_w$.56	.56			

Table 3: Calibrated parameter values. **Stickiness**: CE refers to the benchmark calibration; AL refers to the estimates of Amato and Laubach (2002). **Indexation**: ALf represents indexation of wages to the forecast of productivity; ALp represents indexation of wages to past productivity.

Phillips Curves Parameters		
$\omega_w = \nu\phi$	$\omega = \omega_w + \omega_p$	$\zeta_w = \frac{\alpha_w}{1-\tilde{\psi}_w(1-\alpha_w)}$
$\psi_p = \frac{\alpha_p\tilde{\psi}_p}{1-\tilde{\psi}_p(1-\alpha_p)}$	$\xi_p = \frac{1-\alpha_p}{\alpha_p} \frac{1-\alpha_p\beta}{1+\omega_p\theta_p}$	$\kappa_p = \xi_p\omega_p$
$\psi_w = \frac{\alpha_w\tilde{\psi}_w}{1-\tilde{\psi}_w(1-\alpha_w)}$	$\xi_w = \frac{1-\alpha_w}{\alpha_w} \frac{1-\alpha_w\beta}{1+\nu\theta_w}$	$\kappa_w = \xi_w(1+\omega_w)$

Table 4: Parameters of the wage and price Phillips curves, expressed as functions of the structural parameters.

Loss Function	CE	AL
Λ_p	.11	.52
Λ_w	.89	.48
Λ_x	.0026	.0035

Table 5: Calibrated parameter values: loss function. CE and AL denote the values of the welfare weights derived under the benchmark calibration and the estimates of Amato and Laubach (2002) respectively.

Interest Rate Rule	Volcker-Greenspan	Pre-Volcker
ϕ_i	.79	.68
ϕ_π	2.15	1
ϕ_x	.93	.27

Table 6: Calibrated parameter values: interest rate rule. The Volcker-Greenspan column refers to estimates in Galí *et al.* (2000) for the period 1979:III to 1996:IV, while the Pre-Volcker column refers to estimates for the period 1960:I to 1979:II.

Empirical Impulse Responses

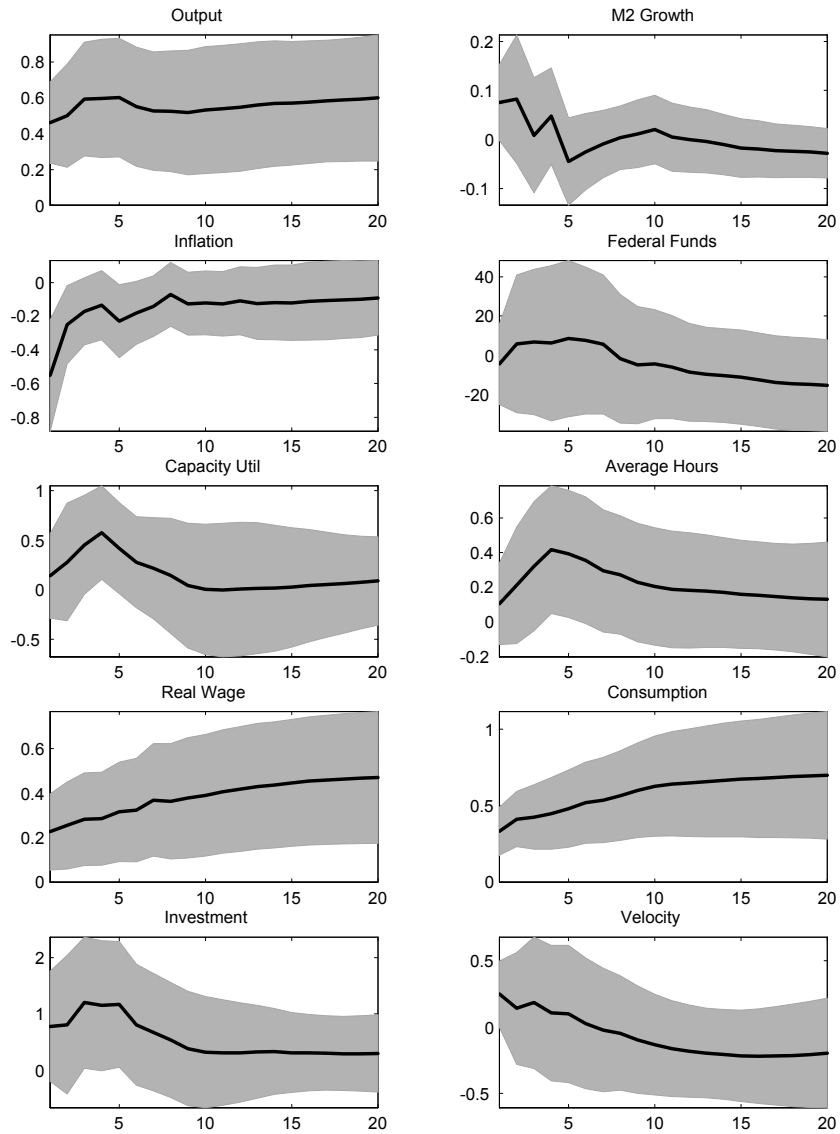


Figure 1: Impulse responses to an innovation in technology (from Altig *et al.*, 2002). The shaded area denotes 95% confidence bands.

Growth Rate of Productivity, Actual and Forecasted




Figure 2: Impulse responses of the actual (γ_t^a) and forecasted ($\gamma_{t|t-1}^a$) growth rate of productivity to a 1% persistent technology shock (ε_t^γ), under full and partial information. Productivity growth is expressed as an annualized percentage rate.

