

Managing Disinflation under Uncertainty

Mewael F. Tesfaselassie* Eric Schaling†

This Version June 2008

Abstract

In this paper we analyze disinflation policy when a central bank has imperfect information about private sector inflation expectations but learns about them from economic outcomes, which are in part the result of the disinflation policy. The form of uncertainty is manifested as uncertainty about the effect of past disinflation policy on current output gap. Thus current as well as past policy actions matter for output gap determination. We derive the optimal policy under learning (DOP) and compare it two limiting cases—certainty equivalence policy (CEP) and cautionary policy (CP). It turns out that under the DOP inflation stay between the levels implied by the CEP and the CP. A novel result is that this holds irrespective of the initial level of inflation. Moreover, while at high levels of inherited inflation the DOP moves closer to the CEP, at low levels of inherited inflation the DOP resembles the CP.

Keywords: Learning, Inflation Expectations, Disinflation Policy, Separation Principle, Kalman Filter, Optimal Control

JEL Codes: C53, E52, F33

*Corresponding author: Kiel Institute for the World Economy, Duesternbrooker Weg 120, 24105 Kiel, Germany. E-mail: mewael.tesfaselassie@ifw-kiel.de, Tel: +49 431 8814 273.

†South African Reserve Bank Chair, University of Pretoria and CentER for Economic Research, Tilburg University, The Netherlands; Pretoria 0002, South Africa. E-mail: eric.schaling@up.ac.za, Tel: +27 12 420 3820.

1 Introduction

How should central bank's manage a disinflation process? The received view in the literature—as expressed by King (1996) at the Kansas City Fed symposium on Achieving Price Stability—seems to be for a gradual timetable, with inflation targets consistently set below the public's inflation expectations. As King puts it, “the aim was not to bring inflation down to below 2 percent by the next month, or even the next year. It was to approach price stability gradually ... some four to five years ahead”. However, King also raises the possibility that a central bank may try to convince the private sector of its commitment to price stability by choosing to reduce inflation towards the inflation target quickly. He calls this ‘teaching by doing’. Then the choice of a particular inflation rate influences the speed at which expectations adjust to price stability.

King shows how the optimal speed of disinflation depends crucially on whether the private sector immediately believes in the new low inflation regime or not. If they do, the best strategy is to disinflate quickly, since the output costs are zero. Of course, if expectations are slower to adapt, disinflation should be more gradual as well. Teaching by doing effects have also been analyzed by Hoerberichts and Schaling (2000) and Schaling (2003). They find that allowing for ‘teaching by doing’ effects always speeds up the disinflation vis-à-vis the case where this effect is absent. Thus, their result is that ‘speed’ in the disinflation process does not necessarily ‘kill’ in the sense of creating large output losses.

In this paper we analyze optimal disinflation policy when the central bank faces uncertainty regarding the prevailing level of inflation expectations and uses data from the economy to learn about them. The process of learning involves updating in real time using standard Kalman filtering methods. We find that when the central bank internalizes the effect of its current disinflation policy on future uncertainty about inflation expectations, it disinflates more than implied by a policy of certainty equivalence but less than implied by a cautionary policy. Under learning, the optimal policy is a nonlinear function of the state of the economy and the central bank's belief about inflation expectations. It turns out that, given its belief, the optimal policy stays close to a certainty equivalence policy when the inherited level of inflation is

high. When the inherited level of inflation is low, the optimal policy stays close a policy that implies caution (as first shown by Brainard (1967), but now extended to a dynamic context). In our case, a cautionary policy disinflates more than implied by the certainty equivalence policy.

Regarding the focus on learning and control, our paper is related to other studies that have analyzed the role of parameter uncertainty in optimal monetary policy (see e.g., Bertocchi and Spagat (1993), Balvers and Cosimano (1994), Wieland (2000a), Wieland (2000b), Ellison and Valla (2001), Yetman (2003), Ellison (2006), and Svensson and Williams (2007)).¹ However, these studies typically assume the presence of uncertainty about the effects of *current* policy actions on the economy.² Also, a common feature of most of these studies is that the linear economic process subject to central bank control is static. By contrast, in our model, imperfect information about inflation expectations is reflected as uncertainty about the effects of *past* policy actions. Thus, in our case the lag of the policy instrument is crucial for the dynamics of the economy.

The remainder of the paper is organized as follows. Section 2 presents a simple model and discusses private sector behavior regarding the credibility of the central bank's inflation target. It also discusses belief updating on the part of the central bank. In section 3 we derive the optimal degree of disinflation under alternative scenarios—certainty equivalence, cautionary and dynamically optimal policies and present sensitivity analysis to changes in the key parameters. In section 4 we discuss convergence of limit beliefs and policies. Finally, section 5 gives concluding remarks.

¹Formally, the numerical methods for solving optimal control under parameter uncertainty originate in the *dual control* literature (see e.g., Prescott (1972)). The dual control literature has shown that the so-called separation principle may not hold, and a trade-off between estimation and control arises because current actions influence estimation (learning) and provide information that may improve future performance. See e.g., Wieland (2000b) for a detailed discussion.

²As our focus is on parameter uncertainty, we abstract from other forms of uncertainty, such as model uncertainty (see e.g., Cogley and Colacito and Sargent (2005)), which are also important for monetary policymakers.

2 The Model

King (1996) discusses disinflation policy using a simple macroeconomic model, which combines nominal wage and price stickiness and slow adjustment of expectations to a new monetary policy regime. The model has three key equations—aggregate supply, monetary policy preferences and inflation expectations. Aggregate supply exceeds the natural rate of output when inflation is higher than was expected by agents when nominal contracts were set. This is captured by a simple short-run Phillips curve (see also Cogley and Colacito and Sargent (2005)).³

$$z_t = \pi_t - \pi_t^e + u_t \tag{1}$$

Where π_t is the rate of inflation, z_t is the output gap and π_t^e indicates that the expectation of inflation is the subjective expectation (belief) of private agents. As in King (1996), this belief does not necessarily coincide with rational expectations. The model is not restrictive as long as inflation expectations are in part influenced by past monetary policy (see e.g., Bomfim and Rudebusch (2000) and Yetman (2003)).

