



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Monetary Economics 52 (2005) 703–725

Journal of
MONETARY
ECONOMICS

www.elsevier.com/locate/jme

Monetary policy for inattentive economies[☆]

Laurence Ball^a, N. Gregory Mankiw^b, Ricardo Reis^{c,*}

^a*Department of Economics, Johns Hopkins University, Baltimore, MD 21218, USA*

^b*Department of Economics, Harvard University, Cambridge, MA 02138, USA*

^c*Department of Economics and Woodrow Wilson School, Princeton University, Princeton, NJ 08544, USA*

Received 11 October 2002; received in revised form 1 March 2005; accepted 3 March 2005

Abstract

We offer a contribution to the analysis of optimal monetary policy. We begin with a critical assessment of the existing literature, arguing that most work is based on implausible models of inflation–output dynamics. We then suggest that this problem may be solved with some recent behavioral models, which assume that price setters are slow to incorporate macroeconomic information into the prices they set. A specific such model is developed and used to derive optimal policy. In response to shocks to productivity and aggregate demand, optimal policy is price level targeting. Base drift in the price level, which is implicit in the inflation targeting regimes currently used in many central banks, is not desirable in this model. When shocks to desired markups are added, optimal policy is flexible targeting of the price level. That is, the central bank should allow the price level to deviate from its target for a while in response to these supply shocks, but it should eventually return the price level to its target path. Optimal policy can also be described as an elastic price standard: the central bank allows the price level to deviate from its target when output is expected to deviate from its natural rate.

© 2005 Elsevier B.V. All rights reserved.

JEL classification: E50; E52; E30

Keywords: Monetary policy; Price-level targeting; Sticky information; Phillips curve

[☆]We are grateful to Jeffrey Fuhrer, Charles Goodhart and Yves Nosbusch for helpful comments. Reis is grateful to the Fundacao Ciencia e Tecnologia, Praxis XXI for financial support.

*Corresponding author.

E-mail address: rreis@princeton.edu (R. Reis).

1. Introduction

What policy rule should a central bank follow? Recent years have seen a resurgence of theoretical research on this classic question. Most of this work has built on “new Keynesian” models of the output–inflation trade-off derived from forward-looking models of staggered price adjustment. Unfortunately, these models make implausible predictions about the effects of monetary policy: for example, they imply that a policy change that gradually reduces inflation causes an output boom. There is therefore good reason to be skeptical about what the literature tells us about the effects of alternative policies.

This article tries to make progress toward determining which policies are optimal by studying this question in a model that more closely fits the facts about monetary policy. In particular, we draw on recent behavioral models of the output–inflation trade-off based on the assumption that agents are slow to incorporate information about macroeconomic conditions, even if the information is publicly available. Recent work has shown that such models capture the inertia that is central to inflation dynamics in modern economies. These models should provide more reliable insights into the policy choices facing central banks.

Section 2 briefly reviews the two literatures on which our work builds, the work on optimal monetary policy and the work on behavioral macroeconomics. Section 3 presents a specific model, which builds on the “sticky information” model of [Mankiw and Reis \(2002\)](#), and Section 4 discusses the determinants of welfare in this model. Sections 5–7 derive the optimal policy rules in the model. Section 8 compares our results with those obtained from the standard new Keynesian Phillips curve. Section 9 concludes.

Our central result is that price level targeting is the optimal policy in the model. Inflation targeting—the currently popular policy of allowing base drift in the price level—is suboptimal. When the economy is hit by shocks to aggregate demand or productivity, strict price level targeting is optimal: policymakers should return the price level to a pre-determined path as quickly as possible. However, if there are persistent shocks to firms’ markups, the optimal rule allows temporary deviations from the long-run price target. In this way, the prescriptions of our model are similar to the practice of many central banks, which allow temporary deviations from policy rules in response to “supply shocks.” One can also describe optimal policy as the elastic price standard proposed by [Hall \(1984\)](#). Under this policy, the price level can deviate from target as long as output is expected to deviate from its natural rate.

2. Motivation

This research arises from two recent literatures—one on optimal rules for monetary policy and one on behavioral approaches to the Phillips curve. We believe that the second may hold the hope of remedying some deficiencies in the first. The natural place to start our analysis is with a brief overview of these two broad literatures.

2.1. *The sorry state of monetary policy analysis*

In recent research on policy rules, a canonical approach has emerged. The researcher sets up a model based on monopolistically competitive producers, with monetary nonneutrality arising because of some friction in price setting. Policymakers are assumed to choose policy to minimize a loss function, which is either assumed or derived as an approximation to agents' utility functions. With this framework in hand, the researcher can derive the optimal policy rule and evaluate policy proposals, such as targeting inflation or the price level. [Goodfriend and King \(1997\)](#) have referred to this literature as the New Neoclassical Synthesis, while [Clarida et al. \(1999\)](#) have called it the New Keynesian Science of Monetary Policy. Whatever the label, there is no doubt that this approach to monetary-policy analysis is large and growing.

Yet policymakers should be wary of the prescriptions that this literature has yielded. The results in any such analysis depend crucially on the assumed model of the Phillips curve, which determines how inflation is related to output movements and expectations. The results should be believed only if the assumed Phillips curve is credible. Unfortunately, that is rarely the case.

The most common approach is to use the new Keynesian Phillips curve. This model is based on the assumption of time-contingent price adjustment and is derived from the seminal articles of [Taylor \(1980\)](#), [Rotemberg \(1982\)](#), and [Calvo \(1983\)](#). The model's appealing microeconomic foundations have made it the workhorse of much of the modern literature on monetary policy. The articles by [Goodfriend and King \(1997\)](#) and [Clarida et al. \(1999\)](#) are two examples.

There is, however, a problem with this approach: Because the new Keynesian Phillips curve lacks any source of inflation inertia, it makes absurdly counterfactual predictions about the effects of monetary policy. According to this model, disinflations can result in booms rather than recessions ([Ball, 1994](#)). Similarly, the model fails to explain why monetary policy shocks have a delayed and gradual effect on inflation ([Mankiw, 2001](#)). One should be suspect of policy recommendations arising from any model that is patently inconsistent with both the econometric evidence and the views of central bankers about their influence over the economy.

An alternative approach to using the new Keynesian Phillips curve is to use a more traditional accelerationist Phillips curve. This approach, which can be justified with the assumption of backward-looking expectations, has the advantage that it is consistent with the conventional wisdom about the effects of monetary policy and with standard empirical analyses of inflation (e.g., [Gordon, 1997](#); [Staiger et al., 1997](#)). For these reasons, [Ball \(1999\)](#) and [Svensson \(1997\)](#) take an accelerationist Phillips curve as a starting point for their analyses.

Yet this approach is also problematic. Although inflation inertia is a feature of the U.S. monetary regime of the past several decades, this was not true when the economy operated under a gold standard ([Barsky, 1987](#); [Alogoskoufis and Smith, 1991](#)). When analyzing alternative policy rules, it is imperative that expectations be allowed to adjust to the new regime. Regardless of how well it fits the recent data, the accelerationist Phillips curve is a suspect tool for monetary-policy analysis. When

analyzing alternative monetary policy rules, ignoring the Lucas critique is not just an aesthetic faux pas; it produces results that are, literally, incredible.