Benchmark Model

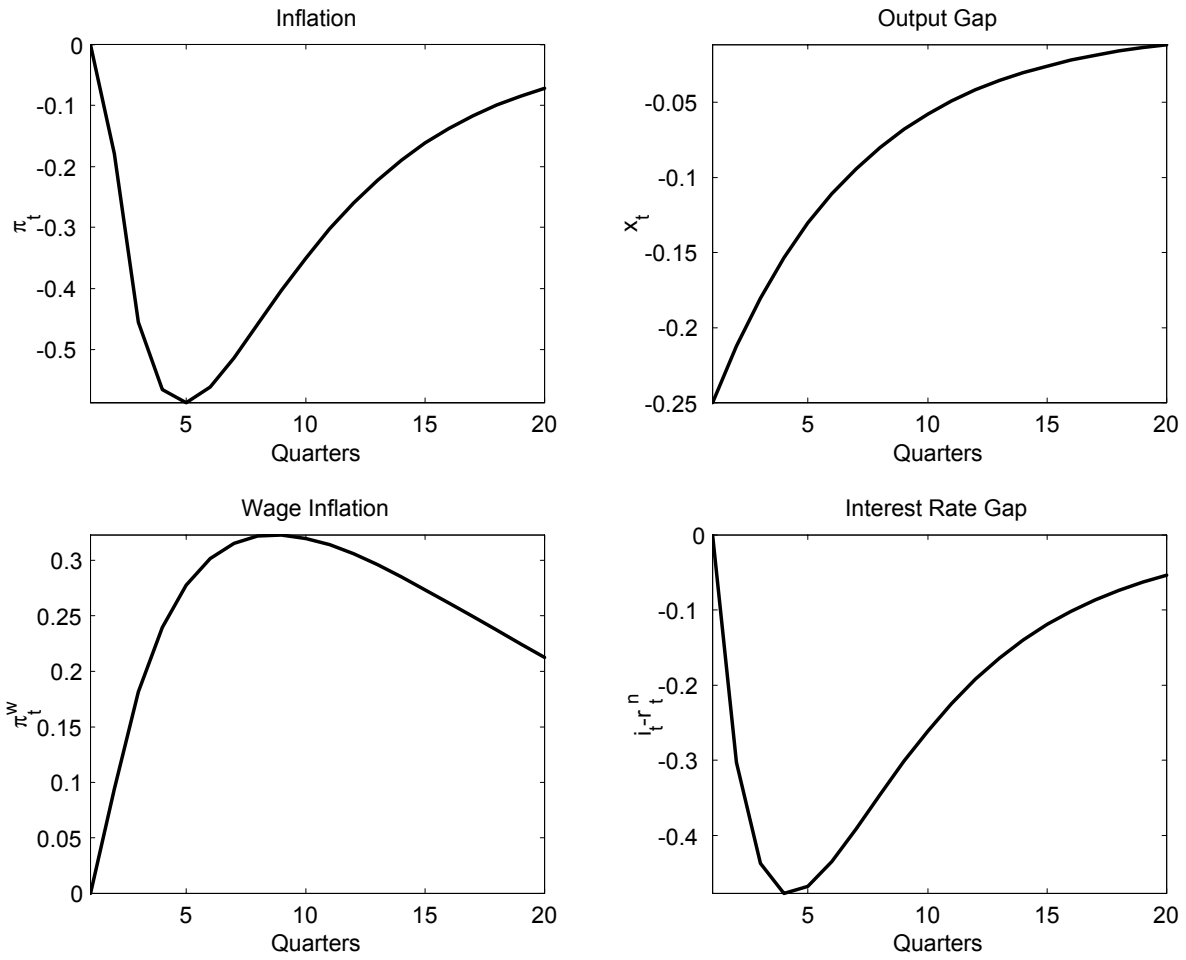


Figure 3: Simulated impulse responses to a 1% ε_t^γ shock for the benchmark model; monetary policy is assumed to target the one period ahead forecast of the output gap. Inflation and the interest rate gap are expressed as annualized percentage rates, while the output gap is expressed in percentage points per quarter.