The regime change is represented by a new inflation target $\pi^* = 0$, which is announced to the public at the end of $t - 1$. The new target is lower than the initial steady state inflation, denoted by π_0 .

The central bank's objective as of period t is to choose a sequence of current and future inflation rates $\{\pi_\tau\}_{\tau=t}^\infty$ so as to minimize its intertemporal loss

$$E_{t-1}^c \sum_{\tau=t}^{\infty} \delta^{\tau-t} L(\pi_\tau, z_\tau) \tag{2}$$

where

$$L(\pi_t, z_t) = \frac{1}{2}(z_t - z^*)^2 + \frac{\alpha}{2}(\pi_t - \pi^*)^2 \tag{3}$$

³In their analysis of U.S. monetary policy experimentation in the 1960s, Cogley and Colacito and Sargent (2005) use a model similar to ours but with unemployment instead of output.

and E_{t-1}^c denotes expectations conditional on the central bank's information at the end of period $t - 1$. The parameter $\alpha \geq 0$ is the relative weight on inflation stabilization while δ is the discount factor ($0 < \delta < 1$).

Inflation expectations

King (1996) analyzes two extreme cases of inflation formation: (1) a completely credible policy regime where private sector expectations adjust immediately to the inflation target (since the announcement is fully credible)—this is the case of rational or model consistent expectations; (2) ‘endogenous forecasting’, where the private sector expectations depend on monetary policy choices (that is on actions, not just on words) made in the new regime.

In general, expectations are affected both by the inflation target and by actual inflation performance. After experiencing high inflation for a long period of time, there may be good reasons for the private sector not to believe the disinflation policy fully (See also Bomfim and Rudebusch (2000) and Schaling (2003)). In light of this, we assume that private sector inflation expectations follow a simple rule, that is a linear function of the inflation target and the lagged inflation rate

$$\pi_t^e = \gamma\pi_{t-1} + (1 - \gamma)\pi^* = \gamma\pi_{t-1} \tag{4}$$

where $0 \leq \gamma \leq 1$ captures the degree of credibility of the new regime. The closer γ to 0, the higher is the credibility of the regime change.

We introduce uncertainty by supposing that the central bank can not observe private sector expectations directly. Moreover, we assume the central bank does not know the credibility parameter γ and can not observe (even ex post) the shock u_t , so that it can not infer private sector inflation expectations from (4). In period t , the central bank observes z_t only after it has chosen π_t and the shock u_t has realized. Under this scenario, the unobservability of inflation expectations is manifested as parameter uncertainty—the central bank does not know the degree of credibility, as measured by γ . It follows that optimal monetary policy affects (and is affected by) the dynamics of belief updating about γ . In other words *control* and *estimation* of

the economy are interrelated.⁴

Belief Updating

Let $y_{t-1} = -\pi_{t-1}$. Substituting (4) into (1), the actual dynamics of the Phillips curve is given by

$$z_t = \pi_t + \gamma y_{t-1} + u_t \quad (5)$$

$$y_t = -\pi_t \quad (6)$$

The information set at the end of period t is $\Omega_t = \{z_t, z_{t-1}, \dots\}$. Under parameter uncertainty, the central bank's belief about γ , before setting π_t , can be characterized by a prior mean $c_{t-1} = E(\gamma|\Omega_{t-1})$ and prior variance $p_{t-1} = E(\gamma - c_{t-1})^2$. After π_t is chosen and z_t realizes, the central bank updates its belief to c_t and p_t . Updating takes a standard recursive structure,

$$c_t = c_{t-1} + y_{t-1} p_{t-1} F_{t-1}^{-1} (z_t - \pi_t - c_{t-1} y_{t-1}) \quad (7)$$

$$p_t = p_{t-1} - p_{t-1}^2 y_{t-1}^2 F_{t-1}^{-1} \quad (8)$$

where $F_{t-1} = p_{t-1} y_{t-1}^2 + \sigma_u^2$. These two equations represent the learning channel through which the current policy action, π_t , affects future beliefs about γ , i.e., c_{t+j}, p_{t+j+1} for $j = 0, 1, 2, \dots$. The filtering process maps the sequence of prediction errors into a sequence of revisions; and the term $y_{t-1} p_{t-1} F_{t-1}^{-1}$ on the right hand side of (7) and (8) is usually referred to as the Kalman gain, which is a nonlinear function of period $t-1$ policy π_{t-1} .⁵

⁴This bounded rationality assumption follows, among others, Marcet and Sargent (1988) and Evans and Honkapohja (2001) in that in forecasting private sector inflation expectations, the central bank acts like an econometrician. Evans and Honkapohja (2001) and others have studied determinacy and learnability of rational expectations equilibria when the private sector has to learn about key parameters and the central bank follows a simple monetary policy rule. Here we are interested in how uncertainty and learning affect the central bank's optimal control problem (see for e.g. Beck and Wieland (2002) and Tesfaselassie et al. (2006)).

⁵See Tesfaselassie (2005) for a detailed derivation.

3 Optimal Disinflation Policy

We distinguish three policy scenarios—certainty equivalence policy, cautionary policy and dynamically optimal policy. The three policies differ in their approaches to parameter uncertainty and learning. The certainty equivalence policy and the cautionary policy ignore the non-linear updating equations, and so policy is conducted under passive learning and the policy rules are linear in the state variable y_{t-1} . The certainty equivalence policy is an extreme case, where the prior variance is set to zero ($p_{t-1} = 0$). The dynamically optimal policy takes account of the updating equations and thus represents an active learning policy. In that case, the policy rule is a non-linear function of y_{t-1} and can be solved for only numerically.