The problems with both the new Keynesian and the accelerationist Phillips curves are increasingly understood. This has led some authors to suggest a compromise. According to the so-called “hybrid” Phillips curve, a subset of agents have backward-looking inflation expectations and the rest have rational expectations (e.g., Gali and Gertler, 1999). Yet this compromise may yield the worst of both worlds. Like the new Keynesian Phillips curve, the hybrid model yields an immediate jump in inflation in response to monetary policy shocks (unless all agents are backward looking). Like the accelerationist Phillips curve, the hybrid model fails to explain the absence of inflation inertia under earlier monetary regimes (unless no agents are backward looking). That is, by taking a weighted average of two flawed models, the hybrid model of the Phillips curve ends up with the flaws of each.

2.2. *The promise of behavioral models*

Behavioral economics might offer a way out of this conundrum. Behavioral economics is, broadly defined, the growing subfield that incorporates into economic theory the flaws in human decision-making that are ignored in the standard model of rational man. It finds its roots in Herbert Simon’s suggestion that people are “satisficers” rather than rational maximizers. Many economists have suggested that departures from rationality may be important for issues in macroeconomics (Akerlof, 2002; Sargent, 1993). In particular, several recent articles on the micro-foundations of the Phillips curve have proposed that monetary nonneutrality arises because people are slow to process widely available macroeconomic information.

When Milton Friedman (1968) introduced the natural rate hypothesis in his AEA presidential address, he proposed a theory of monetary nonneutrality that, to modern ears, sounds remarkably behavioral. According to Friedman, when the money supply rises unexpectedly, the price level rises, pushing down the real wage. Employers hire more because the cost of labor has fallen. Employees are willing to work more because they focus on the nominal wage and infer (incorrectly) that the reward for working has risen. Friedman did not explain this asymmetry in information between firms and workers. For him, it seemed natural to assume that workers lacked full information about macroeconomic conditions. This ignorance was the centerpiece of Friedman’s proposed explanation for the short-run Phillips curve.

Work that followed on Friedman’s ideas tried to incorporate rational economic man into the story. The task has proved difficult. Lucas (1973) suggested a model in which producers observe the prices of what they sell but not of what they buy and, as a result, need to solve signal-extraction problems to infer relative prices. Critics of the Lucas model wondered why a rational man would fail to use the widely available information on the consumer price index, which is published monthly. Fischer (1977) told a version of Friedman’s story in which workers were locked into long-term contracts setting the nominal wage. Critics of the Fischer model wondered why a rational man would ever agree to such an inefficient contract *ex ante* or fail to renegotiate it *ex post*.

In recent years, there have been several attempts to come up with better models of inflation dynamics by returning to Friedman's idea that some people fail to incorporate all available macroeconomic information into their decision-making. According to the new Keynesian Phillips curve, prices are sticky, but inflation can respond instantly to changes in monetary policy. In practice, inflation responds sluggishly. This fact suggests that, for some reason, price setters are not quite awake or smart or informed enough to process all available information about monetary policy and react immediately to it.

One approach to modeling this inattentiveness is to use the tools of information theory, as explicated, for instance, in the textbook by Cover and Thomas (1991). Drawing on these tools, Sims (2003) suggests modeling humans as having a limited channel for absorbing information. That is, the human brain is imperfect in the same way as a computer with a slow internet connection would be. Woodford (2003a) uses this idea to build a model of inflation–output dynamics. In his model, because price setters learn about monetary policy through a limited-information channel, it is as if they observe monetary policy with a random error and have to solve a signal-extraction problem along the lines of Lucas (1973).

Ball (2000) proposes another approach to the problem. He suggests that when forming expectations of any variable, people optimally use all information in the past values of that variable, but fail to incorporate information from other variables. That is, expectations are based on optimal univariate forecasts. Ball shows that this approach can explain why the accelerationist Phillips curve fits the recent data well, while the earlier data conform more closely to a classic Phillips curve.

Mankiw and Reis (2002), like Ball and Woodford, impose a constraint on the information that people use when forming expectations. They assume that in each period there is a fixed probability that a person updates his information set; otherwise, he continues to set prices based on outdated information. Mankiw and Reis compare this sticky-information model with the standard new Keynesian Phillips curve and conclude that it is more consistent with standard views about the effects of monetary policy. The Mankiw–Reis model has strong formal ties to the earlier efforts of Lucas and Fischer, but its bald informational assumption rejects the attempt to reconcile the Phillips curve with the axiom of rationality.

A weakness of all three of these approaches is that they fail to explain why people do not incorporate widely available news about monetary policy into their plans. In reality, it is easy to find out what the central bank is doing, but it is often hard to figure out what it means. That is, the real problem is not obtaining information but processing it. Unfortunately, economists do not have the tools to model imperfect information processing. The approaches of Woodford, Ball, and Mankiw and Reis are all based on the hope that a model of imperfect information acquisition may serve as a rough substitute.

A strength of all three approaches is that they can explain inflation inertia. In these models, when monetary policy changes, most price setters are unaware of it, so they keep marking up prices as if no policy change had occurred. That is, all three models generate inflation inertia by positing a type of inattentiveness on the part of price setters. Because these models of inattention fit the most basic facts about inflation

dynamics, they hold out the greatest promise for delivering credible advice for monetary policy.

The specific model we use here is the sticky-information model proposed by Mankiw and Reis. One reason is tractability. A nice feature of the Mankiw–Reis model is that, given the dates when people get their information, they have conventional rational expectations. This feature allows the use of rational-expectations tools to solve for the resulting equilibrium. That is, we can use many standard and powerful modeling techniques, while making the behavioral assumption that price setters are inattentive.

3. The model

We begin our analysis by presenting the model, including agents' objectives, market structure, price setting, and the role of aggregate-demand policy.

3.1. Agents and market structure

We assume a simple variation on the standard “yeoman farmer economy” studied by authors such as Rotemberg and Woodford (1997) and Woodford (2003b, Chapter 3).¹ The economy contains a continuum of agents indexed by i , with i distributed uniformly on the unit interval. Each agent uses his own labor to produce a differentiated good, sells his good to other agents, and buys goods to consume.

The utility of agent i in period t is given by

$$U_{it} = \frac{C_{it}^{1-\sigma} - 1}{1-\sigma} - \frac{L_{it}^{1+\psi}}{1+\psi}, \quad (1)$$

where L_{it} is the agent's labor supply in period t and C_{it} is a CES aggregator of the agent's consumption of different goods

$$C_{it} = \left[\int_0^1 (C_{it}^j)^{(\gamma-1)/\gamma} dj \right]^{\gamma/(\gamma-1)}, \quad (2)$$

where C_{it}^j is agent i 's consumption of the good produced by agent j . The parameters measuring risk aversion (σ) and the marginal disutility of labor supply (ψ) are non-negative, while the elasticity of substitution between different goods (γ) is larger than one. The agent's production function is

$$Y_{it} = A_i L_{it}, \quad (3)$$

where Y is his level of output and A is an aggregate productivity shifter that follows an arbitrary stochastic process.

As is well-known (e.g., Woodford, 2003b, Chapter 2), utility maximization with this form of utility implies that the demand for each good depends on aggregate

¹A more detailed derivation of the model can be found in the appendix of Ball et al. (2003).

spending and the good's relative price. In logs, the demand function is

$$y_{it} = y_t - \gamma(p_{it} - p_t), \quad (4)$$

where y_i is the log of output by farmer i , y is the log of aggregate output, p_i is the log of the price charged for good i , and p is the log of the price index for the aggregate consumption good C .

