Learning by Disinflating*

Alina Barnett† — Bank of England

Martin Ellison‡ — University of Oxford

March 31, 2011

Abstract

Disinflationary episodes are a valuable source of information for economic agents trying to learn about the economy. This paper is especially interested in how a policymaker can themselves learn by disinflating. The approach differs from the existing literature, which typically focuses on the learning of private agents during a disinflation. We build a model where both the policymaker and private agents learn, and ask what happens if the policymaker has to disinflate to satisfy a new central bank mandate specifying greater emphasis on inflation stabilisation. In this case, our results show that inflation may fall dramatically before it gradually rises to its new long run level. The potential for inflation to undershoot its long run level during a disinflationary episode suggests that caution should be exercised when assessing the success of any change in the policymaker's mandate.

JEL classification: D83, E52, E58

Keywords: Disinflation, Escape Dynamics, Learning, Monetary Policy

*We thank Bill Branch, Ken Kasa, Andrew Scott and Jouko Vilmunen for helpful comments and suggestions. Martin Ellison acknowledges support from an ESRC Research Fellowship, "Improving Monetary Policy for the 21st Century" (RES-000-27-0126). The views expressed in this paper do not necessarily reflect those of the Bank of England.

†Monetary Assessment and Strategy Division, Bank of England, Threadneedle Street, London, EC2R 8AH, United Kingdom, Tel: +44 20 7601 4444, E-mail: Alina.Barnett@bankofengland.co.uk

‡Department of Economics, University of Oxford, Manor Road Building, Oxford, OX1 2UQ, United Kingdom, Tel: +44 1865 284249, E-mail: martin.ellison@economics.ox.ac.uk

1 Introduction

The US disinflation of the late 1970s and early 1980s has been instrumental in shaping modern views of the monetary transmission mechanism and the efficacy of monetary policy. It is now widely accepted that disinflations create downturns, and that the standard New Keynesian model needs amending to avoid disinflations being expansionary. The need for amendments was first noted by Ball (1994) in his famous "disinflationary booms" paper, and led to models by Ball (1995) and Erceg and Levine (2003) in which private agents only slowly learn the objectives of a policymaker with imperfect credibility. Other contributions to this literature include Ireland (1995), Goodfriend and King (2005), Nicolae and Nolan (2006), Evans and Nicolae (2010) and Cogley, Matthes and Sbordone (2011).

The focus of the existing literature on disinflations is on the role of private agents learning. In this paper we argue that it is equally important to examine what the policymaker themselves learns during a disinflation. In our view of the world, policymakers are continuously learning how the economy operates and striving to improve their policies. Their experiences during disinflations are then likely to provide crucial information on key economic concepts such as the sacrifice ratio and the long and variable lags of the monetary transmission mechanism. Put simply, we believe that a serious model of disinflations needs to encompass learning by both private agents and the policymaker. Only a few papers such as Sargent, Williams and Zha (2006) and Primiceri (2006) share our belief that learning by policymakers is important, and none of these papers specifically looks at the policymaker learning by disinflating.

To make our point, we construct a model in which both the policymaker and private agents have non-trivial learning problems. The policymaker is assumed to learn while solving the joint estimation and policy problem described by Sargent (1999) in *The Conquest of American Inflation*. Private agents learn according to a standard adaptive learning algorithm of the type closely associated with Evans and Honkapohja (2001). Our closest antecedent is the model used by Cho and Kasa (2008) to study currency crises, although aspects of our analysis also overlap with Cho, Williams and Sargent (2002) and Williams (2004). To induce learning by disinflating we analyse what happens if politicians decide to appoint a Rogoff (1985) conservative central banker who places greater emphasis on inflation stabilisation. In other words, we examine how the economy reacts to a change in the policymaker's mandate when both private agents and the policymaker are learning. To the best of our knowledge, we are the first to explicitly examine how a policymaker learns by disinflating.

The short run effects of appointing a conservative central banker are potentially dramatic in our model. If the policymaker is learning then inflation may fall rapidly to a very low level shortly after the conservative central banker is appointed, before eventually rising back to a new long run equilibrium. Such undershooting of inflation is potentially hazardous if it gives politicians a false sense of how successful changing the central bank's mandate has been. The idea that external events can cause sudden and dramatic shifts in the macroeconomy has also been explored by McGough (2006) in the context of shocks to the natural rate of output that cause the economy to escape from its self confirming equilibrium, a paper that shares some techniques with ours. The long run effects of appointing a conservative central banker in our model are lower and less volatile inflation, as predicted in the working paper of Gerali and Lippi (2002).

The paper is structured as follows. Section 2 presents our model of learning by a policymaker and private agents, and models how politicians change the central bank's mandate. The long and short run effects of appointing a conservative central banker are derived and discussed in Section 3 and Section 4 respectively. Section 5 presents simulation results that quantify the likely magnitude and duration of any inflation undershooting. A final Section 6 concludes.

2 Model

The fundament of our model is the Sargent (1999) framework in which the policymaker learns through solving a joint estimation and policy problem. To it we add adaptive learning by private agents à la Evans and Honkapohja (2001), and the possibility of politicians changing the mandate of the policymaker by appointing a Rogoff (1985) conservative central banker. The exposition of our model begins with an explanation of the true structure of the economy that links policy and private expectations to macroeconomic outcomes, and proceeds by describing the estimation and optimisation problem of the policymaker. How we model the appointment of a conservative central banker is next, followed by a definition of the learning of private agents to complete the model.