In the next two sections we consider the cases of certainty equivalence and cautionary policy. In both cases the central bank disregards the effect of current policy actions on future estimation and control. In other words, by ignoring the non-linear updating equations for c_t and p_t , the central bank treats control and estimation separately. Learning is in effect passive in the sense that, the central bank optimizes assuming its actions will not affect future beliefs but updates its beliefs once new data arrives (Sargent (1999)).

The Certainty Equivalence Policy

Under certainty equivalence the central bank ignores parameter uncertainty, being fully confident about its prior $c_{t-1} = c$. Its belief about γ is thus given by the pair $(c_{t-1}, p_{t-1}) = (c, 0)$. The sequence of events is as follows.

Certainty Equivalence: Timing of events in period t

Stage 1	Stage 2	Stage 3	Stage 4
private sector sets π_t^e	central bank sets $E_{t-1}^c \pi_t^e$	central bank chooses $\pi_t = \pi(y_{t-1}, c; \cdot)$	u_t realizes, determining z_t

The minimization problem is

$$\min_{\{\pi_\tau\}_{\tau=t}^{\infty}} E^c \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} L(z_\tau, \pi_\tau) | y_{t-1} \right] \quad (9)$$

subject to the linear constraint (5). Since under CE the control problem is linear-quadratic, the solution for the optimal level of π_t is similar to that under perfect knowledge. The certainty-equivalence rule simply replace γ with its conditional mean c . As is shown in Appendix, the solution for π_t is given by ('CE' denotes certainty equivalence)

$$\pi_t^{CE} = \frac{z^* - \frac{1}{2}\delta\mu_1^{CE}}{1 + \alpha + \delta\mu_2^{CE}} + \frac{cy_{t-1}}{1 + \alpha + \delta\mu_2^{CE}} \quad (10)$$

where

$$\begin{aligned} \mu_2^{CE} &= \frac{-1 - \alpha + \delta c^2 + \sqrt{4\alpha\delta c^2 + (1 + \alpha - \delta c^2)^2}}{2\delta} \\ \mu_1^{CE} &= \frac{2cz^*(\alpha + \delta\mu_2^{CE})}{1 + \alpha - c\delta + \delta\mu_2^{CE}} \end{aligned}$$

Note that in general, $\mu_2^{CE} > 0$. When $\alpha = 0$, the case of strict output targeting, $\mu_1^{CE} = \mu_2^{CE} = 0$. Then, from the optimal rule (10), $\pi_t = z^* + cy_{t-1}$. For $\alpha > 0$, π_t moves less than one-to-one with cy_{t-1} .

The Cautionary Policy

A cautionary policy recognizes parameter uncertainty $p > 0$. In a seminal paper, Brainard (1967) raised the issue of parameter uncertainty and optimal policy. Using a simple static model where there is no opportunity for learning, Brainard showed that optimal policy that allows for parameter uncertainty induces caution, in the sense that the policy instrument changes by a smaller amount compared to the that implied by the CEP.⁶ Within our dynamic model, the role of $p > 0$ can be seen by decomposing $E_{t-1}^c(z_t - z^*)^2$ into the square of the conditional mean $E_{t-1}^c(z_t - z^*)$ and the conditional variance F_{t-1} .

$$\begin{aligned} E_{t-1}^c(z_t - z^*)^2 &= (E_{t-1}^c z_t - z^*)^2 + F_{t-1} \\ &= (\pi_t + cy_{t-1} - z^*)^2 + py_{t-1}^2 + \sigma_u^2 \end{aligned} \quad (11)$$

⁶See also Tesfaselassie et al. (2006) and the references therein.

The expected loss due to output variability has an additional term, py_{t-1}^2 . Since $y_t = -\pi_t$, a lower value of $|\pi_t|$ reduces the conditional variance of z_{t+1} . Thus, parameter uncertainty matters for optimal monetary policy.

Cautionary Policy: Timing of events in period t

Stage 1	Stage 2	Stage 3	Stage 4
private sector sets π_t^e	central bank sets $y_t = E_{t-1}^c \pi_t^e$	central bank chooses $\pi_t = \pi(y_{t-1}, c, p; \cdot)$	u_t realizes, determining z_t

As before, the central bank continues to ignore the fact that current policy can affect future beliefs and by construction treats c and p as fixed parameters, implying that the only state variable from the central bank's point of view is y_t . The first order condition with respect to π_t will thus take the same form as (10). The difference is that now μ_2 is a function of p as well as c (see Appendix).

$$\pi_t^{CP} = \frac{z^* - \frac{1}{2}\delta\mu_1^{CP}}{1 + \alpha + \delta\mu_2^{CP}} + \frac{cy_{t-1}}{1 + \alpha + \delta\mu_2^{CP}} \quad (12)$$

where 'CP' denotes cautionary policy and

$$\begin{aligned} \mu_2^{CP} &= \frac{-1 - \alpha + \delta(c^2 + p) + \sqrt{4\delta(p + (c^2 + p)\alpha) + (1 + \alpha - \delta(c^2 + p))^2}}{2\delta} \\ \mu_1^{CP} &= \frac{2cz^*(\alpha + \delta\mu_2^{CP})}{1 + \alpha - c\delta + \delta\mu_2^{CP}} \end{aligned}$$

As $\mu_2^{CP} > \mu_2^{CE}$, we have $\pi_t^{CP} < \pi_t^{CE}$ implying that given its initial belief, the central bank disinflates by more under the cautionary policy. Moreover, the larger is p , the larger is the disinflation move.

The intuition behind a less accommodating policy under the cautionary policy lies in the additional loss from $p\pi_t^2$. Given $p > 0$, the central bank must choose π_t lower than π_t^{CE} so that the effect of p on future output variability is less magnified. In the limiting case where $p = 0$ the cautionary policy collapses to the certainty equivalence policy.