We assume that a government levies a proportional sales tax τ_t on all goods, which follows some stationary stochastic process. The tax revenues are used to finance equal lump sum rebates to all agents.

These assumptions determine the price for good i that maximizes farmer i 's utility:

$$p_{it}^* = p_t + \alpha(y_t - y_t^N) + u_t, \quad (5)$$

where $\alpha = (\psi + \sigma)/(1 + \gamma\psi)$, y^N is the natural level of output, and u reflects random variation in taxes. The natural level of output is defined as the level when prices are flexible (so $p_i = p_i^*$ for all i) and the tax rate is at its average level $\bar{\tau}$. It is given by

$$y_t^N = \frac{(1 + \psi) \log(A_t) - \log(\gamma/[(1 - \gamma)(1 - \bar{\tau})])}{\psi + \sigma}. \quad (6)$$

Note that y^N varies with the productivity shock A . Finally, the shock u is given by $u = \log((1 - \bar{\tau})/(1 - \tau))/(1 + \gamma\psi)$. The intuition is simple: as in standard macro models with imperfect competition (e.g., Romer, 2001, Chapter 6), an agent's desired price increases one-for-one with the aggregate price level and depends positively on total spending in the economy. The price decreases with a positive shock to productivity and the natural rate (y^N) and increases with a rise in indirect taxation (u).

We introduce random taxation because it causes variation in farmer's "mark-ups"—the gap between the desired price and the marginal cost of producing output. The key feature of these shocks is that they change the equilibrium level of output under flexible prices without changing the efficient level of output—that is, they cause variation in the level of distortions. In this way, the shocks differ from productivity shocks, which cause changes in equilibrium output that are efficient. Markup shocks are becoming a standard feature of models used to analyze monetary policy (e.g., Woodford, 2003b; Clarida et al., 2002; Steinsson, 2003). Shifts in markups can be justified in different ways—for example, there can be shifts in the degree of collusion in an industry, or shifts in the aggressiveness of wage bargainers in models with a labor market. We introduce variable taxation because it is the simplest way to generate time-varying markups. However, one can interpret our results as applying to markup variations arising from other sources as well.²

Finally, we again follow Woodford (2003b) and assume the existence of complete financial markets. This implies perfect risk-sharing in consumption across all agents.

²In Ball and Mankiw (1995), for example, supply shocks arise from the interaction of menu costs and asymmetries in the distribution of relative shocks. The key effect of these shocks, however, is to change average markups, because firms with positive relative shocks respond more or less than firms with negative shocks.

The consumption level of all agents is the same (C_{it} is the same for all i), although the levels of labor supply and output can vary.

3.2. Price setting and sticky information

Here we follow [Mankiw and Reis \(2002\)](#). Each period, a fraction λ of farmers, chosen randomly, receives complete information on the state of the economy. Prices are perfectly flexible in the sense that prices are adjusted by all farmers in each period. However, prices are set based only on the last information received by each farmer. We take a first-order approximation to optimal price-setting, which yields certainty-equivalent behavior. Thus, when farmer i sets his price in period t , he sets it equal to his expectation of the optimal price:

$$p_{it}^k = E_{t-k} p_{it}^*, \tag{7}$$

where $t - k$ is the last period when the farmer received information.

Again taking a first-order approximation, we measure the log of the price level, p , by the average of the individual prices p_i . Thus p is the average of prices set based on information at all past dates, weighted by the proportion of firms that last received information in each period. Mankiw and Reis aggregate the rule for individual price setting, (5), to derive the behavior of the price level in the absence of markup and productivity shocks. Extending their analysis to include these shocks is straightforward (see the appendix of [Ball et al., 2003](#)). The bottom line can be expressed as an equation either for the price level, or for aggregate inflation $\pi_t = p_t - p_{t-1}$:

$$p_t = \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j E_{t-j} [p_t + \alpha(y_t - y_t^N) + u_t], \tag{8}$$

$$\begin{aligned} \pi_t &= \frac{\alpha\lambda}{1 - \lambda} (y_t - y_t^N) + \frac{\lambda}{1 - \lambda} u_t \\ &+ \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j E_{t-1-j} [\pi_t + \alpha(\Delta y_t - \Delta y_t^N) + \Delta u_t]. \end{aligned} \tag{9}$$

These equations show that the price level and its rate of change depend on current output and the shocks in the model, as well as past expectations of prices, output, and shocks at the various times farmers receive information. Not surprisingly, an adverse technology shock or a rise in the desired markup raises the current price level.

The second version of the equation is the sticky information Phillips curve that Mankiw and Reis propose as a replacement for the new Keynesian Phillips curve arising from the Calvo model of staggered price adjustment. Mankiw and Reis show that their model performs better at fitting the stylized facts about the output–inflation trade-off than does the new Keynesian equation. In particular, it captures the fact that a monetary contraction that gradually reduces inflation also reduces output, whereas the Calvo model implies that gradual disinflations are expansionary.

3.3. The demand side and policy

The equation for the price level, (8), summarizes the supply side of the model. We close the model with the simplest possible demand side: the quantity equation. That is, we assume a cash-in-advance constraint that implies that nominal spending is proportional to the money supply, or

$$y_t = m_t - p_t, \quad (10)$$

where m is the log of the money supply.

We assume that policymakers control the money supply up to a white-noise control error. In every period t , they choose a money-supply target \hat{m}_t , and

$$m_t = \hat{m}_t + e_t, \quad (11)$$

where e_t is a serially uncorrelated control error. \hat{m}_t is set in period $t - 1$, before the productivity and markup shocks for t are observed. These assumptions capture the facts that policy decisions affect the economy with a lag and that there are exogenous shifts in aggregate demand that policy cannot offset contemporaneously. We interpret the monetary control error as standing in for other kinds of demand shocks, such as shifts in fiscal policy or in confidence.

In principle, one can solve for the optimal money-supply rule in this model. A different way of analyzing the problem is more appealing, however. Using Eqs. (8)–(11), one can derive a reduced form for the price level p_t as a linear function of the target money supply, \hat{m}_t , variables determined before t , and shocks at t (the monetary control error and the unexpected components of the productivity and markup shocks).³ Since the shocks are unforecastable, this implies a linear relation between \hat{m}_t and the expectation of the price level $E_{t-1}p_t$, given the predetermined variables. Thus the choice at $t - 1$ of a target money supply is implicitly a choice of the expected price level at t , that is, a price-level target. We can therefore interpret $E_{t-1}p_t$ as the policy instrument, and solve for the optimal rule for this target as a function of all information through period $t - 1$. One can then find the money-supply rule that implements the optimal price-targeting rule, but we will not emphasize this form of policy.

Note that our model of the demand side is simpler than those in most recent work on policy rules. It is common to assume that policymakers control an interest rate, and to relate the interest rate to output through an IS or aggregate-spending equation. If we added such an equation to our model, we could still interpret the expected price level as the policy instrument; a rule for the price target would implicitly define an interest rate rule that implements it. We choose not to introduce an IS curve because there is currently no consensus about the right specification.⁴

³This derivation is provided in Ball et al. (2003).

⁴The debate about the IS curve parallels the debate about the Phillips curve, with some authors advocating forward-looking specifications and others advocating backward-looking specifications. Gabaix and Laibson's (2002) work on lagged consumption responses to the stock market suggests that applying the sticky-information assumption to spending decisions and thus the IS curve may be a fruitful line of

4. Welfare

The next step is to consider welfare in our model, which will provide the foundation for the policy analysis in the next section.