Model Comparison: Informational Frictions

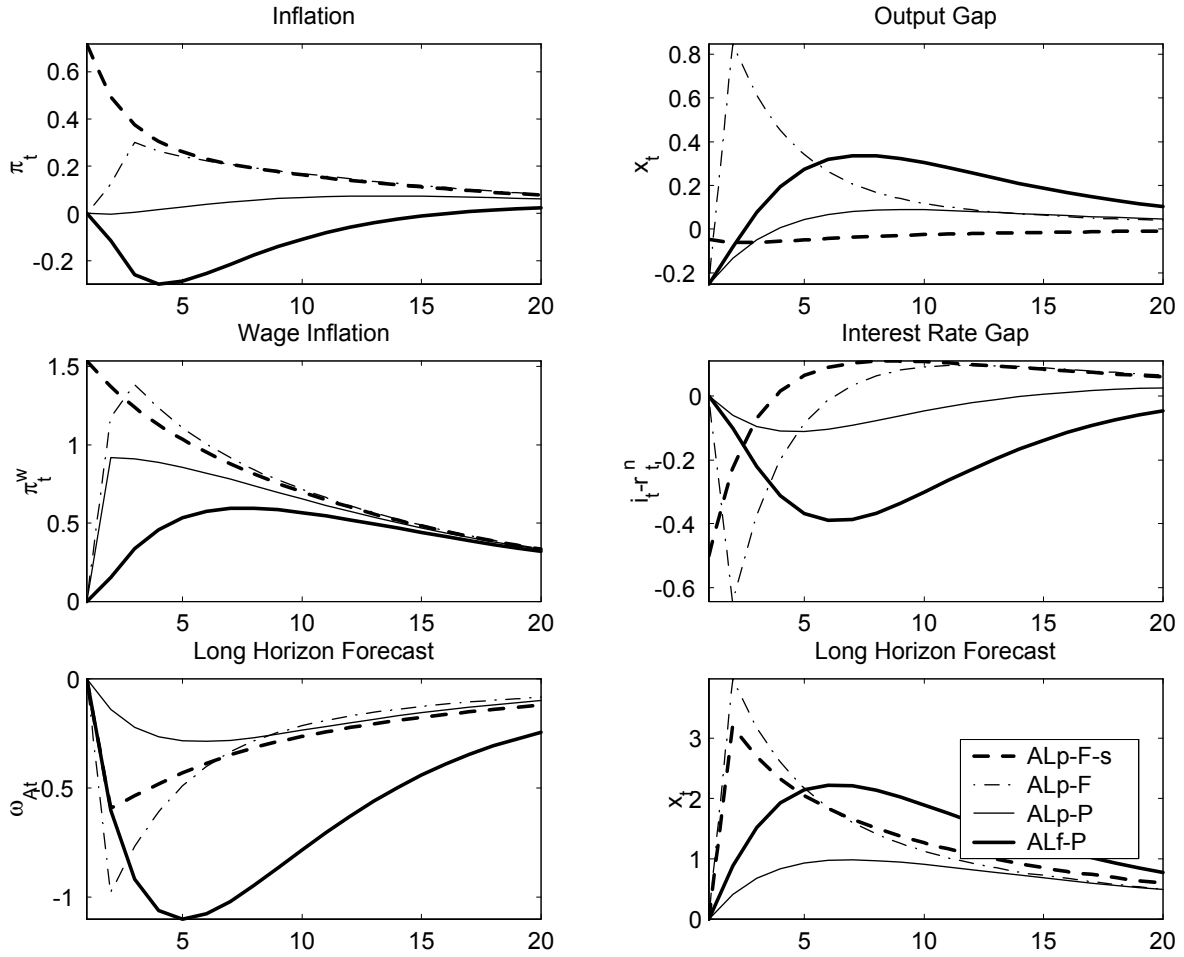


Figure 4: Simulated impulse responses to a 1% ε_t^γ shock under the Volcker-Greenspan interest rate rule (table 6). ALp and ALf refer to the type of wage indexation (table 3); F and P refer to Full and Partial information respectively; s refers to the simple model with no information delays. The thick line (ALf-P) is the benchmark model. Inflation and the interest rate gap are expressed as annualized percentage rates, while the level variables are expressed as percentage deviations from their steady state value; the long horizon forecasts are expressed in percentage points per quarter.

Empirical Comparison

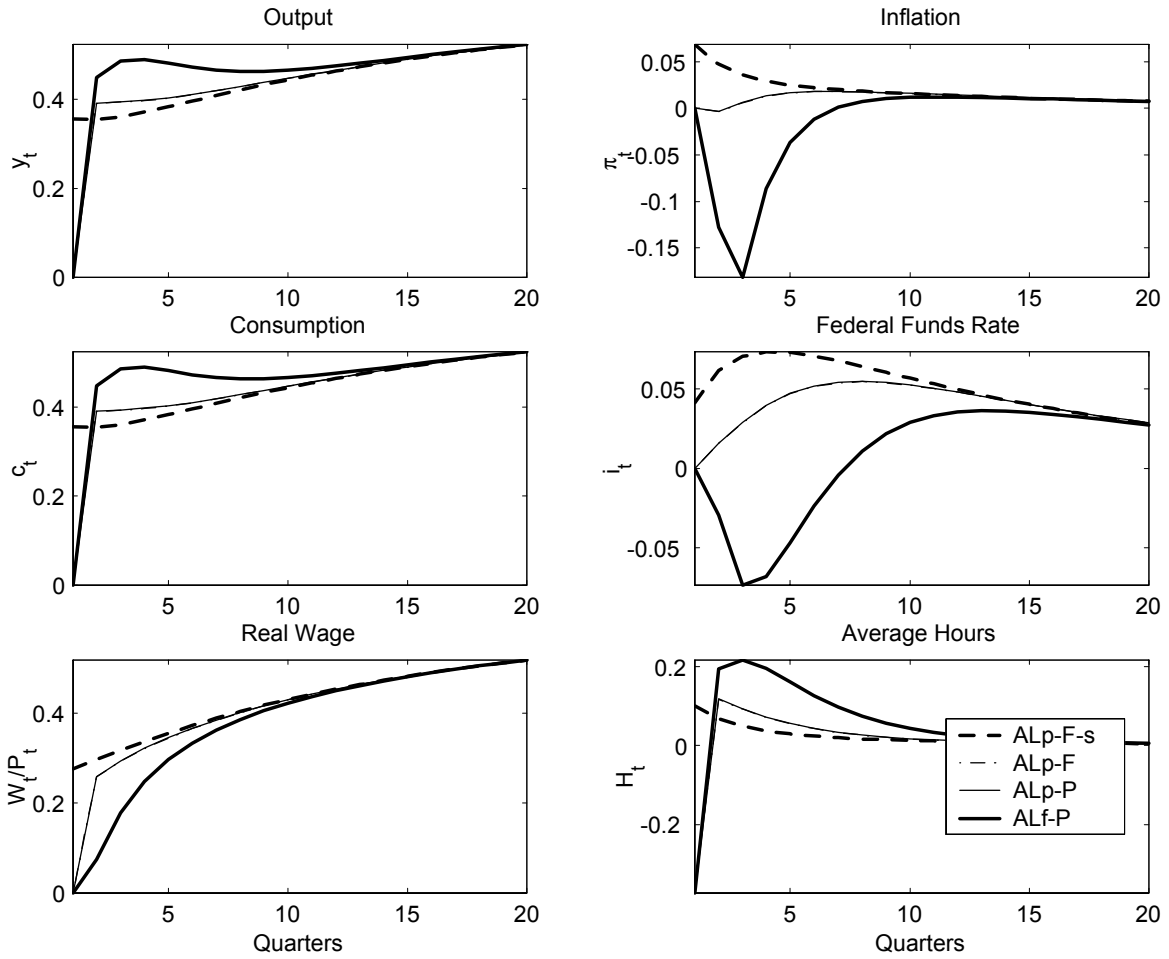


Figure 5: Simulated impulse responses to the unconditional shock (ε_t) under the Volcker-Greenspan interest rate rule (table 6). ALp and ALf refer to the type of wage indexation (table 3); F and P refer to Full and Partial information respectively; s refers to the simple model with no information delays. The thick line (ALf-P) is the benchmark model. Inflation and the interest rate are expressed as annualized percentage rates, while the level variables are expressed as percentage deviations from their steady state value.