Unlike the case of certainty equivalence, $\lim_{\alpha \rightarrow 0} \mu_2^{CP} \neq 0$

$$\lim_{\alpha \rightarrow 0} \mu_2^{CP} = \frac{-1 + \delta(c^2 + p) + \sqrt{4\delta p + (1 - \delta(c^2 + p))^2}}{2\delta} \quad (13)$$

which is different from zero unless $p = 0$. This is an important result for the following reason. Suppose $\alpha = 0$. Under perfect knowledge, the optimal policy is to accommodate inflation expectations π_t^e , whatever the level may be. That is $\pi_t = z^* + \pi_t^e$. This rule also applies under certainty equivalence since $\pi_t = z^* + c\pi_{t-1}$, where the central bank accommodates its forecast of inflation expectations. By contrast, the cautionary policy does not fully accommodate the central bank's forecast of inflation expectations ($\pi_t = z^* + \frac{1}{1+\phi\delta}c\pi_{t-1} < z^* + c\pi_{t-1}$), where $\phi \equiv \lim_{\alpha \rightarrow 0} \mu_2^{CP}$.

Dynamically Optimal Policy

We now examine how disinflation policy is affected by learning considerations. The dynamic control problem is

$$\min_{\{\pi_\tau\}_{\tau=t}^{\infty}} E^c \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} L(z_\tau, \pi_\tau) | (y_{t-1}, c_{t-1}, p_{t-1}) \right] \quad (14)$$

subject to three constraints—the linear Phillips curve (5) and the two non-linear updating equations (7) and (8). Under fully optimal policy, there are three state variables: y_{t-1} , c_{t-1} and p_{t-1} .

Dynamically Optimal Policy: Timing of events in period t

Stage 1	Stage 2	Stage 3	Stage 4
private sector sets π_t^e	central bank sets $E_{t-1}^c \pi_t^e$	central bank chooses $\pi_t = \pi(y_{t-1}, c_{t-1}, p_{t-1})$	u_t realizes, determining z_t

The Bellman equation associated with the dynamic programming problem (14) is

$$\begin{aligned} V(c_{t-1}, p_{t-1}, y_{t-1}) &= \min_{\pi_t} \left\{ L(z_t, \pi_t) + \delta E_{t-1}^c V(c_t, p_t, y_t) \right\} \\ &= \min_{\pi_t} \left\{ \frac{1}{2} E_{t-1}^c (z_t - z^*)^2 + \frac{\alpha}{2} \pi_t^2 \right. \\ &\quad \left. + \delta \int V(c_t, p_t, y_t) f(z_t | c_{t-1}, p_{t-1}, y_{t-1}, \pi_t) dz_t \right\} \quad (15) \end{aligned}$$

where $E_{t-1}^c(z_t - z^*)^2$ is now decomposed as follows

$$E_{t-1}^c(z_t - z^*)^2 = (\pi_t - c_{t-1}y_{t-1} - z^*)^2 + p_{t-1}y_{t-1}^2 + \sigma_u^2 \quad (16)$$

The terms on the right hand side of (15) represent the tradeoff between control and estimation. The first two terms are current expected reward while implicit in the last term are two opposing components—one is the effect of π_t on $L(z_{t+1}, \pi_{t+1})$ (note that c_t and p_t depend on π_{t-1} but not on π_t) and the other is the expected improvement in payoffs from $t + 2$ onwards due to better information about the unknown parameter (via the effect of π_t on p_{t+1}). The first component implies that, given $p_t > 0$, as with the CP, the DOP reduces the conditional variance of z_{t+1} by decreasing the level of π_t . But $\pi_{t-1} > 0$ means that $p_t < p_{t-1}$, which reduces expected losses and π_t does not have to decrease by as much as it does in the CP. Therefore, this channel leads to gradual disinflation compared to the CP and the gradualist policy is enhanced the larger the initial level of inflation, which helps reduce parameter uncertainty considerably.

Unlike the CEP and CP, the DOP is a non-linear function of the state variables and can be solved for only numerically. As shown by Easley and Kiefer (1988) and Kiefer and Nyarko (1989) an optimal feedback rule exists and the value function is continuous and satisfies the Bellman equation. Policy and value functions can be obtained using an iterative algorithm based on the Bellman equation and starting with an initial guess.

We solve for the optimal policy under learning using numerical dynamic programming (see e.g. Wieland (2000a)). Then we compare disinflation policy under the DOP with those of CEP and CP. First we show results for a baseline parameters where $\alpha = 0.5$, $\sigma_u^2 = 1$, $\delta = 0.95$ and $z^* = 0.25$.

Figure 1 shows that for various combinations of initial beliefs about the mean and variance of the unknown parameter, the DOP is in general more accommodative to inflation expectations than the CP but less accommodative than the CEP. The initial position of inflation determines whether DOP stays closer to the CP or to the CEP. For large deviations of initial inflation from zero the DOP is similar to the CEP. An intuitive explanation for this result is that, from the updating equations, the central bank recognizes that the larger the deviation of π_{t-1} from zero, the