4.1. Woodford's approximation

We define welfare in a period as the average level of utility across all farmers. We assume the discount factor approaches one, so policymakers seek to minimize the unconditional expectation of welfare.⁵ Following similar steps to Woodford (2003b), a second-order approximation of farmers' utility functions is

$$-Var(y_t - y_t^N) - \chi E[Var_i(y_{it} - y_t)] + \text{t.i.p.}, \quad (12)$$

where $\chi = (\psi + \gamma^{-1})/(\psi + \sigma)$ and t.i.p stands for terms independent of policy.⁶

The two terms in this expression are intuitive. The first is the variability of output around the natural level. (Note that the natural level of output is not efficient, because of the distortion caused by the average markup. However, variability around y^N differs from variability around efficient output only by a constant independent of policy.) The second term captures the cross-sectional variability of output across different firms. Variability at the firm level is inefficient because it creates variability in labor supply around the efficient level. Throughout our analysis we define the "optimal" monetary policy as the one that minimizes the expression in Eq. (12).

Another way of writing the objective function is useful. The demand equation (4) implies that a farmer's output y_i is proportional to his relative price $p_i - p$. Thus the variance of y_i is proportional to the variance of $p_i - p$. Thus we can write welfare as

$$-Var(y_t - y_t^N) - \omega E[Var_i(p_{it} - p_t)] + \text{t.i.p.}, \quad (13)$$

where $\omega = \chi\gamma^2$. That is, welfare depends on the variance of the output gap and the cross-sectional variance of relative prices.

4.2. Welfare with Mankiw–Reis price-setting

To analyze the model, we need to determine how the terms in the loss function depend on aggregate variables. Here is the first point at which our results for the sticky-information model depart from those for the Calvo model. In the Calvo model, Woodford (2003b) shows that the cross-sectional variability of prices is determined by current and lagged values of inflation squared. In contrast, the

(footnote continued)

pursuit. In the model we consider here, however, the stickiness of information affects only the yeoman farmer's price setting decision.

⁵While the assumption that the discount factor approaches one simplifies the exposition, most of our results can be obtained with a discount factor smaller than one.

⁶The appendix of Ball et al. (2003) formally derives this second-order approximation.

following result gives an expression for cross-sectional price variability in terms of aggregate variables in our model:

Lemma 1.

$$\text{Var}_t(p_{it} - p_t) = \sum_{j=1}^{\infty} \eta_j (p_t - E_{t-j}(p_t))^2, \quad (14)$$

where

$$\eta_j \equiv \left(\frac{\lambda(1-\lambda)^j}{(1-(1-\lambda)^j)(1-(1-\lambda)^{j+1})} \right).$$

Thus, the variance of relative prices depends on the squared deviations of the price level from the levels expected at all past dates. The weights η_j decline as j increases.

This result is derived in the appendix of Ball et al. (2003). The intuition is straightforward. Surprises in the aggregate price level create micro price variability because prices for a given period are set based on information at various times in the past. In period t , surprises since $t-j$ affect the prices of agents who have received information since then, but not the prices of other agents.

The bottom line, therefore, is that welfare is approximated by a weighted average of Eq. (14) and the variance of $(y - y^N)$. In the analysis below, we experiment with various weights on these two terms. Fortunately, the results are not very sensitive to the weights we assume. In particular, most of our analysis holds even if the weight on (14) is zero, so that the loss function is $\text{Var}(y - y^N)$, the simple loss function assumed in much work on optimal monetary policy.

5. Optimal policy with demand and productivity shocks

We now turn to our central question: what is the optimal policy rule in the model? In this section, we simplify things by eliminating the markup shock u from the model. In this case, the model has two shocks, the productivity shock and the demand shock (i.e., monetary control error). The optimal policy rule turns out to be very simple. The next section discusses the complications arising from the introduction of markup shocks.

5.1. Results

Recall that we can describe policy as a rule for the price target, $E_{t-1}p_t$. In our model,

Proposition 1. *With only productivity and demand shocks, the set of optimal policies is the set of rules $E_{t-1}p_t = K_t$, where K_t follows a deterministic path known at the beginning of time.*

To interpret this result, note first that one optimal policy is the one in which K_t is a constant. In this case, policy targets a fixed price level, the purest form of “price level targeting.” According to the Proposition, however, it is optimal also to allow the target to change over time, as long as the changes are deterministic. The target price level could grow at a constant rate or follow some more complicated (but predictable) path.

Proposition 1 is proved in the appendix. We note here that it is not surprising that any deterministic path for the price level is as good as any other. In the model, the non-neutrality of monetary policy arises from surprises that it creates that are not observed immediately by all price setters. A deterministic price path creates no surprises and thus has no real effects.

Having shown what policies are optimal, it is easy to point out some that are not optimal:

Corollary 1. *Inflation targeting—a policy that sets $E_{t-1}\pi_t$ equal to a constant—is suboptimal. So is nominal income targeting—a policy that sets $E_{t-1}(p_t + y_t)$ equal to a constant.*

This result follows from the facts that only price-level targeting is optimal, and that inflation and nominal-income targeting differ from price-level targeting. It is well-known that inflation targeting allows “base drift”: it means that the target for the price level is adjusted one-for-one with past shocks to the price level. Thus the price target does not follow a deterministic path, as required by Proposition 1. Nominal income targeting would be equivalent to price-level targeting if there were only demand shocks. However, in the case of productivity shocks, nominal income targeting implies that the shocks have persistent effects on both the price level and output, so the expected price level is not deterministic. It is optimal for output alone to respond to a productivity shock, as this keeps output in line with its natural level y_t^N .

5.2. Discussion

These results are significant in light of the current practice of policymaking. There is a strong trend in the world towards inflation targeting, with dozens of countries adopting some version of this policy since the early 1990s. Many observers suggest that U.S. policy is close to inflation targeting, although the target is not explicit. In our model, however, this popular policy is sub-optimal. Welfare is higher if policymakers target the price level rather than the inflation rate—that is, if they eliminate base drift in prices.

To understand this result, it is useful to compare the dynamic effects of shocks under price-level and inflation targeting. We focus here on demand shocks; the story for productivity shocks is similar, since demand and productivity shocks have symmetric effects on the output gap $y - y^N$. Fig. 1 shows the responses of the aggregate price level and aggregate output to a unit demand shock under the two policies, assuming $\alpha = 0.1$ and $\lambda = 0.25$, the baseline values in Mankiw and Reis. The figure also shows the response of the cross-sectional variance of prices. In the