Case Study: The New Economy

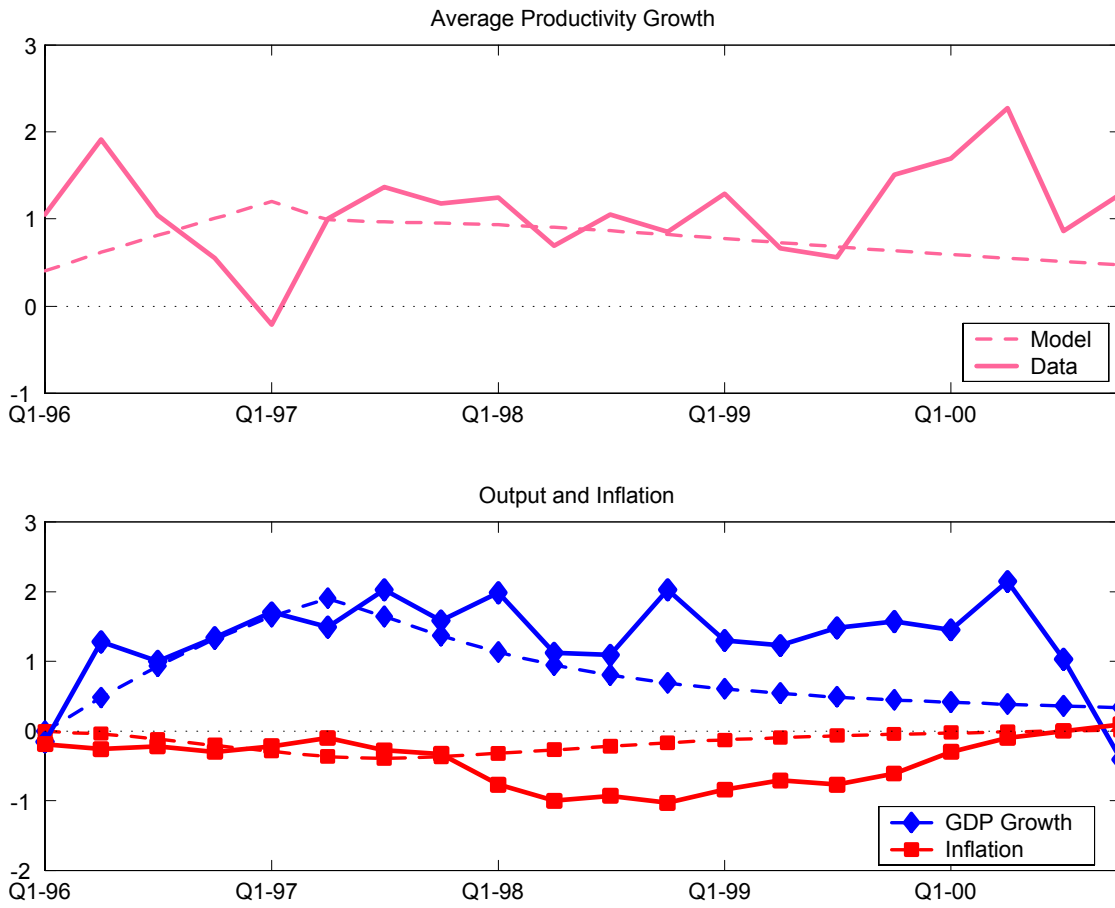


Figure 6: Simulated (dotted line) and actual (solid line) evolution of productivity growth, inflation and GDP growth for the period 1996:I to 2000:IV, under the Volcker-Greenspan interest rate rule. The model is hit in 1996:I with a productivity shock that decays like an AR(1) process with the benchmark parameter 0.8. All variables are expressed as growth rates over the previous year, in deviation from their average over the year 1995.

Case Study: The Productivity Slowdown




Figure 7: Simulated (dotted line) and actual (solid line) evolution of productivity growth, inflation and GDP growth for the period 1973:III to 1978:II, under the Volcker-Greenspan interest rate rule. The model is hit in 1973:III with a productivity shock that decays like an AR(1) process with the benchmark parameter 0.8. All variables are expressed as growth rates over the previous year, in deviation from their average over the entire business cycle 1971:I-1973:II.