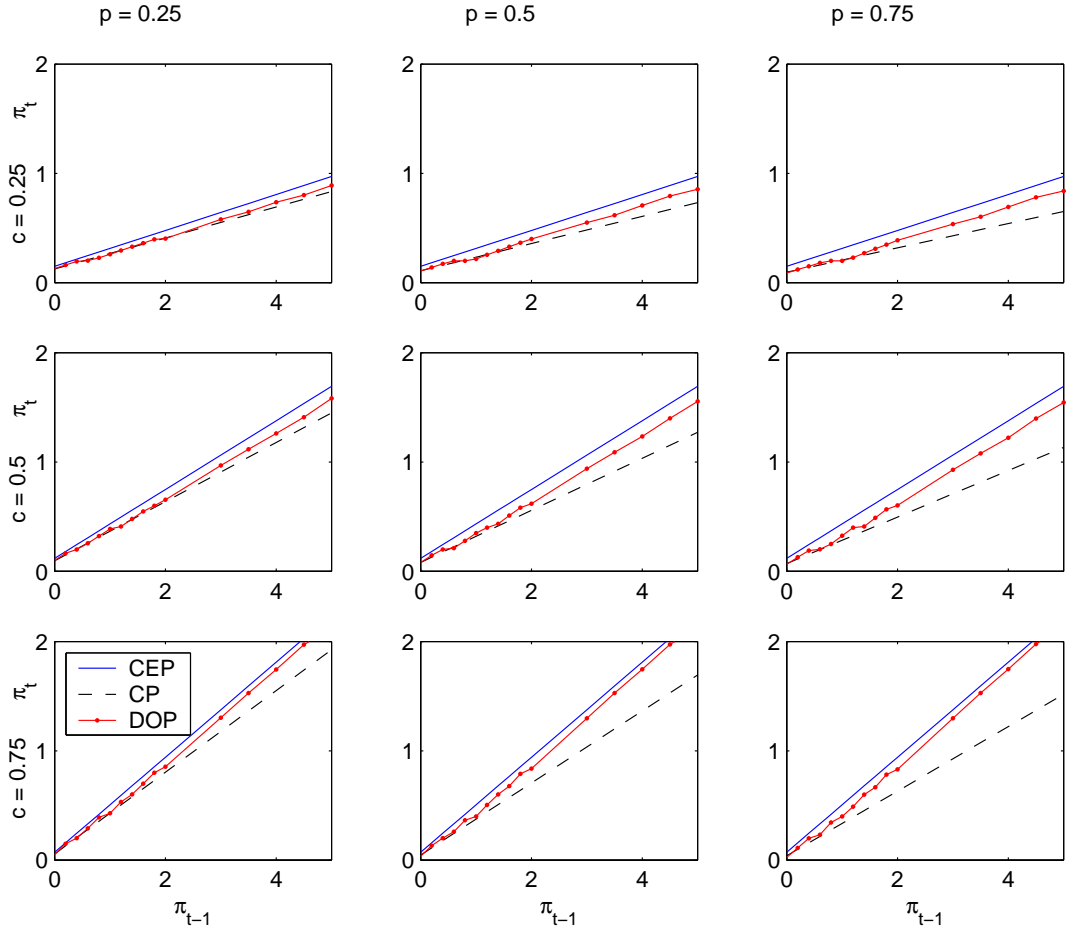


Figure 1: CEP, CP and DOP for baseline parameters ($\alpha = 0.5, \sigma_u^2 = 1, \delta = 0.95$ and $z^* = 0.25$).

smaller p_t , which in turn reduces the conditional variance of z_{t+1} . In anticipation of this, the central bank does not have to disinflate as much as the CP would imply. Therefore, the optimal policy remains closer to the CEP. On the other hand, when π_{t-1} is small, p_t remains close to p_{t-1} . Expecting only a marginal reduction in the degree of uncertainty in period $t + 1$, the central bank gives more weight to reducing expected future losses from the immediate future relative to reducing estimation errors in the more distant future. Thus, it disinflates more aggressively, moving towards the CP. However, this effect is weaker the larger is the initial parameter uncertainty (i.e., the larger p_{t-1}).

Sensitivity Analysis

Below we show results when the variance of the exogenous shocks to output and the discount factor take different values than the baseline values.

Smaller variance of shocks

Note that, the CEP and the CP are independent of the variance of the output shock. Figure 2 compares the DOP for two levels of the variance of the shock to output gap ($\sigma_u^2 = 1$ and $\sigma_u^2 = 0.1$). In that case, the DO policy is closer to the CEP when $\sigma_u^2 = 0.1$ than when $\sigma_u^2 = 1$.

The intuition for this effect is that the output gap z_t is more stable under $\sigma_u^2 = 0.1$ than under $\sigma_u^2 = 1$. Given y_{t-1} , the updating equation for c_t implies that the forecast error $z_t - \pi_t - c_{t-1}y_{t-1}$ is more informative about the unknown parameter the smaller the variance of z_t due to exogenous shocks and the larger the variance of z_t due to estimation errors. This effect is also apparent from the updating equation for p_t , which is positively related to σ_u^2 .⁷

Shorter Policy Horizon

Changes in δ affect all types of policies.⁸ For instance, the smaller δ , the more heavily future losses are discounted and the shorter the central bank's policy horizon. In that case, the effect of parameter uncertainty on future expected losses is less of a concern to the central bank. Thus, the DOP and the CP will move towards the CEP, implying that all policies call for a more gradual disinflation process.

Figure 3 shows the effect of a decrease in the discount factor on the DOP relative to the CP (the degree of gradualism of the DOP relative to the CP). The DOP induces less relative gradualism at the lower value of the discount factor if the initial level of inflation is large. However, the differences between the DOP and the CP seem to disappear at low to moderate rates of initial inflation, and at small values of the

⁷The change in the DOP is more muted when the weight on inflation increases.

⁸When $\alpha = 0$, the CEP collapse to a static optimal policy, thus independent of δ .