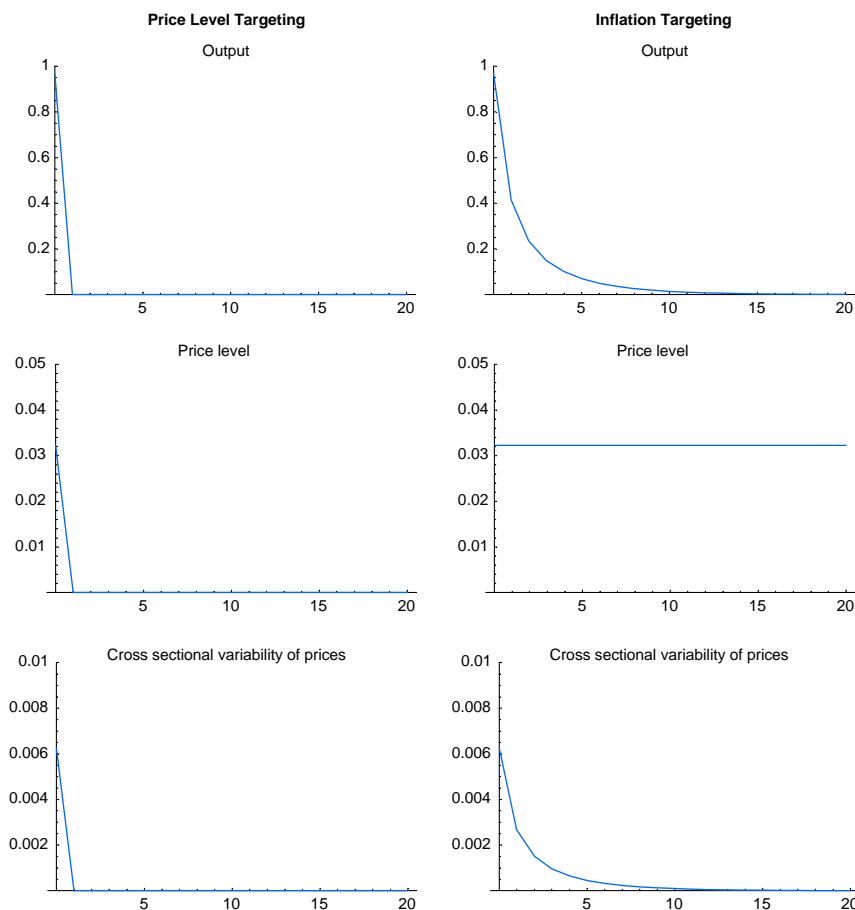


Fig. 1. Impact of a demand shock with price level or inflation targeting.

period when a shock occurs, all the effects are the same under price-level and inflation targeting, reflecting the fact that policy cannot respond contemporaneously to shocks. After the initial period, price targeting implies that all the effects disappear completely. In contrast, the effects on both output and relative prices die out slowly under inflation targeting. This persistence in the effects of shocks increases both components of the social loss function, the variance of the output gap and the variance of relative prices.

To interpret these results, suppose the economy starts in a steady state in which all prices are zero, the output gap is zero, and this situation is expected to continue forever. A demand shock occurs in period zero; it inevitably raises output at zero and it raises the price level as well because firms who receive information in that period adjust their current prices. But these effects disappear in subsequent periods under price-level targeting. In these periods, some prices are set by firms who have not

observed the shock, and therefore are fixed at zero. Price-level targeting means the monetary authority must induce informed firms to set prices of zero as well, so the aggregate price level returns to zero. In this case relative-price variability is eliminated, since informed and uninformed firms set the same prices. In addition, to induce the informed firms to choose the same price as the uninformed, monetary policy must produce an output gap of zero, since only a zero output gap implies a desired relative price of zero. Thus price targeting eliminates the output effects of the shock after period zero as well as the relative-price effects.

In contrast, inflation targeting means the price level must remain at a positive level after the shock raises it in period zero. Because the uninformed still choose prices of zero in subsequent periods, the monetary authority must induce the informed to choose positive prices to keep the aggregate price level positive. This creates variability across the prices of informed and uninformed firms, and it requires a positive output gap to induce the informed to set positive relative prices.

Underlying these results is the fact that demand shocks affect real variables in the sticky information model by creating surprises in the price level. These surprises create both relative price variability (see Eq. (14)) and output movements (note from Eq. (8) that the output gap is zero if the price level equals the level expected in all previous periods). To stabilize the economy, the monetary authority needs to minimize price surprises. It does so by reversing deviations of the price level from its expected path as quickly as possible. It should not exacerbate uncertainty about the price level by allowing base drift.

6. Optimal policy with markup shocks

The optimal policy is more complicated in the presence of markup shocks. In this case, strict price level targeting is no longer optimal. Instead, policymakers must make special accommodations for markup shocks, the nature of which depends on the serial correlation of the shocks.

We assume the markup shock follows an arbitrary stationary process with MA representation $u_t = \sum_{j=0}^{\infty} \rho_j \varepsilon_{t-j}$, where ε_t is a white-noise innovation. Given this notation, we have

Proposition 2. *With productivity, demand, and markup shocks, the set of optimal policies is the set of rules:*

$$E_{t-1}p_t = K_t + \sum_{j=1}^{\infty} \phi_j \varepsilon_{t-j}, \tag{15}$$

where

$$\phi_j \equiv \frac{\rho_j}{\alpha^2 \omega + (1 - \lambda)^{j+1} / (1 - (1 - \lambda)^{j+1})},$$

and K_t is again deterministic.

This proposition is proved in the appendix.

To interpret Proposition 2, consider first the case of white-noise markup shocks. In this case all the ρ_j 's are zero, so all the ϕ_j 's are zero: the optimal policy is strict price level targeting as before. If the markup shock is serially correlated, however, then the price target is adjusted in response to past shocks. If all the ρ_j 's are positive, then an increase in the markup raises the price target for awhile. However, as long as the markup is stationary, the price level is also stationary: it eventually returns to the target path given by K_I .

Fig. 2 illustrates this result for particular parameterizations of the model. The figure shows how the price level responds to a unit shock to the innovation ε . In this simulation, again we set $\alpha = 0.1$ and $\lambda = 0.25$. We assume the markup shocks follow an AR(1) process with a coefficient of $\bar{\rho}$ so $\rho_j = \bar{\rho}^j$. We experiment with two values for $\bar{\rho}$, 0.8 and 0.4. Finally, we set the parameter ω equal to one, which means that output variance and relative-price variance have the same weight in the social loss function.

In Fig. 2, the price level jumps up at the time of the markup shock, which is inevitable since policy cannot respond contemporaneously. In contrast to the cases of demand and productivity shocks, the price level remains high after the initial period; it either follows a hump-shaped pattern (for $\bar{\rho} = 0.8$) or falls monotonically (for $\bar{\rho} = 0.4$). In either case, the price level eventually converges back to its long-run target.

The optimal response to markup shocks reflects a trade-off between the goals of stabilizing relative prices and stabilizing output. As with demand and productivity shocks, relative-price variance is minimized by a strict price level target, as this causes informed firms to set prices at the same level as uninformed firms. Such a policy requires a large fall in output, however: a higher markup means firms desire higher relative prices for a given output level, so low output is needed to induce firms to choose a zero relative price. Given this trade-off, optimal policy allows output to fall somewhat but not enough to stabilize relative prices fully. It induces informed firms to set positive prices as long as the shock persists.

Our model implies that the optimal response to markup shocks depends on the serial correlation of the shocks. Unfortunately, it is difficult to calibrate this serial correlation. Like other authors, we would like to interpret markup shifts as metaphors for the “supply shocks” that influence real-world inflation, but there is no consensus about the nature of these shocks. If we interpret the model literally, so markup shocks arise from tax changes, it is natural to interpret these changes as highly persistent. Similarly, if markup shocks are interpreted as changes in union aggressiveness, it is plausible to assume persistence in union attitudes. However, some kinds of real-world supply shocks may have more transitory effects on desired prices. We need a better understanding of supply shocks to make progress on designing optimal policies.