Optimal Policy: Unconditional Shock

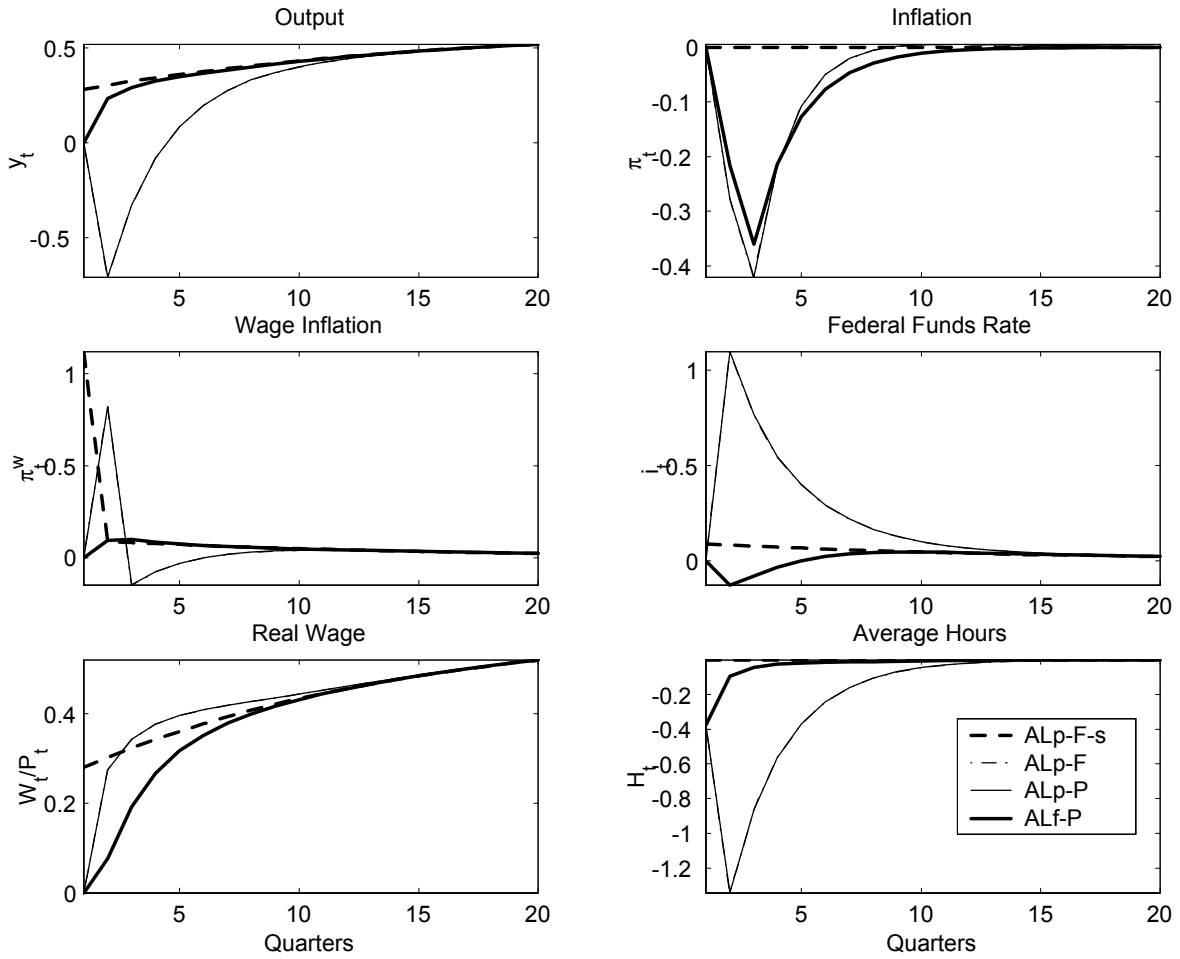


Figure 8: Simulated impulse responses to the unconditional shock (ε_t) under the Volcker-Greenspan interest rate rule (table 6). ALp and ALf refer to the type of wage indexation (table 3); F and P refer to Full and Partial information respectively; s refers to the simple model with no information delays. The thick line (ALf-P) is the benchmark model. Inflation and the interest rate are expressed as annualized percentage rates, while the level variables are expressed as percentage deviations from their steady state value.