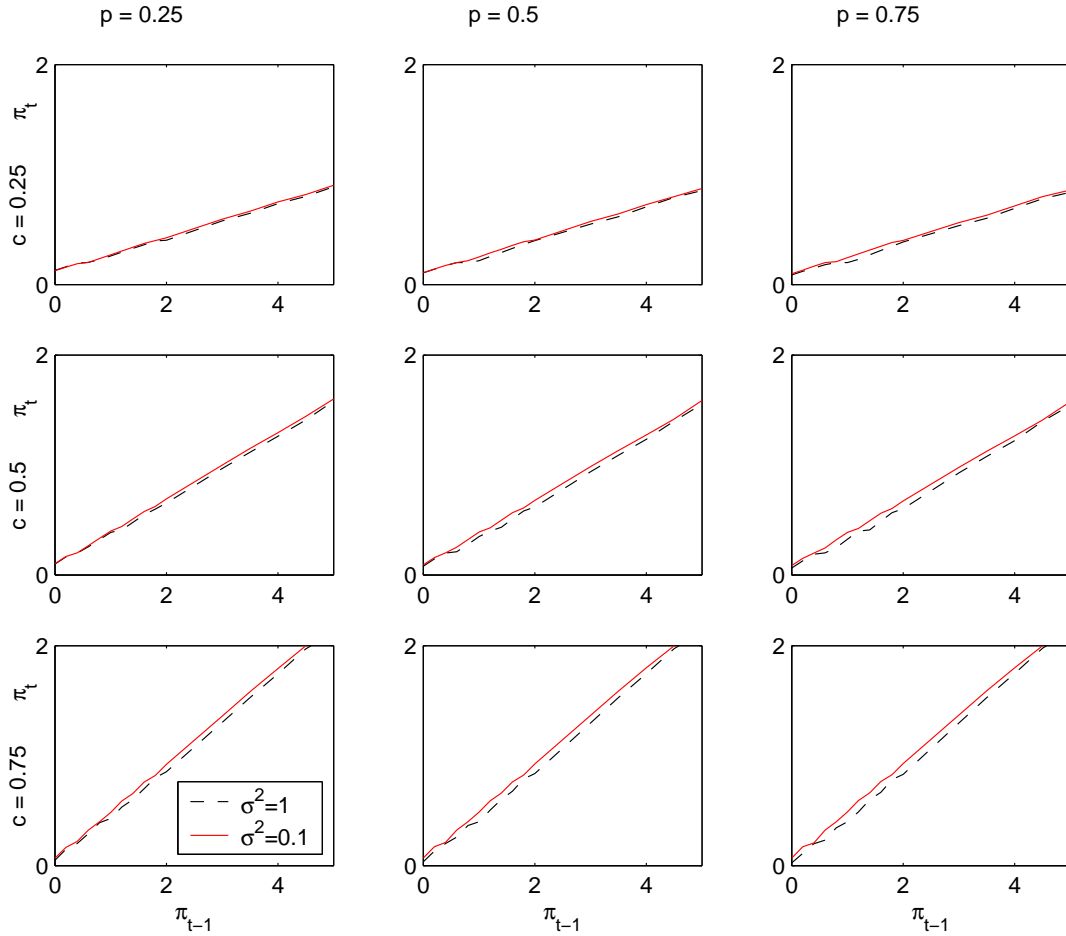


Figure 2: DOP under alternative values of σ_u^2 ($\sigma_u^2 = 1$ vs $\sigma_u^2 = 0.1$).

parameter estimates.

4 Speed of Learning and Convergence

We now turn to the dynamics of inflation and central bank belief and their convergence in the limit. The question is whether in the limit inflation approaches its new target under alternative policies. Could the central bank end up having a wrong limit belief about γ , which would lead to incorrect limit policy, whereby inflation stabilizes at a level different from its target?

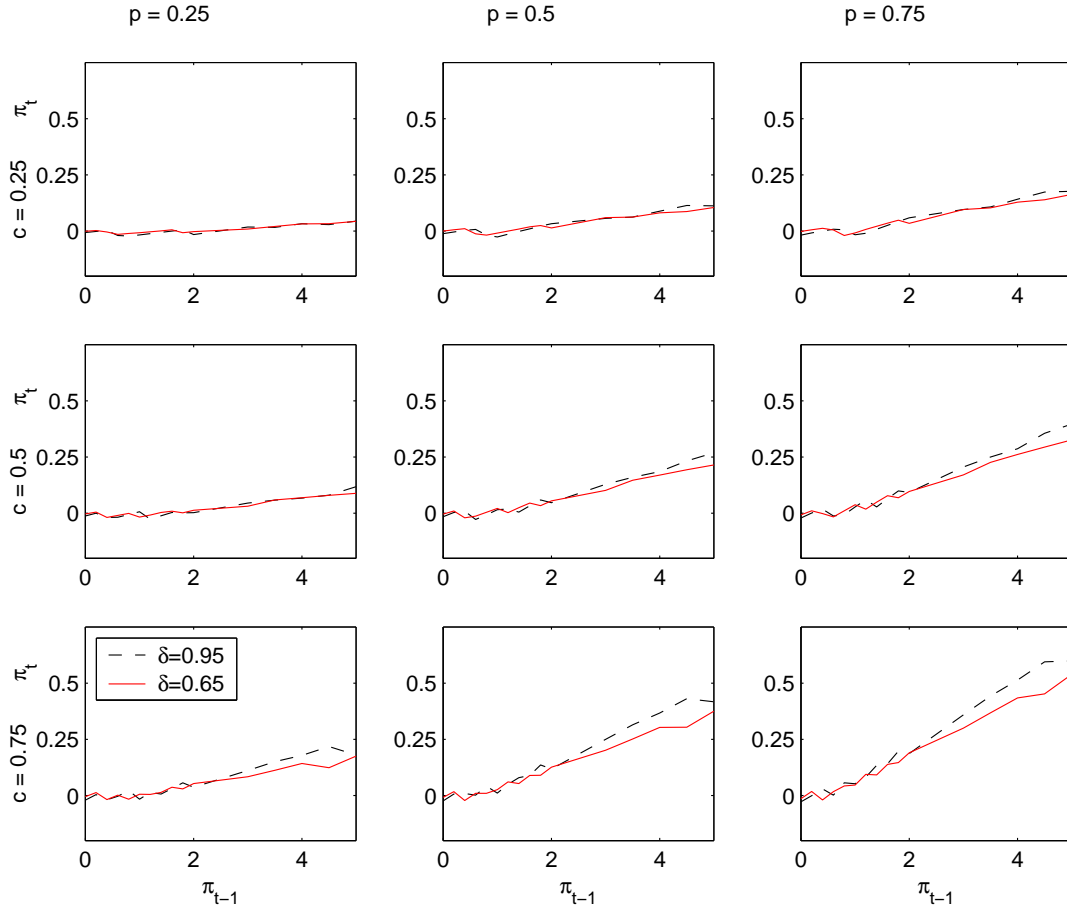


Figure 3: Degree of activism (DOP relative to CP) ($\delta = 0.95$ vs $\delta = 0.65$).