2.1 Structure of the economy

Macroeconomic outcomes are determined by the inflation expectations of private agents π_t^e and the level of inflation intended by the policymaker x_t. They are subject to real and nominal shocks. In equation (1), unemployment U_t is a function of the natural rate of unemployment

U^*, inflation that is unexpected by private agents, $\pi_t - \pi_t^e$, and a real shock $\nu_{1t} \sim \mathcal{N}(0, \sigma_1^2)$. In equation (2), inflation π_t depends on the level of inflation intended by the policymaker x_t and a nominal shock $\nu_{2t} \sim \mathcal{N}(0, \sigma_2^2)$.[1]

$$U_t = U^* - \theta(\pi_t - \pi_t^e) + \nu_{1t} \tag{1}$$

$$\pi_t = x_t + \nu_{2t} \tag{2}$$

The structure of the economy implies that unemployment deviates from its natural rate whenever there is a gap between the inflation intended by the policymaker and the inflation expected by private agents. Such a gap arises naturally in our model as the policymaker and private agents are both learning. It means that disinflations will be costly.

2.2 Learning by the policymaker

The estimation problem of the policymaker is taken from Sargent (1999). Accordingly, the policymaker understands that unemployment and inflation are interrelated, but does not know the true structure of the economy (1). Instead, they believe there is a trade-off between unemployment and inflation that may be shifting over time. The policymaker learns each period by estimating the relationship between unemployment and inflation:

$$U_t = \gamma_0 + \gamma_1 \pi_t + \xi_t. \tag{3}$$

using discounted least squares. This is equivalent to the policymaker employing a constant gain learning algorithm to form and update its beliefs $\gamma_t = (\gamma_{0t}, \gamma_{1t})'$ about the parameters in (3), with the constant gain reflecting the policymaker's concern that there may be shifts in the trade off between unemployment and inflation. The beliefs of the policymaker are updated according to:

$$\gamma_{t+1} = \gamma_t + a_g \left(R_t^{-1} \begin{pmatrix} 1 \\ \pi_t \end{pmatrix} (U_t - \gamma_{0t} - \gamma_{1t}\pi_t) \right), \tag{4}$$

$$R_{t+1} = R_t + a_g \left(\begin{pmatrix} 1 \\ \pi_t \end{pmatrix} \begin{pmatrix} 1 & \pi_t \end{pmatrix} - R_t \right), \tag{5}$$

where a_g is the constant gain or exponential rate at which the policymaker discounts past data. R_t is a 2×2 matrix that describes the precision of the current coefficient estimates.

[1] ν_{2t} can also be interpreted as a control error made by the policymaker.

2.3 The policy problem

The optimisation problem of the policymaker follows Sargent (1999), and is to minimise a weighted average of the quadratic deviations of unemployment and inflation from their target levels of zero. The weight β_t assigned to inflation stabilisation is mandated by politicians. The policy problem is constrained by the policymaker's perceived structure of the economy (3), the policymaker's current estimates $(\gamma_{0t}\ \gamma_{1t})'$ of the parameters in it, and tacit acknowledgement that inflation is not completely under the control of the policymaker:

$$\min_{x_t} E_t \sum_{s=0}^{\infty} \delta^s (U_{t+s}^2 + \beta_t \pi_{t+s}^2),$$

$$\text{s.t.}$$

$$U_{t+s} = \gamma_{0t} + \gamma_{1t}\pi_{t+s} + \xi_{t+s},$$

$$\pi_{t+s} = x_{t+s} + v_{2t+s}.$$

The fully optimal policy maximises the policymaker's objective, subject not only to the beliefs of the policymaker and the current mandate, but also subject to the way beliefs are updated (4) and any changes expected in future mandates. Including these constraints would make the policy problem highly nonlinear, so we follow Sargent (1999) and invoke anticipated utility maximisation as the policymaker's decision criterion. The policymaker then treats its beliefs and mandate as fixed at their current values, now and into the future, even though it knows they may subsequently change. Kreps (1998) argues that maximising anticipated utility in this way is a simple and robust strategy for a policymaker facing a complex dynamic problem such as ours.[2] The solution of the policy problem determines intended inflation as a function of the current beliefs and mandate of the policymaker:

$$x_t = -\frac{\gamma_{0t}\gamma_{1t}}{\beta_t + \gamma_{1t}^2}. \tag{6}$$

The policymaker sets intended inflation (8) at the beginning of the period, and updates their beliefs through (4) at the end of the period. Note that certainty equivalence holds with respect to the nominal shock ν_{2t} in the policy but not the learning problem.

[2] The numerical simulations of Cogley, Colacito and Sargent (2007) also suggest that maximising anticipated utility gives a good approximation to the fully optimal policy.

2.4 Changing the policymaker's mandate

To induce disinflations in our model, we allow politicians to change the mandate of the policymaker by appointing a Rogoff (1985) conservative central banker. This is the central innovation of our paper. We allow for time variation in β_t, the weight on inflation stabilisation in the policymaker's mandate, whereas the existing literature follows Sargent (1999) and keeps β_t constant. We assume that politicians decide on the timing, severity and abruptness of the change in mandate by choosing the parameters β_L, β_H, s and ϕ in the function:

$$\beta_t = \beta_L + \frac{\beta_H - \beta_L}{1 + e^{-\phi(t-s)}}, \tag{7}$$

so the change in mandate is centred on time s and involves an increase in the weight assigned to inflation stabilisation from β_L to β_H. The parameter ϕ governs the abruptness of the change in mandate, with higher values of ϕ corresponding to more abrupt changes. As $\phi \to \infty$ the change becomes instantaneous. Figure 1 shows an example of how β_t evolves when $\beta_L = 1, \beta_H = 1.5, s = 4$ and $\phi = 10$.