Optimal Policy: Trend Shock

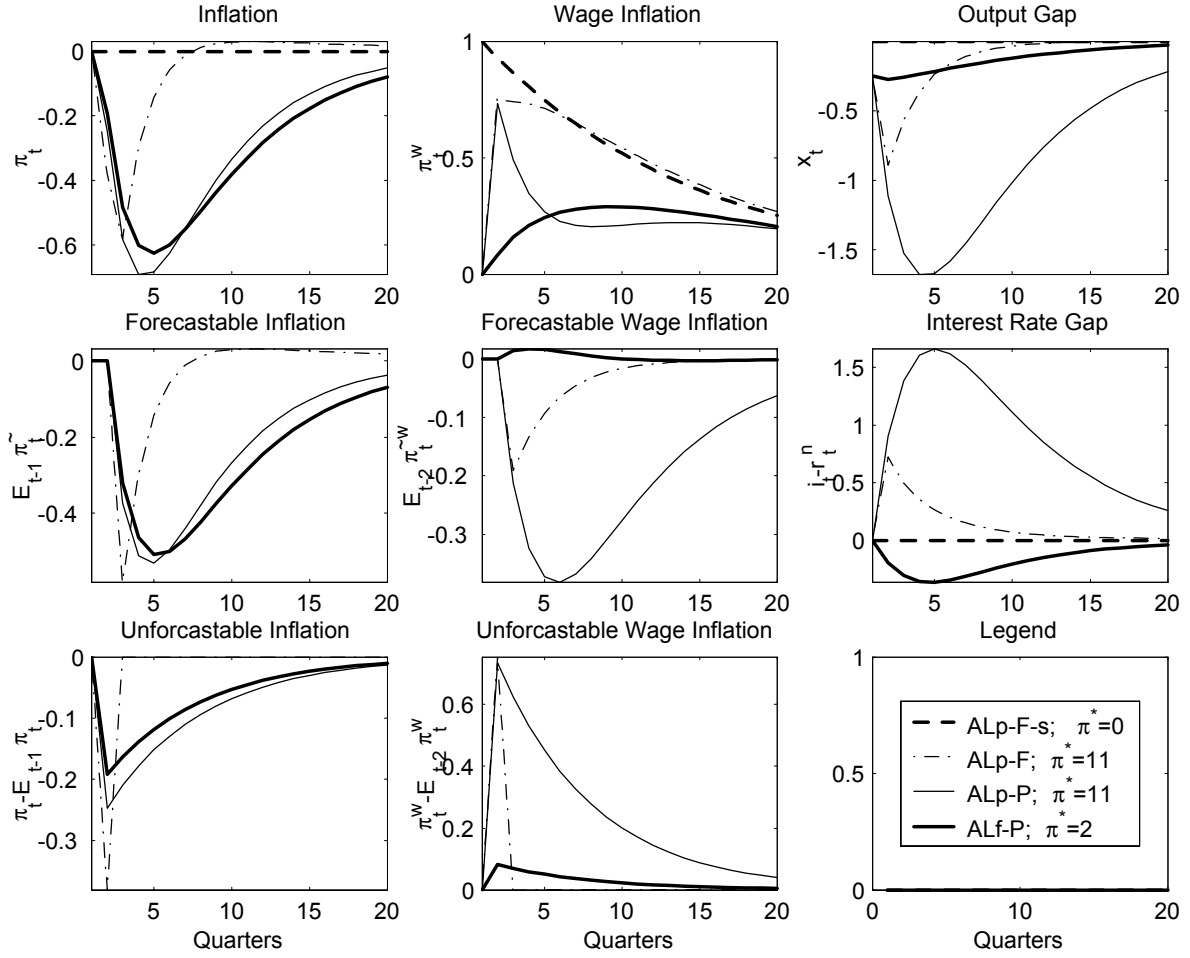


Figure 9: Simulated impulse responses to a 1% ε_t^γ shock under the optimal policy. ALp and ALf refer to the type of wage indexation (table 3); F and P refer to Full and Partial information respectively; s refers to the simple model with no information delays. The thick line (ALf-P) is the benchmark model. π^* is the “inflation equivalent”. Inflations and the interest rate gap are expressed as annualized percentage rates, while the output gap is expressed in percentage points per quarter.

Policy Comparison: Targeting Rules

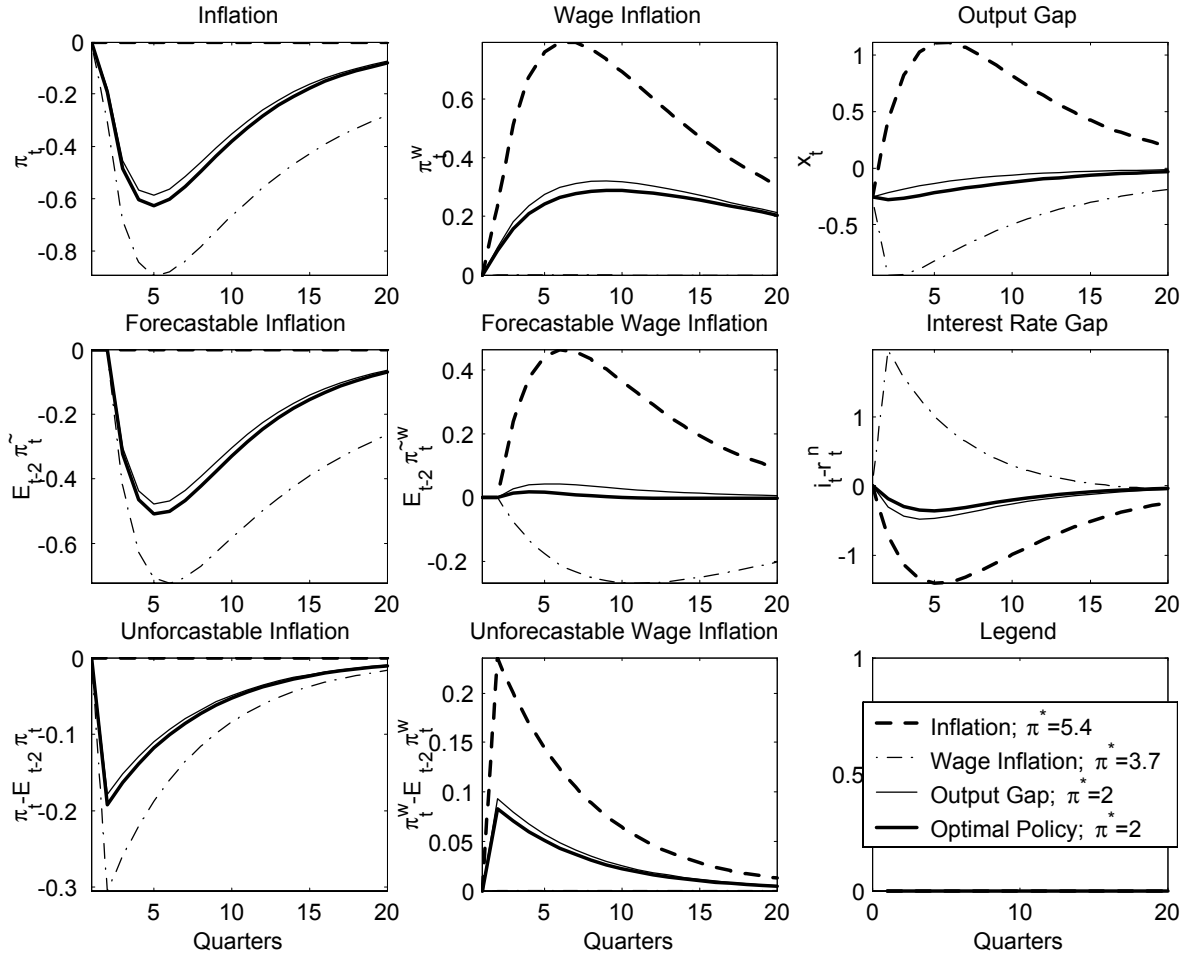


Figure 10: Simulated impulse responses to a 1% ε_t^γ shock under the optimal policy; benchmark calibration. The different line styles refer to inflation targeting, nominal wage inflation targeting, output gap forecast targeting and the optimal policy, as indicated by the legend. π^* is the “inflation equivalent”. Inflation and the interest rate gap are expressed as annualized percentage rates, while the output gap is expressed in percentage points per quarter.