Our model suggests that policymakers should treat price-level movements in a special way if they arise from certain kinds of supply shocks. In this respect, the model's prescriptions match policymaking in many countries. It is common for inflation targeters to allow deviations from their normal targets when supply shocks

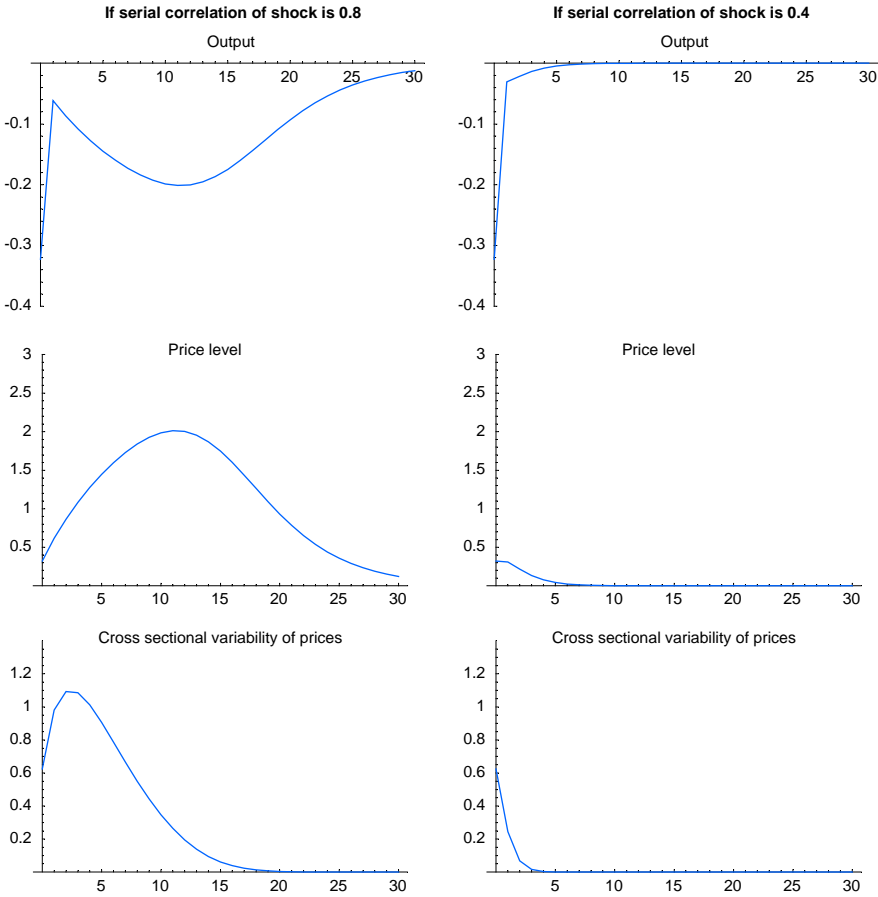


Fig. 2. Impact of a markup shock under the optimal policy.

occur. Sometimes, this is accomplished by explicit “caveats” to policy, as in New Zealand. More often, policymakers allow temporary deviations from the target by focusing on an “underlying” or “core” measure of inflation that strips out the effects of supply shocks. Our analysis suggests that policymakers are on the right track in making exceptions to their policies for supply shocks—although the basic policy target should be the price level, not inflation as in current practice.

7. A simple description of optimal policy

So far, we have described optimal policy as a function of the exogenous driving variables. There is, however, a simpler description of optimal policy in terms of the endogenous variables:

Proposition 3. *For an economy that experiences shocks to demand, productivity, and markups, optimal monetary policy can be described as:*

$$E_{t-1}p_t = K_t - \frac{1}{\alpha\omega} E_{t-1}(y_t - y_t^N). \quad (16)$$

That is, monetary policy commits to a deterministic path for the price level, represented here by K_t , and allows the expected price level to deviate from the target path only if output is expected to deviate from its natural rate.⁷

This proposition encompasses our previous results. In the case of shocks to demand or productivity, discussed in Section 4, optimal policy has the property that $E_{t-1}(y_t - y_t^N) = 0$. Thus, Proposition 3 includes Proposition 1 as a special case. In the case of shocks to the markup, the expected price level varies as a complicated function of the shocks, as shown in Proposition 2, but it does so in a way that precisely parallels the expected output gap. This can be seen in Fig. 2, where the paths of output and the price level after period zero are reflections of each other.

One remarkable fact about Proposition 3 is that it closely resembles the monetary policy rule proposed by Hall (1984), who dubbed it an “elastic price standard.” Hall does not present a formal model to back up his proposal, but his informal and intuitive arguments are consistent with the model we have explored here. One difference is that Hall advocates a constant target (K_t), whereas in our model a rising target for the price level would serve equally well. In either case, the expected output gap determines how much the central bank can allow the price level to deviate from its target.

8. Comparison to the new Keynesian Phillips curve

We have argued that the sticky-information Phillips curve of Mankiw and Reis is a better tool for monetary-policy analysis than the popular new Keynesian Phillips curve, because it has more realistic implications about the interactions of output and inflation. But does the choice of a Phillips curve make a difference for our results about optimal policy? To address this issue, we now reconsider our policy problem using the new Keynesian Phillips curve—the relation derived from Calvo’s (1983) model of staggered price adjustment with full information. As is well known, this Phillips curve (with no discounting) is

$$\pi_t = E_t\pi_{t+1} + \kappa(y_t - y_t^N) + u'_t, \quad \kappa > 0, \quad (17)$$

⁷Clarida et al. (1999, p. 1704) derive an equation for optimal policy that closely resembles this proposition in an analysis based on the Calvo model of price adjustment. This similarity of optimal policies under Calvo and Mankiw–Reis price setting is not robust, however. The result of Clarida et al. is based on the assumption that the monetary authority can set the money supply based on contemporaneous information. Yet if policymakers can affect the economy only with a lag (that is, if there are control errors), as we have assumed, then optimal monetary policy in the Calvo model can no longer be described by such a simple rule. In the next section, we explore more fully the similarities and differences between the two models.

where u'_t is proportional to the markup shock u_t . As discussed by Mankiw and Reis (2002), a critical difference between this equation and the Phillips curve arising from the sticky-information model is the timing of expectations. Rather than past expectations of current inflation, it is the current expectation of next period's inflation that enters the new Keynesian specification. This feature of the model is mainly responsible for its counterfactual predictions.

8.1. Implications for the welfare function

Note first that the choice of a Phillips curve affects the determinants of aggregate welfare. As shown in Section 3, the canonical model of monopolistic competition yields an approximate loss function with two terms, the variance of the output gap and the variance of relative prices. This result is robust, but the determinants of the variance of relative prices depends on the specification of price adjustment. In the Mankiw–Reis model, the variance of relative prices depends on the variance of unexpected changes in the price level (see Eq. (14)). In the Calvo model, it is determined by current and lagged values of inflation squared. Thus the two models produce loss functions that have one term in common, the variance of the output gap, and one term that differs across models. Thus one would expect the models to deliver different results about optimal policies. However, the main differences discussed below do not depend on the difference in loss functions. The differences in results arise even when the weight on the variance of relative prices in the loss function is zero, so that the loss function is simply the variance of the output gap in both models.