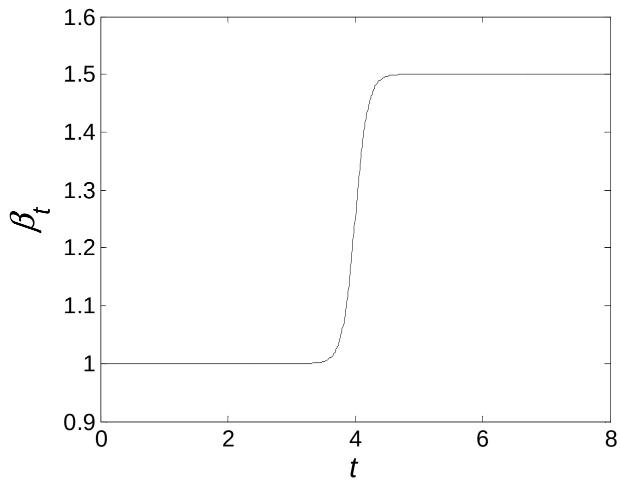

Figure 1: Example of change in the policymaker's mandate

2.5 Learning by private agents

We assume that private agents learn using a standard adaptive learning algorithm. This deviates from Sargent's (1999) assumption of rational expectations, and is motivated by the empirical evidence that disinflations create downturns and unemployment. If private agents have rational expectations in our model then unemployment (1) is i.i.d. and disinflations are costless. Cho and Kasa (2008) assume adaptive learning by private agents to similarly generate output losses after a currency crisis. Following their lead, private agents in our model learn according to:

$$\pi^e_{t+1} = \pi^e_t + a_p(\pi_t - \pi^e_t), \tag{8}$$

where a_p is the constant gain at which they discount past data. Private agents discount past data in implicit acknowledgement that the structure of the economy may be changing over time. The higher the value of a_p the more they believe there are changes, so the more they weight current data and discount past data.

3 Long run dynamics

The dynamics of the model are described by the structure of the economy (1)-(2), the learning of the policymaker (4)-(5), the solution of the policy problem (6), the policymaker's mandate (7) and the learning of private agents (8). This is a system of ordinary difference equations in unemployment, inflation, the beliefs of the policymaker, the mandate of the policymaker and expected inflation. We characterise its dynamics by applying stochastic approximation techniques taken from Cho, Williams and Sargent (2002). To do so, the equations for the learning of the policymaker are written as:

$$\frac{\gamma_{t+1} - \gamma_t}{a_g} = R_t^{-1} \begin{pmatrix} 1 \\ x_t + v_{2t} \end{pmatrix} \left(U^* - \theta(x_t + v_{2t} - \pi^e_t) + \nu_{1t} - \gamma_{0t} - \gamma_{1t}(x_t + v_{2t})\right), \tag{9}$$

$$\frac{R_{t+1} - R_t}{a_g} = \begin{pmatrix} 1 \\ x_t + v_{2t} \end{pmatrix} \begin{pmatrix} 1 & x_t + v_{2t} \end{pmatrix} - R_t, \tag{10}$$

which can be interpreted as a discrete-time approximation of a continuous-time process perturbed by shocks v_{1t} and v_{2t}. Taking the limit as $a_g \to 0$, the approximation error tends to zero and a weak law of large numbers ensures that the stochastic element becomes negligible. The learning of the policymaker can therefore be represented by a system of deterministic ordinary differential equations:

$$\dot{\gamma}_t = R_t^{-1} \begin{pmatrix} U^* - \theta(x_t - \pi^e_t) - \gamma_{0t} - \gamma_{1t}x_t \\ (U^* - \theta(x_t - \pi^e_t) - \gamma_{0t} - \gamma_{1t}x_t)x_t - (\theta + \gamma_{1t})\sigma_2^2 \end{pmatrix}, \tag{11}$$

$$\dot{R}_t = \begin{pmatrix} 1 & x_t \\ x_t & x_t^2 + \sigma_2^2 \end{pmatrix} - R_t, \tag{12}$$

known as the *mean dynamics.* They describe the expected evolution of policymaker's beliefs. Similar calculations for the limiting behaviour of learning by private agents (10) as $a_p \to 0$

define the mean dynamics of inflation expectations:

$$\dot{\pi}_t^e = x_t - \pi_t^e. \tag{13}$$

The mean dynamics of the beliefs of the policymaker (11)-(12) and the inflation expectations of private agents (13) together form a system of ordinary differential equations in γ_t, R_t and π_t^e. Their long-run outcome is the fixed point at which $\dot{\gamma}_t = \dot{R}_t = \dot{\pi}_t^e = 0$:

$$\begin{aligned}
\bar{\gamma}_0 &= U^* \left(1 + \frac{\theta^2}{\beta_t}\right), \\
\bar{\gamma}_1 &= -\theta, \\
\bar{U} &= U^*, \\
\bar{\pi} &= \bar{\pi}^e = \bar{x} = \frac{\theta U^*}{\beta_t},
\end{aligned}$$

a self confirming equilibrium (SCE) in the terminology of Sargent (1999). It is unique in our model.[3] The self confirming equilibrium has the same long run inflation bias that appears in Nash equilibrium in our model if the policymaker and private agents are assumed to have rational expectations. It follows that long run inflation falls from $\theta U^*/\beta_L$ to $\theta U^*/\beta_H$ after politicians change the policymaker's mandate to place greater weight on inflation stabilisation. Our model therefore exhibits the same long run reduction in inflation bias that prompted Rogoff (1985) to argue for delegating policy to a conservative central banker.

4 Short run dynamics

The fall in long run inflation after a change in the policymaker's mandate occurs as beliefs adjust to new levels consistent with the self confirming equilibrium. For the policymaker, beliefs change in the long run from $\left(U^* \left(1 + \theta^2 \beta_L^{-1}\right), -\theta\right)'$ to $\left(U^* \left(1 + \theta^2 \beta_H^{-1}\right), -\theta\right)'$. For private agents, expected inflation falls in the long run from $\theta^2 U^* \beta_L^{-1}$ to $\theta^2 U^* \beta_H^{-1}$. The adjustment in beliefs comes about as the policymaker and private agents learn the changing self confirming equilibrium, with the stability result in footnote 3 guaranteeing that beliefs eventually converge to their new long run values. What happens while beliefs are adjusting determines the short run dynamics of the model. The focus of our paper on learning by disinflating means we are particularly interested in whether adjustment is a smooth monotonic process or whether

[3] The eigenvalues of the Jacobian of the system of equations (9)-(11) all have negative real parts, so the self confirming equilibrium is stable.

something more interesting happens before beliefs converge to their new values. To proceed, we ask if beliefs under learning are likely to deviate substantially from the transition path implied by the mean dynamics of the model. In the terminology of Sargent (1999), any such deviations are known as escape episodes.