Starting from period 0 the sequence of estimation, control and updating is $(c_0, p_0) \rightarrow (\pi_1, z_1) \rightarrow (c_1, p_1) \rightarrow (\pi_2, z_2)$ and so on. As the dynamics of (π_t, z_t) and (c_t, p_t) are interrelated, even if we start with the same priors (c_0, p_0) , the dynamics of estimation and control will depend on the type of policy followed by the central bank.

Note that as in King (1996), the central bank has perfect control over inflation. That means, absent exogenous shocks to inflation, there is a possibility of incomplete learning about inflation expectations if inflation is stabilized too quickly. As it induces low variations in π_t the CP is most susceptible to the danger of incomplete learning. By adopting a slower speed of disinflation, the CEP and the DOP increase estimation precision and thereby improving future control of the economy.

Of course in reality, inflation is subject to shocks outside the control of the central bank. When there is an additive control error, actual inflation is the sum of intended monetary policy π_t^I and an exogenous control error ν_t , that is, $\pi_t = \pi_t^I + \nu_t$. The econometric model is still given by (1). Actual inflation is now stochastic even if π_t^I is fixed by the central bank. Thus, in the limit c_t converges to γ . Even if these additive control error has a very small variance, inflation will never settle down with time, implying that in the limit, beliefs converge to the true parameter. As the variance of the additive control error increases, the central bank focuses more on current control of the economy and less on future estimation. It follows that the central bank has an incentive to speed up the disinflation process by reducing π_t^I more rapidly. Given policy, learning tends to be slow as the variance of the additive control error diminishes. The implication is that, if the additive control error is insignificant and the central bank improves its control of inflation, the speed of learning will depend more on its disinflation policy.

One possible extension of the analysis is to let γ be time-dependent, for example a random walk $\gamma_t = \gamma_{t-1} + \eta_t$ as in Beck and Wieland (2002) or an autoregressive process $\gamma_t = \rho\gamma_{t-1} + \eta_t$, $0 < \rho < 1$ as in Balvers and Cosimano (1994). It is easy to conjecture that in this case, learning will be perpetual, as the underlying parameter changes all the time. This may reduce the incentives for learning, and move the DOP towards the CP, implying larger disinflation than the case of fixed unknown parameter.

5 Concluding Remarks

The paper analyzes disinflation policy when the central bank has imperfect information about private sector inflation expectations, thus extending King (1996), which supposes perfect observability of inflation expectations. The central bank learns about inflation expectations from past economic outcomes, which are in part the result of past policy decisions. Due to the dependence of inflation expectations on past policy decisions, the problem facing the central bank is one of parameter uncertainty, that is, uncertainty about the effect of past policy on the current level of output. Formally, the dynamic control problem differs from other studies on

learning and control, where the assumed uncertainty is about the effect of current actions on current economic outcomes, and lagged control variable is absent from the dynamic process.

We compare three policy scenarios under which disinflation policy may proceed—certainty equivalence policy (CEP), cautionary policy (CP) and dynamically optimal policy (DOP). The CEP and CP represent passive learning but while the CEP ignores parameter uncertainty the CP policy assumes that current uncertainty about inflation expectations will remain unchanged in the future. Given the state of the economy, the DOP disinflates by more than the CEP but by less than the CP. A novel result is that, unlike the case of uncertainty about current policy effect, our result holds irrespective of the initial state of the economy (characterized by past level of inflation).

It turns out that, given the central bank's belief about inflation expectations, the DOP moves closer to the CEP when past inflation is high. By contrast, when past inflation is low, the DOP stays close to the CP implying more caution. In general, the danger with the CP is that if inflation drops sharply and stabilizes too soon, the central bank might fail to learn about inflation expectations, leading to poor policy performance in the distant future. By taking into account the effect of policy on inflation expectations, the DOP and the CEP are less prone to the danger of incomplete learning.

References

- Balvers, R. and Cosimano, T. (1994), Inflation Variability and Gradualist Monetary Policy, *The Review of Economic Studies* **61**, 721–738.
- Beck, G. and Wieland, V. (2002), Learning and Control in a Changing Economic Environment, *Journal of Economic Dynamics and Control* **26**, 1359–1377.
- Bertocchi, G. and Spagat, M. (1993), Learning, Experimentation and Monetary Policy, *Journal of Monetary Economics* **32**, 169–183.
- Bomfim, A. and Rudebusch, G. (2000), Opportunistic and Deliberate Disinflation Under Imperfect Credibility, *Journal of Money, Credit, and Banking* **32**, 707–721.
- Brainard, W. (1967), Uncertainty and the Effectiveness of Policy, *American Economic Review* **57**, 411–425.
- Cogley, T. and Colacito, R. and Sargent, T. (2005), Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas, 2005 Meeting Papers, *Society for Economic Dynamics*.
- Easley, D. and Kiefer, N. (1988), Controlling a Stochastic Process with Unknown Parameters, *Econometrica* **56**, 1045–1064.
- Ellison, M. (2006), The Learning Cost of Interest Rate Reversals, *Journal of Monetary Economics* **53**, 1895–1907.
- Ellison, M. and Valla, N. (2001), Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions, *Journal of Monetary Economics* **48**, 153–171.
- Evans, G. and Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press, NJ.
- Hoerberichts, M. and Schaling, E. (2000), Why Speed Doesn't Kill: Learning to Believe in Disinflation, *Research Paper Department of Economics RAU* no. 0009.
- Kiefer, N. and Nyarko, Y. (1989), Optimal Control of an Unknown Linear Process with Learning, *International Economic Review* **30**, 571–586.