Policy Comparison: Relative Stickiness

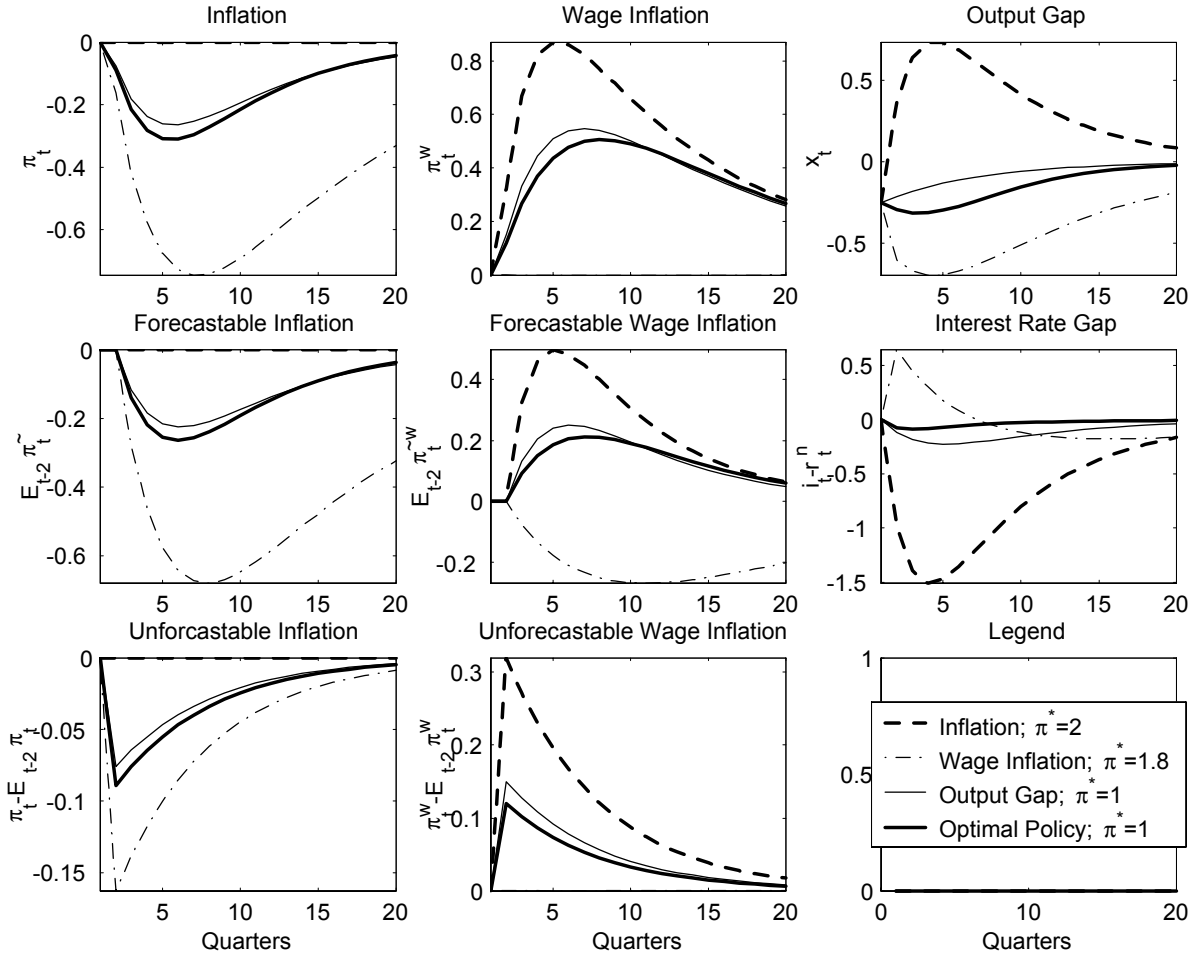


Figure 11: Simulated impulse responses to a 1% ε_t^γ shock under the optimal policy, with $\alpha_w = \alpha_p = 0.66$. The different line styles refer to inflation targeting, nominal wage inflation targeting, output gap forecast targeting and the optimal policy, as indicated by the legend. π^* is the “inflation equivalent”. Inflation and the interest rate gap are expressed as annualized percentage rates, while the output gap is expressed in percentage points per quarter.