8.2. Implications for base drift

We now consider optimal policy in a model that is identical to ours except that the Calvo sticky-price assumption replaces the Mankiw–Reis sticky-information assumption, so that the Phillips curve becomes Eq. (17) above. The following result shows a key difference between the normative implications of the two models:

Proposition 4. *With Calvo price adjustment, the optimal policy rule when there are demand or productivity shocks is: $E_{t-1}\pi_t = \delta\pi_{t-1}$, where $0 < \delta < 1$.*

That is, the Calvo model supports flexible inflation targeting rather than price level targeting. Put differently, allowing base drift in the price level in response to demand and productivity shocks is optimal with Calvo price setting, while it is not with Mankiw–Reis price setting. The proof of the proposition is in the appendix of Ball et al. (2003).

Proposition 4 implies not only a unit root in the price level but also persistence in inflation. In particular, the inflation rate is AR(1) with parameter δ . We believe that the following intuition explains the optimality of inflation persistence in the Calvo model: When the economy experiences a positive demand shock, output and the price level both rise. If the monetary authority responds to this shock by committing to higher inflation in the future, then price setters will respond to the shock by raising

their prices more aggressively. This more aggressive price response implies a smaller initial output response, keeping the economy closer to its natural rate. In the long run, however, the only effect of inflation is relative price variability, so optimal policy always returns to zero inflation.

As in our main model, the nature of optimal policy is more complicated when there are persistent markup shocks. Just as in our model, optimal policy involves a gradual adjustment of the price level in response to a shock. Yet now the price level is non-stationary: it remains forever higher after a shock that raises prices.

The result that inflation rather than price-level targeting is optimal under Calvo pricing may appear surprising in light of the previous literature. A number of authors study models that include the Calvo Phillips curve and conclude that a price level target is optimal. The discrepancy between our results and previous ones are explained by differences in the policy problem being solved. There are two main differences.

First, some previous authors, such as King and Wolman (1999), study models in which it is feasible to stabilize the price level perfectly. That is, there is a perfect relation between a policy instrument and the current price level; there is no such thing as a demand shock that policy cannot offset contemporaneously, as with our assumption of monetary control errors. In the King–Wolman model, the policy of keeping the price level absolutely constant is optimal given the types of shocks they assume. A perfectly constant price level can be interpreted as implementing either a constant price-level target or a zero inflation target. That is, there is no meaningful distinction between price and inflation targeting if there are no shocks that raise the possibility of base drift.⁸

Vestin (1999) compares price and inflation targeting in versions of the Calvo model in which these policies are different, and concludes that price targeting is optimal. The problem he solves, however, is that of the optimal policy rule under “discretion.” That is, he asks whether it is optimal to appoint a policymaker whose objective function is the variance of prices or the variance of inflation, given that the policymaker cannot precommit to his preferred policy because it is not dynamically consistent. In contrast, we study optimal policy under commitment: we assume that policymakers can implement any rule they want and ask which rule maximizes social welfare.⁹

8.3. Implications for trend inflation

A final difference between our results and those for the Calvo model concern the welfare consequences of trend inflation. As shown above, trend inflation is neutral in our model, and so optimal policy can take the form of an increasing price level target

⁸By contrast, Goodfriend and King (1997) briefly consider the issue of imperfect control in a model with staggered price setting. Their analysis confirms our conclusion that base drift in the price level is optimal in this type of model.

⁹Like Vestin, Svensson’s (1999) well-known paper on inflation versus price level targeting emphasizes the case of discretion rather than commitment, but with a different specification for the Phillips curve.

rather than a constant target. In contrast, the optimal inflation target in the Calvo version of the model must be zero. That is, moving from Calvo to Mankiw–Reis price setting means introducing the possibility of trend inflation as well as eliminating base drift in the price level. Here, the difference in loss functions in the two models is crucial. In the Calvo model, welfare depends on the variance of inflation around zero rather than around some mean level, so that zero inflation is optimal. This reflects the fact that even steady inflation causes inefficient relative-price variability under Calvo's assumption of sticky prices.

9. Conclusion

According to the model developed here, optimal monetary policy can be described as flexible targeting of the price level. The central bank should announce a target path of the price level and then commit itself to returning to this path in response to shocks. In response to shocks to demand or productivity, it should return the price level as quickly as possible to its target. In response to shocks to markups, which here reflect a type of supply shock, the central bank should return the price level to its target more gradually over time. Described differently, optimal policy allows the price level to deviate from its target only if output is expected to deviate from its natural rate.

The optimality of targeting the price level rather than the inflation rate is a common theme in the recent literature on monetary policy. For example, [Hall \(1984\)](#) and [Hall and Mankiw \(1994\)](#) argue that a price level target would aid personal financial planning by making the cost of living far in the future more predictable. [Svensson \(1999\)](#) and [Vestin \(1999\)](#) argue for price level targeting on the grounds that it would help solve some of the time-inconsistency problems associated with discretionary monetary policy. [Cover and Pecorino \(2001\)](#) claim that a price target is stabilizing because any inflationary shock automatically causes a decrease in expected inflation and thus an increase in the real interest rate. The analysis in this article is very different. Our households have no money illusion in financial planning, our monetary policymaker can commit to a policy rule, and we omit the effect of expected inflation on aggregate demand. Nonetheless, the bottom line for policy is the same: central banks should target the price level, not the inflation rate.

These results leave open an intriguing question: Why have central banks adopted inflation targeting if price level targeting has all these desirable properties? One possible answer is that central bankers know something about the world that is missing in these theories of optimal policy. But another possibility is that central bankers have been misled by assuming, incorrectly, that some features of the recent monetary regime are structural, while in fact these features would change if policy did.

Price dynamics are the natural place to look for such a mistaken assumption. The recent data are well described by a backward-looking Phillips curve. If this equation for price dynamics were invariant to policy, then inflation targeting would be optimal, and price level targeting would be unattractive ([Ball, 1999](#)). But a radical change in

the monetary policy rule, such as a target for the price level, would most likely alter the reduced-form equation for the Phillips curve. Whether the particular behavioral model we have examined in this article correctly captures the shift that would occur is open to debate. But the results presented here suggest the issue is well worth pursuing, for the implications for monetary policy could not be more profound.

Appendix A

Proof of Proposition 1. Under the proposed policy, since K_t is deterministic, $E_{t-j}p_t = K_t$ for all $j \geq 1$. Taking expectations at $t - 1$ of the equation for the price level (8) we obtain

$$0 = E_{t-1}(y_t - y_t^N) + \sum_{j=1}^{\infty} (1 - \lambda)^j E_{t-j}(y_t - y_t^N).$$

This difference equation has the solution $E_{t-1}(y_t - y_t^N) = 0$, so any shocks at $t - 1$ or before have no effect on the first term in the welfare function. Policy is therefore minimizing the variance of $(y_t - y_t^N)$, which varies only with shocks at time t , which policy can do nothing about.

From Lemma 1 and given the proposed policy:

$$\begin{aligned} \text{Var}_i(p_{it} - p_t) &= \sum_{j=1}^{\infty} \eta_j (p_t - K_t)^2 \\ &= (p_t - E_{t-1}p_t)^2 \left(\sum_{j=1}^{\infty} \eta_j \right). \end{aligned}$$

It then follows that the cross-sectional dispersion of prices is minimized, since unanticipated changes in the price level occur only with respect to innovations at time t , and policy, which is set at $t - 1$, can do nothing about these. Thus with respect to the second term in the welfare function, policy is also reaching the optimum. Thus, K_t minimizes the full welfare function and so it is the optimal policy.