The formal analysis of escape episodes in economics was pioneered by Williams (2004), who shows how to derive the most likely path that beliefs take if they deviate substantially from their mean dynamics. The method involves solving an optimal control problem to identify the most likely series of perturbations that cause beliefs to escape a neighbourhood around the self-confirming equilibrium. The metric for "likely" is the likelihood function of the shocks needed to perturb beliefs away from the path implied by their mean dynamics. The solution of the optimal control problem is the path of least resistance for beliefs to escape, known as the dominant escape path. Mathematically, it solves the following control problem:

$$\bar{\Psi} = \inf_{\dot{v}} \int_0^t \dot{v}(\varphi)' Q(\Gamma(\varphi), R(\varphi), \beta(\varphi))^{-1} \dot{v}(\varphi) d(\varphi)$$

$$\text{s.t.}$$

$$\dot{\Gamma} = \bar{g}(\Gamma, R, \beta) + \dot{v},$$

$$\dot{R} = \bar{M}(\Gamma, \beta) - R,$$

$$\dot{\beta}(\varphi) = \frac{(\beta_H - \beta_L)\, \phi e^{-\phi(\psi - s)}}{(1 + e^{-\phi(\psi - s)})^2},$$

$$\Gamma(t) \notin G \text{ for some } 0 < t < T,$$

$$\text{given } \Gamma(0) = \bar{\Gamma}, R(0) = \bar{R}, \beta(0) = \beta_L,$$

where for compactness we write $\Gamma = (\gamma, \pi^e)'$ and define $\bar{g}(\Gamma, R, \beta)$ and $\bar{M}(\Gamma, \beta)$ so that $\dot{\Gamma}$ and $\dot{R}$ are the mean dynamics of beliefs and expected inflation (11)-(13). $Q(\Gamma, R, \beta)$ is the likelihood function of the shocks needed to perturb beliefs by $\dot{v}$ and beliefs have to escape a neighbourhood G around their initial self confirming equilibrium levels $\bar{\Gamma}$ and $\bar{R}$.

To solve the optimal control problem we define the Hamiltonian:

$$H = a\bar{g}(\Gamma, R, \beta) - \frac{1}{2}\, a' Q(\Gamma, R, \beta) a + \lambda(\bar{M}(\Gamma, \beta) - R) + \mu \left(\frac{(\beta_H - \beta_L)\, \phi e^{-\phi(\psi - s)}}{(1 + e^{-\phi(\psi - s)})^2} \right), \quad (14)$$

where a, λ and μ are co-state vectors for the evolution of Γ, R and β.[4] The Hamiltonian is

[4] An analytic expression for $a\prime Q(\gamma, R)a$ is given in Appendix A.

convex so the following first order conditions hold along the dominant escape path:

$$
\begin{aligned}
\dot{\Gamma} &= \bar{g}(\Gamma, R, \beta) - Q(\Gamma, R, \beta)a, \\
\dot{R} &= \bar{M}(\Gamma, \beta) - R, \\
\dot{\beta} &= \frac{(\beta_H - \beta_L)\,\phi e^{-\phi(\psi - s)}}{(1 + e^{-\phi(\psi - s)})^2}, \\
\dot{a} &= -a\frac{\partial \bar{g}(\Gamma, R, \beta)}{\partial \Gamma} + \frac{1}{2}\,a'\frac{\partial Q(\Gamma, R, \beta)}{\partial \Gamma}a - \lambda\frac{\partial \bar{M}(\Gamma, \beta)}{\partial \Gamma}, \\
\dot{\lambda} &= -H_R, \\
\dot{\mu} &= -H_\beta.
\end{aligned}
$$

The first-order conditions form a system of ordinary differential equations that characterise a family of escape paths, with each path being indexed by different initial values of the co-state vectors. The dominant escape path is the member of this family that achieves the escape with the most likely series of belief perturbations. A solution to the optimal control problem can therefore be obtained by searching for the initial values of a, λ and μ that imply the most likely belief perturbations.

4.1 Numerical example

The short run dynamics of the dominant escape path are the most likely interesting thing that can happen to beliefs and expected inflation after a change in the policymaker's mandate. To appreciate what this interesting behaviour implies, we calculate the dominant escape path in our model when $U^* = 5, \theta = -1, \sigma_1 = \sigma_2 = 0.3$ and the policymaker's mandate changes from $\beta_L = 1$ to $\beta_H = 1.5$ at time $s = 1.5$ with abruptness parameter $\phi = 30$. The dynamics of the dominant escape path are not very sensitive to particular parameter values, so these numerical results are representative of the most likely interesting way that beliefs and expected inflation adjust to their new long run values in our model. Figure 2 shows how policymaker's beliefs, intended and expected inflation, unemployment and the policymaker's mandate all evolve along the dominant escape path.