Optimal Policy: The New Economy

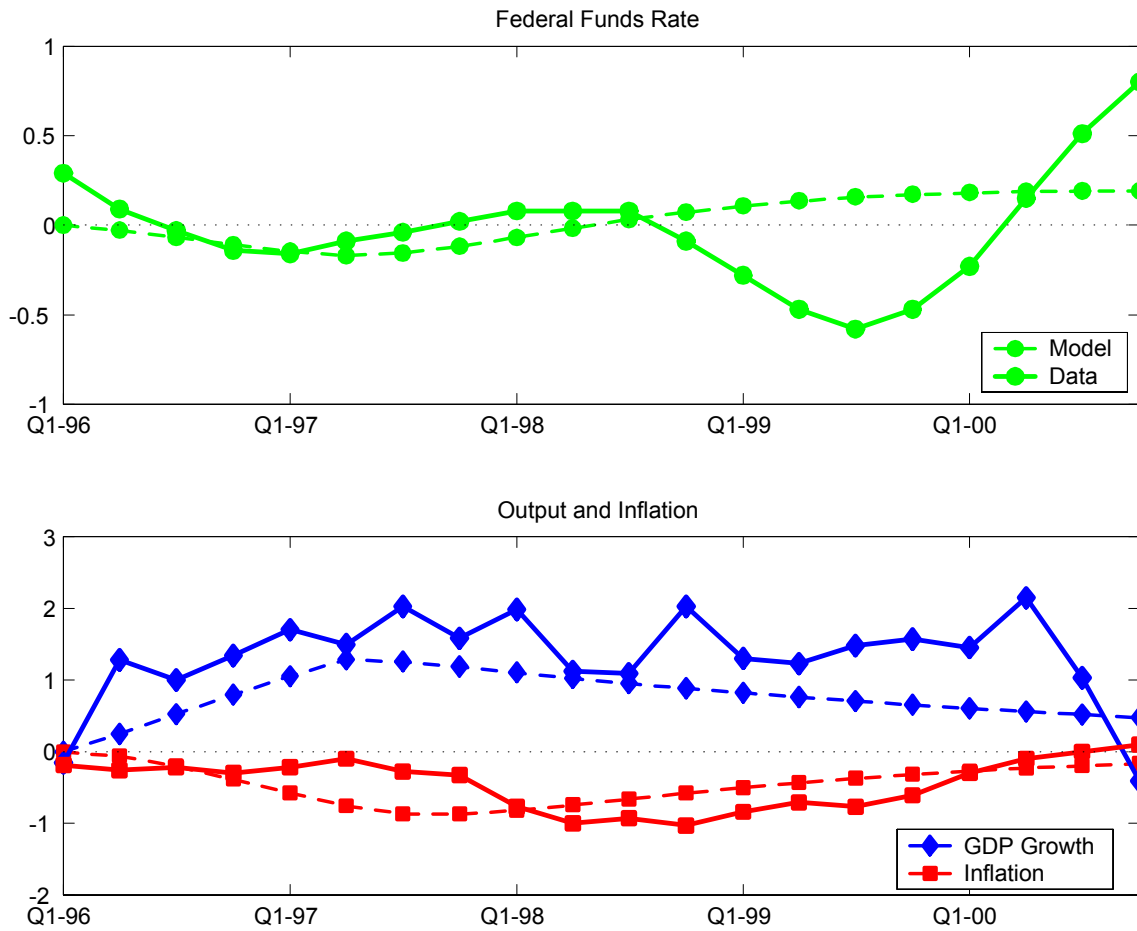


Figure 12: Simulated (dotted line) and actual (solid line) evolution of Federal Funds rate, inflation and GDP growth for the period 1996:I to 2000:IV, under the optimal policy. The model is hit in 1996:I with a productivity shock that decays like an AR(1) process with parameter 0.9. Inflation and GDP growth are expressed as percentage growth rates over the previous year, while the Federal Funds rate is expressed in annualized percentage points per quarter. All variables are in deviation from their average over the year 1995.

The Pre-Volcker Interest Rate Rule




Figure 13: Simulated (dotted line) and actual (solid line) evolution of productivity growth, inflation and GDP growth for the period 1973:III to 1978:II, under the pre-Volcker interest rate rule. The model is hit in 1973:III with a productivity shock that decays like an AR(1) process with the benchmark parameter 0.8. All variables are expressed as growth rates over the previous year, in deviation from their average over the entire business cycle 1971:I-1973:II.

Optimal Policy: The Productivity Slowdown

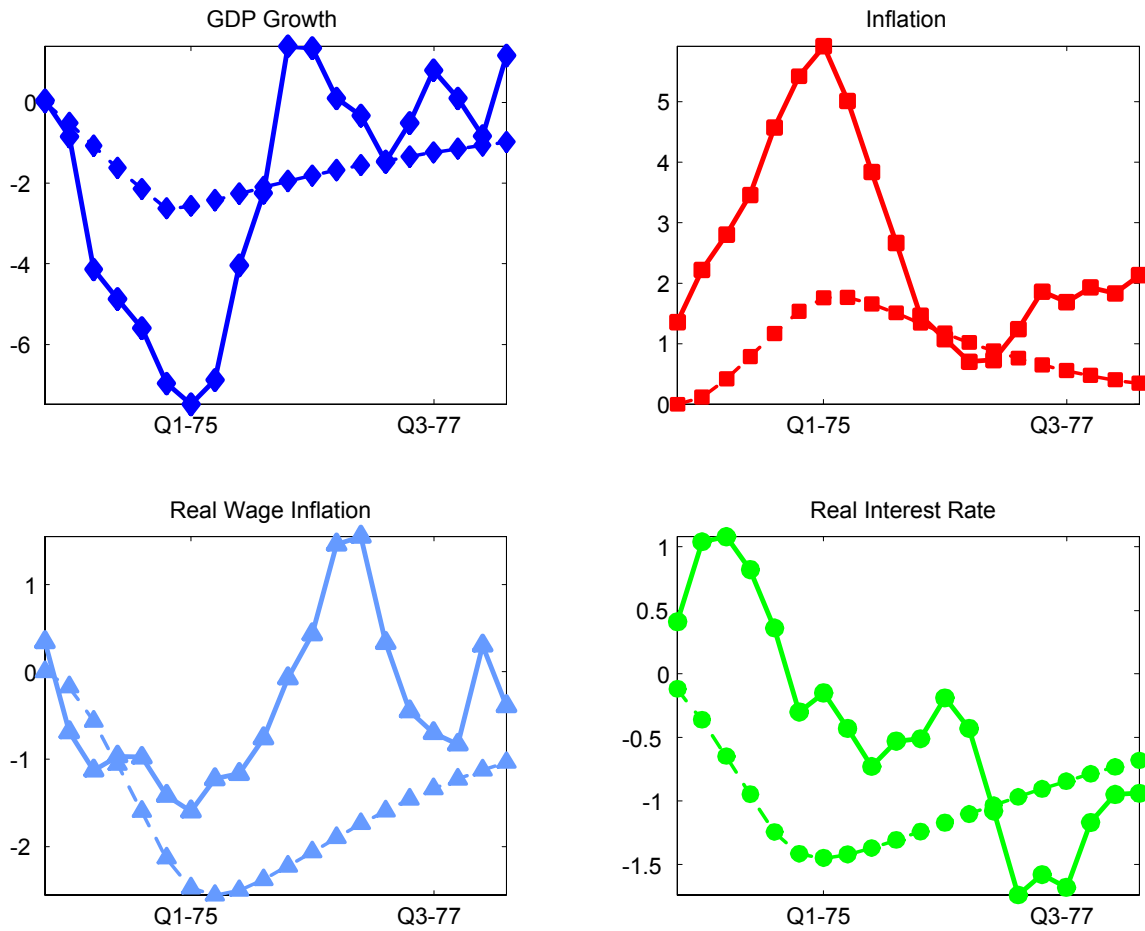


Figure 14: Simulated (dotted line) and actual (solid line) evolution of GDP growth, inflation, real wage inflation and the *ex-post* real interest rate for the period 1973:III to 1978:II, under the optimal policy. The model is hit in 1973:III with a productivity shock that decays like an AR(1) process with the benchmark parameter 0.8. All variables are expressed as growth rates over the previous year, with the exception of the interest rate, which is a four quarters moving average of annualized quarterly rates.