Proof of Propositions 2 and 3. The proof of Propositions 2 and 3 with respect to demand and natural rate shocks follows immediately from Proposition 1. With markup shocks, we represent policy in reaction to a shock as the set of coefficients ϕ_j in the general representation: $E_{t-1}(p_t) = \sum_{j=1}^{\infty} \phi_j \varepsilon_{t-j}$. Similarly, we represent output by the set of φ_j in $y_t = \sum_{j=1}^{\infty} \varphi_j \varepsilon_{t-j}$. Using the notation $A^k = \lambda \sum_{j=0}^k (1 - \lambda)^j$ and the supply curve in (8), then

$$\sum_{j=1}^{\infty} \phi_j \varepsilon_{t-j} = \sum_{j=1}^{\infty} A^j (\phi_j \varepsilon_{t-j} + \alpha \varphi_j \varepsilon_{t-j} + \rho_j \varepsilon_{t-j}).$$

Since this expression must hold for all possible realizations of ε_{t-j} , it then follows that

$$\varphi_j = \frac{1}{\alpha} \left(\frac{1 - A^j}{A^j} \phi_j - \rho_j \right). \quad (\text{A.1})$$

Noting that $p_t - E_{t-i}(p_t) = \sum_{j=0}^{i-1} \phi_j \varepsilon_{t-j}$ and that $\eta_i = 1/A^{i-1} - 1/A^i$, we can then write the objective function as

$$\left[\sum_{j=1}^{\infty} \varphi_j^2 + \omega \sum_{i=1}^{\infty} \left(\frac{1}{A^{i-1}} - \frac{1}{A^i} \right) \sum_{j=0}^{i-1} \phi_j^2 \right] \sigma^2, \quad (\text{A.2})$$

where σ^2 is the variance of ε_t . Minimizing (A.2) subject to (A.1) gives the optimal policy

$$\phi_j = \frac{1}{\alpha^2 \omega + (1 - \lambda)^{j+1} / (1 - (1 - \lambda)^{j+1})} \rho_j.$$

(A key step is to realize that $\sum_{i=j+1}^{\infty} (1/A^{i-1} - 1/A^i) = 1/A^j - 1$.) This proves Proposition 2. As for Proposition 3, we can use the solution above to solve for the output coefficients using (A.1). We find that the relation between output and the price level coefficients is given by $\varphi_j = -\alpha \omega \phi_j$, which proves Proposition 3 when there are markup shocks.

References

- Akerlof, G., 2002. Behavioral macroeconomics and macroeconomic behavior. *American Economic Review* 92, 411–433.
- Alogoskoufis, G., Smith, R., 1991. The Phillips curve, the persistence of inflation and the Lucas critique: evidence from exchange rate regimes. *American Economic Review* 81, 1254–1275.
- Ball, L., 1994. Credible disinflation with staggered price setting. *American Economic Review* 84, 282–289.
- Ball, L., 1999. Efficient rules for monetary policy. *International Finance* 2, 63–83.
- Ball, L., 2000. Near rationality and inflation in two monetary regimes. NBER Working Paper No. 7988.
- Ball, L., Mankiw, N.G., 1995. Relative-price changes as aggregate supply shocks. *Quarterly Journal of Economics* 110, 161–193.
- Ball, L., Mankiw, N.G., Reis, R., 2003. Monetary policy for inattentive economies. NBER Working Paper No. 9491.
- Barsky, R., 1987. The Fisher hypothesis and the forecastability and persistence of inflation. *Journal of Monetary Economics* 19, 3–24.
- Calvo, G.A., 1983. Staggered prices in a utility maximizing framework. *Journal of Monetary Economics* 12, 383–398.
- Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: a new Keynesian perspective. *Journal of Economic Literature* 37, 1661–1707.
- Clarida, R., Gali, J., Gertler, M., 2002. A simple framework for international monetary policy analysis. *Journal of Monetary Economics* 49, 879–904.
- Cover, J.P., Pecorino, P., 2001. Price and output stability under price level targeting. University of Alabama Working Paper WP01-05-01.
- Cover, T.M., Thomas, J.A., 1991. *Elements of Information Theory*. Wiley, New York.

- Fischer, S., 1977. Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of Political Economy* 85, 191–205.
- Friedman, M., 1968. The role of monetary policy. *American Economic Review* 58, 1–17.
- Gabaix, X., Laibson, D., 2002. The 6D bias and the equity premium puzzle. *NBER Macroeconomics Annual* 2001, 257–311.
- Gali, J., Gertler, M., 1999. Inflation dynamics: a structural econometric investigation. *Journal of Monetary Economics* 44, 195–222.
- Goodfriend, M., King, R., 1997. The new neoclassical synthesis and the role of monetary policy. *NBER Macroeconomics Annual* 1997, 231–283.
- Gordon, R., 1997. The time-varying NAIRU and its implications for monetary policy. *Journal of Economic Perspectives* 11, 11–32.
- Hall, R.E., 1984. Monetary strategy with an elastic price standard. In: *Price Stability and Public Policy: A Symposium Sponsored by the Federal Reserve Bank of Kansas City*. Federal Reserve Bank of Kansas City, Jackson Hole.
- Hall, R.E., Mankiw, N.G., 1994. Nominal income targeting. In: Mankiw, N.G. (Ed.), *Monetary Policy*. University of Chicago Press, Chicago.
- King, R.G., Wolman, A.L., 1999. What should monetary policy do if prices are sticky? In: Taylor, J.B. (Ed.), *Monetary Policy Rules*. University of Chicago Press, Chicago, pp. 349–404.
- Lucas, R.E., 1973. Some international evidence on output–inflation trade-offs. *American Economic Review* 63, 326–334.
- Mankiw, N.G., 2001. The inexorable and mysterious tradeoff between inflation and unemployment. *Economic Journal* 111, C45–C61.
- Mankiw, N.G., Reis, R., 2002. Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics* 117, 1295–1328.
- Romer, D., 2001. *Advanced Macroeconomics*, Second ed. McGraw-Hill, Boston.
- Rotemberg, J.J., 1982. Monopolistic price adjustment and aggregate output. *Review of Economic Studies* 44, 517–531.
- Rotemberg, J.J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. *NBER Macroeconomics Annual* 1997, 297–346.
- Sargent, T., 1993. *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- Sims, C.A., 2003. Implications of rational inattention. *Journal of Monetary Economics* 50, 665–690.
- Staiger, D., Stock, J.H., Watson, M.W., 1997. The NAIRU, unemployment, and monetary policy. *Journal of Economic Perspectives* 11, 33–51.
- Steinsson, J., 2003. Optimal monetary policy in an economy with inflation persistence. *Journal of Monetary Economics* 50, 1425–1456.
- Svensson, L.E.O., 1997. Inflation forecast targeting: implementing and monitoring inflation targets. *European Economic Review* 41, 1111–1146.
- Svensson, L.E.O., 1999. Price level targeting vs. inflation targeting: a free lunch? *Journal of Money, Credit and Banking* 31, 277–295.
- Taylor, J.B., 1980. Aggregate dynamics and staggered contracts. *Journal of Political Economy* 88, 1–22.
- Vestin, D., 1999. Price level targeting versus inflation targeting in a forward looking model. Stockholm University, Unpublished.
- Woodford, M., 2003a. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press, Princeton, pp. 25–58.
- Woodford, M., 2003b. *Interest and Prices*. Princeton University Press, Princeton.