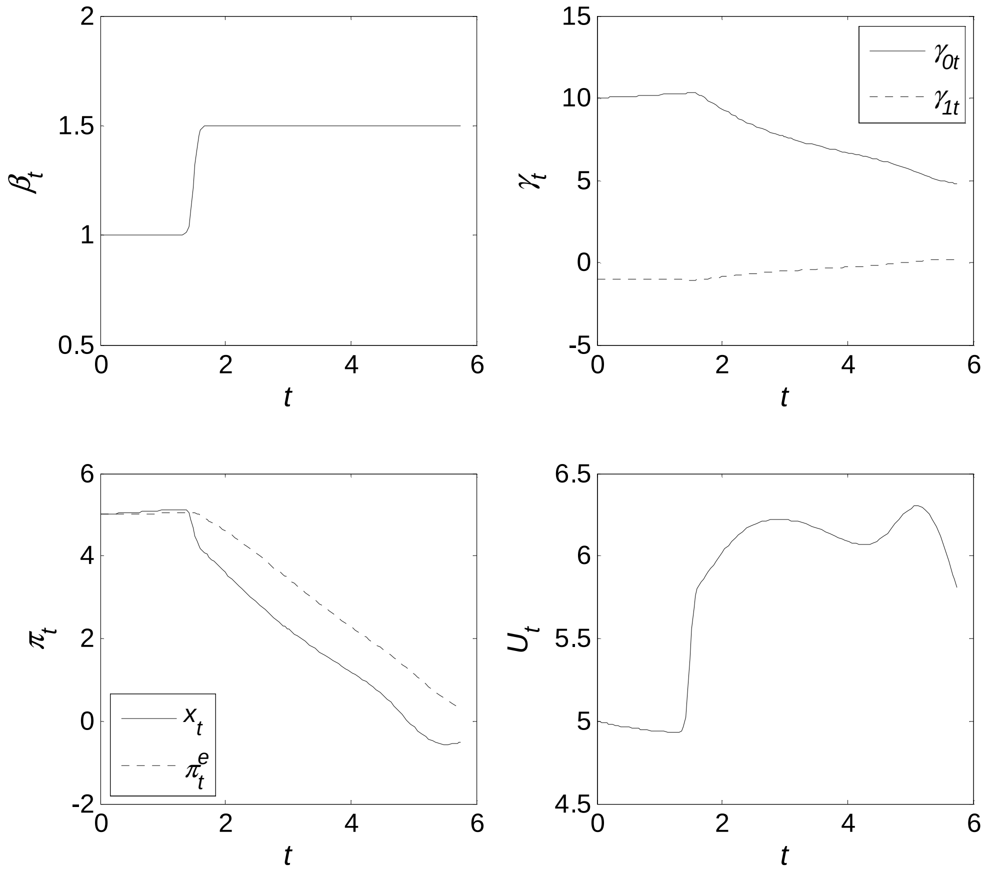

Figure 2: The dominant escape path with a changing policymaker's mandate

The beliefs of the policymaker in the top right hand panel of Figure 2 are initialised at levels $\gamma_t = (10, -1)$ consistent with the self confirming equilibrium for $\beta_t = \beta_L$. However, they start to escape to $\gamma_t = (5, 0)$ as soon as the mandate of the policymaker starts to change. This is an escape episode of the Sargent (1999) type, synonymous with the policymaker abandoning its belief that there is a relationship between unemployment and inflation. The dynamics of the escape episode are so powerful that the beliefs of the policymaker profoundly overshoot the level $\gamma_t = (8\frac{1}{3}, -1)$ of the self confirming equilibrium for $\beta_t = \beta_H$. Only after the escape episode is over will mean dynamics take over and policymaker's beliefs converge to their new long run levels.

The behaviour of intended and expected inflation in the bottom left hand panel of Figure 2 provides evidence for why policymaker's beliefs escape after a change in the policymaker's mandate. The disinflation starts as soon as the mandate places increased emphasis on inflation stabilisation. The initial fall in intended inflation is a surprise to private agents, so unemployment rises in the bottom right hand panel of Figure 2. Shortly afterwards, private agents start revising their inflation expectations downwards and there follows a protracted period during which intended and expected inflation both fall. Private agents forming adaptive expectations continue to be surprised by the falls in inflation, but the gap between intended and expected inflation does not increase so unemployment is high but stable. The policymaker sees that successive reductions in inflation have no incremental effect on unemployment, which causes them to doubt their belief in a relationship between inflation and unemployment. The beliefs of the policymaker then gradually escape, a process only halted when intended inflation has fallen to its target level of zero and the disinflation comes to an end. At this point, mean dynamics take over and inflation will rises steadily back to the value $\theta U^* \beta_H^{-1} = 3\frac{1}{3}$ consistent with the self confirming equilibrium for $\beta_t = \beta_H$. Unemployment similarly returns to its natural rate $U^* = 5$.

Inflation dramatically undershoot its long run value when the policymaker is learning by disinflating. A change in the policymaker's mandate causes inflation to fall to a level close to zero in our model, a result we believe is new in the literature. The tendency for inflation to undershoot suggests that politicians need to exhibit restraint before claiming that low levels of inflation are a consequence of their decision to change the policymaker's mandate. If the policymaker is learning by disinflating then the undershoot in inflation may create an unwarranted sense in politicians that inflation is under control. In reality, inflation is destined to rise again once mean dynamics take over, albeit to a long run level that is lower than before the mandate was changed.

5 Simulation results

The results of the previous section suggest that inflation may undershoot its new long run level after a change in the policymaker's mandate. Whether this happens depends on the probability of beliefs escaping during their transition to levels consistent with the new mandate. To quantify this probability, we simulate the model many times and keep track of how often inflation undershoots its new long run target. We set $U^* = 5, \theta = -1, \sigma_1 = \sigma_2 = 0.3$ as before, and assume that the policymaker's mandate changes from $\beta_L = 1$ to $\beta_H = 1.5$ at time $s = 100$ with abruptness parameter $\phi = 0.25$. The gain coefficients in the learning of the policymaker and private agents are set symmetrically at $a_g = a_p = 0.0275$. We simulate the model 5000 times and present the results in Figure 3.