- King, M. (1996), How Should Central Banks Reduce Inflation? - Conceptual Issues, *Economic Review Federal Reserve Bank of Kansas City* **81**, 25–52.
- Marcet, A. and Sargent, T. (1988), The Fate of Systems with Adaptive Expectations, *American Economic Review* **78**, 168–171.
- Prescott, C. (1972), The Multi-period Control Problem with Uncertainty, *Econometrica* **40**, 1043–1058.
- Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, NJ.
- Sargent, T. (1987), *Dynamic Macroeconomic Theory*, Harvard University Press, MA.
- Schaling, E. (2003), Learning, Inflation Expectations and Optimal Monetary Policy, *Bank of Finland Discussion Paper*, no. 20.
- Svensson, L. and Williams, N. (2007), Bayesian and Adaptive Optimal Policy under Model Uncertainty, *CFS Working Paper Series*, no. 11.
- Tesfaselassie, M. (2005), *Communication, Learning and Optimal Monetary Policy*, CentER Dissertation Series, No.155, Tilburg University, Tilburg.
- Tesfaselassie, M., Schaling, E. and Eijffinger, S. (2006), Learning About the Term Structure and Optimal Rules for Inflation Targeting, *CEPR Discussion Paper*, no. 5896.
- Wieland, V. (2000a), Monetary Policy, Parameter Uncertainty and Optimal Learning, *Journal of Monetary Economics* **46**, 199–228.
- Wieland, V. (2000b), Learning by Doing and the Value of Optimal Experimentation, *Journal of Economic Dynamics and Control* **24**, 501–534.
- Yetman, J. (2003), Probing Potential Output: Monetary Policy, Credibility and Optimal Learning under Uncertainty, *Journal of Macroeconomics* **25**, 311–330.

Appendix: Derivation of CEP and CP

In this appendix we derive the optimal policy for the case of passive learning. When choosing current policy in period t , the central bank assumes that the initial belief (c_{t-1}, p_{t-1}) will remain fixed for all future periods. Consequently, the non-linear updating equations drop out of the set of constraints of the optimization problem.⁹

There are two subcases under passive learning—certainty equivalence policy and CP. Under the cautionary case $F_{t-1} = py_{t-1}^2 + \sigma_u^2$. On the other hand, under the case of certainty equivalence, $p = 0 \Rightarrow py_{t-1}^2 = 0$, and so F_{t-1} is perceived to be independent of y_{t-1} . Thus, the certainty equivalence policy is a limiting case of the CP.

Derivation of the CEP

Under certainty-equivalence the central bank ignores parameter uncertainty, being fully confident about its prior $c_{t-1} = c$. Its current belief is thus characterized as $(c_{t-1}, p_{t-1}) = (c, 0)$. The central bank minimizes (9) subject to (5) and (6). The problem is linear-quadratic, so that the optimal level of π_t is similar to that under perfect knowledge. The certainty-equivalence rule simply replaces γ with its estimate c . We can rewrite the above minimization problem using recursive dynamic programming and then use the standard 'guess and verify' method on the value function.¹⁰ We can write the Bellman equation associated with the minimization of (9) as follows¹¹

$$V(y_{t-1}) = \min_{\pi_t} E_{t-1}^c [L(z_t, \pi_t) + \delta V(y_t)] \quad (\text{A1})$$

subject to (5). Because of the linear-quadratic form of the minimization problem,

⁹Of course, when next period arrives, the bank updates its belief but then expect it to remain fixed from that period on.

¹⁰See for e.g., Tesfaselassie (2005), Chapter 5 and the references therein.

¹¹Note that the value function in the Bellman equation does not have time subscript. This is because in infinite horizon problems, we are interested only in the unique time invariant value function, V , and associated unique, stationary policy rule, that result from repeated iterations on the Bellman equation starting from any bounded continuous V_0 (e.g. $V_0 = 0$). Convergence of the value function is guaranteed due to *contraction mapping theorem* (see e.g. Sargent (1987)).

the value function will be quadratic in the state y_{t-1} .

$$V(y_{t-1}) = \mu_0 + \frac{1}{2}\mu_2 y_{t-1}^2 \quad (\text{A2})$$

where the two coefficients remain to be determined. If (A2) is correct, it follows that

$$E_{t-1}^c V(y_t) = \mu_0 + \frac{1}{2}\mu_2 y_t^2 \quad (\text{A3})$$

or using (6)

$$E_{t-1}^c V(y_t) = \mu_0 + \frac{1}{2}\mu_2 (-\pi_t)^2 \quad (\text{A4})$$

Next, substitute (A4) into (A1) and derive the first order condition with respect to π_t

$$\pi_t = \frac{1}{1 + \alpha + \mu_2 \delta} c y_{t-1} \quad (\text{A5})$$

To identify the value of μ_2 , substitute (A5) into the loss function and match the resulting coefficient of y_{t-1}^2 with the conjectured loss function (A2). A unique solution for μ_2 , such that $0 < \mu_2 < 1$, is given by equation (10) of the main text.

Derivation of the CP

Under the CP, $p > 0$. As before the only state variable from the central bank's point of view is y_{t-1} . The conjecture for the value function is given by (A2). The first order condition with respect to π_t will also take the same form as (A5). Following the steps analogous to the derivation of the CEP, we match the coefficients and arrive at the solution given by equation (12) of the main text. The certainty equivalence case arises if $p = 0$, that is, if we disregard parameter uncertainty.