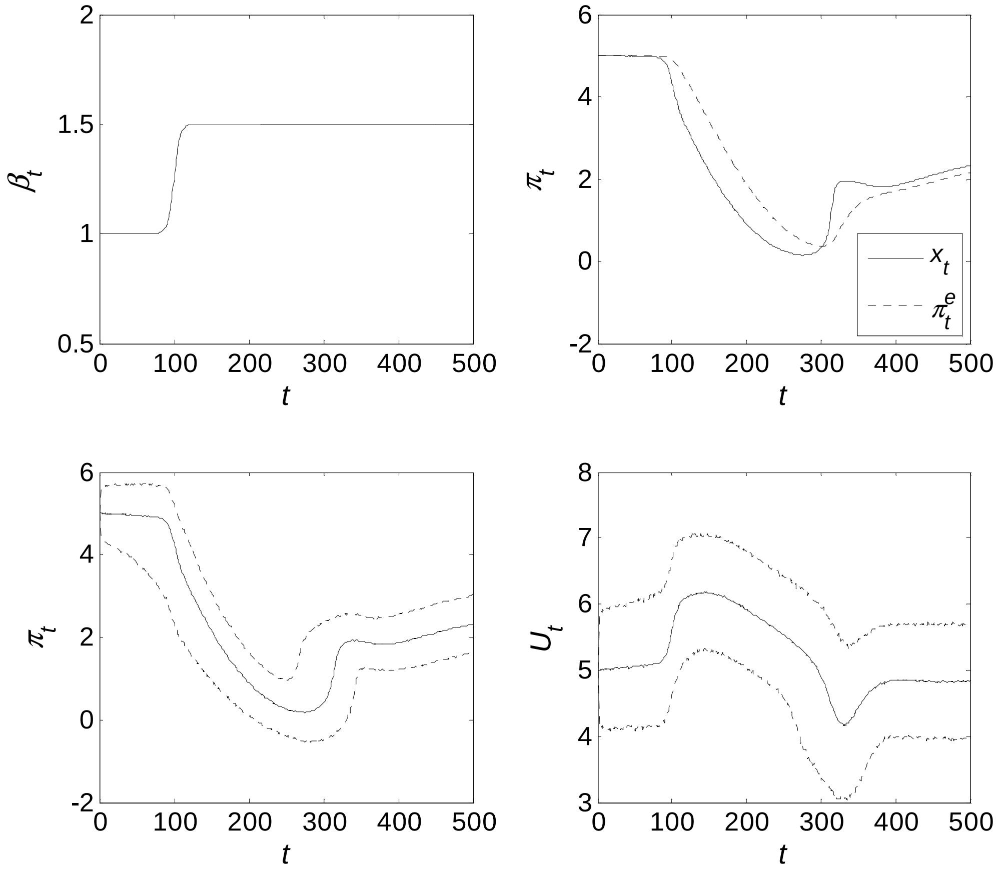

Figure 3: Simulated dynamics with a changing policymaker's mandate

The solid line in the bottom left hand panel of Figure 3 is the median simulated path of inflation. The dashed lines in the same panel are maximum and minimum paths, so inflation undershoots its long run level of $3\frac{1}{3}$ in each of the 5000 simulations. The median behaviour of intended and expected inflation in the top right panel of Figure 3 confirms the predictions of our analytical results. Intended inflation falls as soon as the policymaker's mandate changes, with expected inflation beginning to fall shortly afterwards. The disinflation continues until inflation is close to its target level of zero, at which point inflation begins rising again as mean dynamics start to dominate. The bottom right hand panel of Figure 3 shows unemployment increasing and then returning to its natural rate during the disinflation. The dip in unemployment after period $t = 300$ coincides with a rapid rise in intended inflation, and is caused by the inflation expected by private agents falling below that intended by the policymaker.

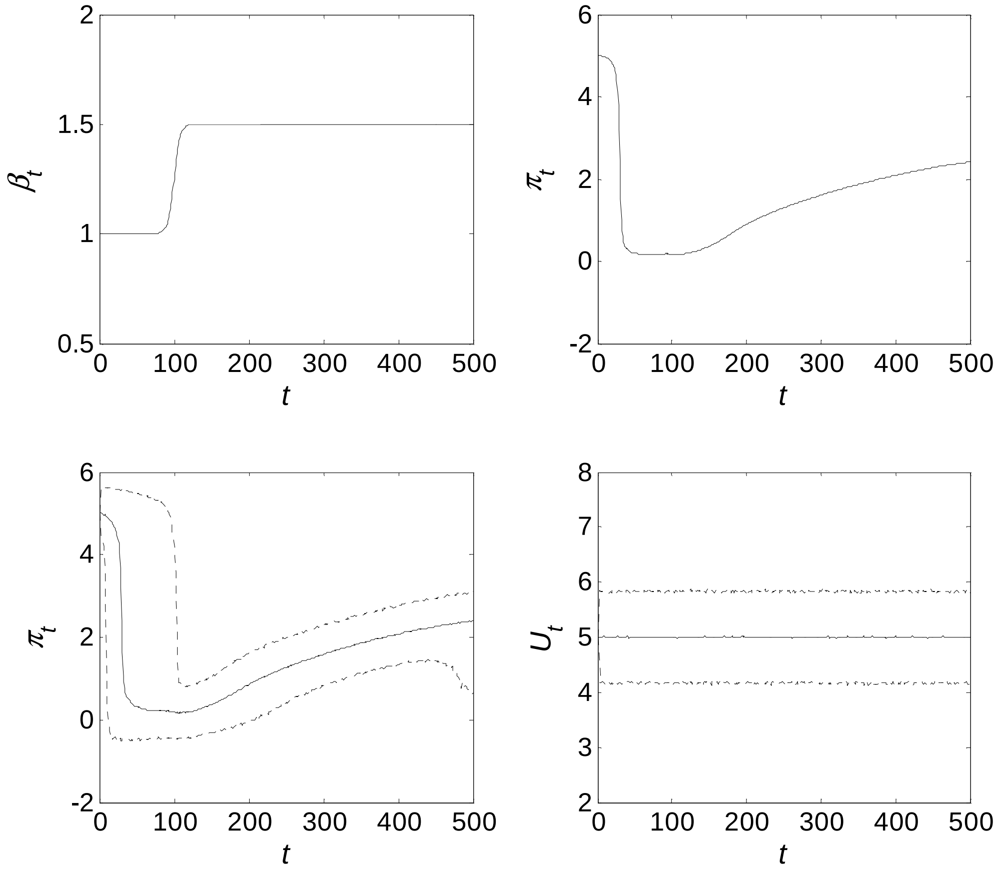

Figure 4: Simulated dynamics under rational expectations

The learning of private agents is important for explaining why unemployment rises when disinflating in our model, but is not instrumental in inflation undershooting its long run level. This is apparent in Figure 4, which presents simulation results for a variant of our model in which private agents have rational rather than adaptive expectations. Unemployment is no longer affected by the disinflation, but inflation still undershoots its long run level in every simulation of the model.

6 Conclusion

The undershooting of inflation in our model occurs because the policymaker is learning by disinflating. This emphasis on learning by the policymaker contrasts with the existing literature, which typically focuses on the learning by private agents. There is considerable evidence that policymakers do learn by disinflating. For example, Bordo et al. (2007) argue that policymakers were learning in three great American disinflations. In the post civil war disinflation, learning led to a temporary suspension of monetary tightening in 1867 as Congress and the US Treasury observed the sharp price deflation having a contractionary impact on the economy. Monetary tightening only resumed in 1869, with policymakers learning to adopt a more gradual approach. In the post World War I disinflation, policymakers instead learned to take a much more aggressive stance and rapidly disinflated in the 1920s.

The Volcker disinflation is closest to our analysis as it was predicated on a desire to reduce inflation, rather than the post civil war and post WWI disinflation which aimed at reducing prices to a level consistent with the gold standard. As such, the danger of inflation undershooting and creating a false sense of how successful changing the central bank's mandate has been is highest during the Volcker disinflation. The evolution of inflation after the Volcker disinflation is shown in Figure 5, where inflation did fall but began to pick up after 1998. It is difficult to decide whether this is the undershooting of inflation predicted by our model, since inflation started to fall again at the beginning of the 2001 recession.

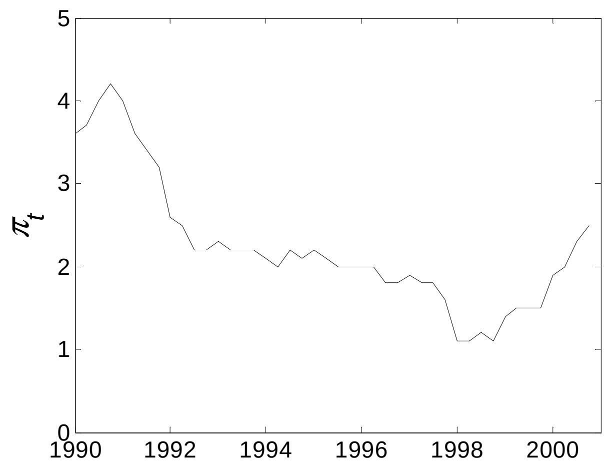

Figure 5: 12 month change in the GDP price deflator

It appears that policymakers did learn during the Volcker disinflation. In particular, policymakers began to argue that there had been a flattening of the Phillips curve, such that changes in output would have less impact on inflation than before. For example, Bernanke (2008) refers to work by Mishkin (2007):

"As he noted, many studies of the conventional Phillips curve find that the sensitivity of inflation to activity indicators is lower today than in the past (that is, the Phillips curve appears to have become flatter); and that the long-run effect on inflation of "supply shocks," such as changes in the price of oil, also appears to be lower than in the past."

Arguing that the Phillips curve has become flatter is what we would expect if the policymaker is learning by disinflating. It is consistent with the policymaker in our model learning to update its estimate γ_{1t} of the slope of the perceived relationship between unemployment and inflation. Further support for the idea that policymakers learnt during the Volcker disinflation is provided by Tillmann (2010), who reports that FOMC forecasts are consistent with the notion that the Phillips curve changed during the second half of the 1990s and that, on average, FOMC members took that into account when submitting their forecasts.

References

[1] Ball, L., 1994, Credible Disinflation with Staggered Price-Setting, American Economic Review, 84:1, 282-89

[2] Bordo, M.D, Erceg, C., Levin, A. and R. Michaels, 2007, Three Great American Disinflation, NBER Working Paper No. 12982

[3] Cho, I-K and K. Kasa, 2008, Learning Dynamics and Endogenous Currency Crises, Macroeconomic Dynamics, 12:2, 257-285

[4] Cho, I-K, Williams, N. and T.J. Sargent, 2002, Escaping Nash Inflation, Review of Economic Studies, 69, 1-40

[5] Cogley, T., Matthes, C. and A.M. Sbordone, 2011, Optimal Disinflation Under Learning, mimeo

[6] Cogley, T., Colacito, R. and T.J. Sargent, 2007, Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas, Journal of Money, Credit and Banking, Blackwell Publishing, 39:s1, 67-99

[7] Erceg, C. J. and A.T. Levin, 2003, Imperfect credibility and inflation persistence, Journal of Monetary Economics, 50:4, 915-944

[8] Evans, G.W. and S.M.S. Honkapohja, 2001, Learning and Expectations in Macroeconomics, Princeton University Press

[9] Evans, L. and Nicolae, A., 2010, The Output Effect of a Transition to Price Stability When Velocity is Time-varying, Journal of Money, Credit and Banking, 42:5, 859-878

[10] Ireland, P. N., 1995, Optimal disinflationary paths , Journal of Economic Dynamics and Control, 19:8, 1429-1448

[11] Gerali, A. and F. Lippi, 2002, On the 'Conquest' of Inflation, mimeo

[12] Goodfriend. M. and R. King, 2005, The Incredible Volcker Disinflation, NBER Working Paper No. 11562

[13] Bernanke, B.S., 2007, Inflation Expectations and Inflation Forecasting, Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, July 10

[14] Kreps, D.M., 1998, Anticipated Utility and Dynamic Choice, in Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997, Cambridge University Press

[15] McGough, B., 2006, Shocking Escapes, Economic Journal, 116:4, 507-528

[16] Mishkin, F.S., 2007, Inflation Dynamics, National Bureau of Economic Research Working Paper No. 13147

[17] Nicolae, A. and C. Nolan, 2006, The impact of imperfect credibility in a transition to price stability, Journal of Money, Credit, and Banking, 38:1, 47–66

[18] Primiceri, G.E., 2006, Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy, The Quarterly Journal of Economics, 121:3, 867-901

[19] Rogoff, K., 1985. The Optimal Degree of Commitment to an Intermediate Monetary Target. Quarterly Journal of Economics. 100:4, 1169-1190

[20] Sargent, T.J., 1999, The Conquest of American Inflation, Princeton University Press

[21] Sargent, T.J., Williams, N. and T. Zha, 2006, Shocks and Government Beliefs: The Rise and Fall of American Inflation, American Economic Review, 96:4, 1193-1224

[22] Tillmann, P., 2010, The Fed's perceived Phillips curve: Evidence from individual FOMC forecasts, Journal of Macroeconomics, 32:4, 1008-10013

[23] Williams, N., 2004, Escape Dynamics in Learning Models, University of Chicago UMI 3006566

A Analytic expression for $a'Q(\Gamma, R, \beta)a$

The likelihood function $Q(\Gamma, R, \beta)$ is used to weight belief perturbations along candidate escape paths. It is equal to the variance-covariance matrix of belief dynamics $\dot{\Gamma}$, and is a fourth moment matrix because belief dynamics are quadratic forms of Gaussian variables. In static models such as ours, Williams (2004) shows that $Q(\Gamma, R, \beta)$ reduces to the logarithm of a moment generating function, meaning the Hamiltonian (14) can be derived analytically. We begin by expressing the second term of the Hamiltonian by the corresponding moment generating function:

$$a'Q(\Gamma, R, \beta)a = \log E \exp \left\langle a \cdot R^{-1}(g(\Gamma, R, \beta, \xi) - \bar{g}(\Gamma, R, \beta)) \right\rangle, \tag{A.1}$$

and using (8) and (9) to obtain an explicit analytic expression:

$$R^{-1}(g(\Gamma, \beta, \xi) - \bar{g}(\Gamma, \beta)) =$$
$$\left(\begin{array}{c} R^{-1} \left(\begin{array}{c} \nu_1 - (\theta + \gamma_1)\nu_2 \\ x\nu_1 + (U^* - (\theta - \pi^e + \gamma_1)x)\nu_2 + \nu_1\nu_2 - (\theta + \gamma_1)\nu_2^2 + (\theta + \gamma_1)\sigma_2^2 \end{array} \right) \\ \nu_2 \end{array} \right)$$

To economise on notation, let R^{-1} and a be defined by:

$$R^{-1} = \begin{pmatrix} R^1 & R^2 \\ R^2 & R^4 \end{pmatrix}, \quad a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}.$$

The right-hand side of (A.1) can be expressed in terms of the underlying shocks v_1 and v_2:

$$\log E \exp \left\langle a \cdot R^{-1}(g(\Gamma, R, \beta, \xi) - \bar{g}(\Gamma, R, \beta)) \right\rangle = \log E[e^{d_0 + d_1 v_1 + d_2 v_2 + d_3 v_1 v_2 + d_4 v_2^2}], \tag{A.2}$$

where the constants $d_0, ..., d_4$ are simple functions of the structural parameters $\{U^*, \theta, \sigma_1, \sigma_2\}$, beliefs γ, expected inflation π^e, the co-state vector a and the precision matrix R:

$$\frac{\gamma_{t+1}-\gamma_t}{a_g} = R_t^{-1}\begin{pmatrix} 1 \\ x_t + v_{2t} \end{pmatrix}(U^* - \theta(x_t + v_{2t} - \pi_t^e) + \nu_{1t} - \gamma_{0t} - \gamma_{1t}(x_t + v_{2t}))$$

$$\begin{aligned}
d_0 &= (\theta+\gamma_1)(a_1R^2 + a_2R^4)\sigma_2^2 \\
d_1 &= a_1R^1 + a_2R^2 + (a_1R^2 + a_2R^4)x \\
d_2 &= -(\theta+\gamma_1)(a_1R^1 + a_2R^2) + (U^* - \theta\pi^e - \gamma_0 - 2\gamma_1 x)(a_1R^2 + a_2R^4) + a_3 \\
d_3 &= a_1R^2 + a_2R^4 \\
d_4 &= -(\theta+\gamma_1)(a_1R^2 + a_2R^4)
\end{aligned}$$

The next step is to factorise v_1 out from the right-hand side of (A.2). The key stage in the factorisation below is the second line, where we exploit the fact that e^{v1} is log-normally distributed with expected value half the variance of v_1.

$$\begin{aligned}
\log E[e^{d_0+d_1v_1+d_2v_2+d_3v_1v_2+d_4v_2^2}] &= d_0 + \log E[E\left(e^{(d_1+d_3v_2)v_1}\middle| v_2\right)e^{d_2v_2+d_4v_2^2}], \\
&= d_0 + \log E[e^{0.5(d_1+d_3v_2)^2}e^{d_2v_2+d_4v_2^2}], \\
&= d_0 + 0.5d_1^2 + \log E[e^{(d_2+d_1d_3)v_2+(d_4+0.5d_3^2)v_2^2}].
\end{aligned}$$

The outcome of factorisation is an expression in only the v_2 shock. The remaining expectation can be solved analytically by defining $k_1 = d_2 + d_1d_3, k_2 = d_4 + 0.5d_3^2 - 0.5$ and completing the square of $k_1x + k_2x^2$. Defining $A = \sqrt{-2k_2}, B = -k_1/A, C = -B^2/2$, we have:

$$\begin{aligned}
E[e^{(d_2+d_1d_3)v_2+(d_4+0.5d_3^2)v_2^2}] &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{(k_1x+k_2x^2)}dx, \\
&= \frac{e^{-C}}{\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-0.5(Ax+B)^2}dx, \\
&= \frac{e^{-C}}{A},
\end{aligned}$$

and the final analytic expression for $a'Q(\Gamma, R, \beta)a$ is:

$$a'Q(\Gamma, R, \beta)a = d_0 + 0.5d_1^2 - \log A - C.$$