

Optimal Disinflation Under Learning^{*}

Timothy Cogley
New York University

Christian Matthes
Universitat Pompeu Fabra and Barcelona GSE

Argia M. Sbordone
Federal Reserve Bank of New York

December 2010

Abstract

We model transitional dynamics that emerge after the adoption of a new monetary-policy rule. We assume that private agents learn about the new policy via Bayesian updating, and we study how learning affects the nature of the transition and choice of a new rule. The model endogenously generates time-varying volatility during the transition. Managing this volatility is the central bank's main challenge. The optimal policy depends on subtle features of the private sector's prior. Nevertheless, two robust conclusions emerge from our examples. First, the central bank can adjust target inflation freely without triggering high volatility. Second, none of our examples rationalizes a gradual reduction in inflation. On the contrary, inflation falls sharply at impact, overshoots the new target, and converges from below.

1 Introduction

We examine the problem of a newly-appointed central bank governor who inherits a high average inflation rate from the past. The bank has no official inflation target and lacks the political authority unilaterally to set one, but it has some flexibility in

^{*}For comments and suggestions, we thank Martin Ellison, Boyan Jovanovic, Thomas Sargent, and seminar participants at the Banque de France, Duke, FRB Atlanta, FRB Philadelphia, the Federal Reserve Board, NYU, Oxford, Pompeu Fabra, Rutgers, the Toulouse School of Economics, and UQAM. The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

choosing how to implement a vague mandate. We assume that the new governor's preferences differ from his predecessor and that he wants to disinflate. We want to find an optimal Taylor-type rule and to determine whether the optimal policy generates a gradual or sharp decline in inflation.

Sargent (1982) studies an analogous problem in which the central bank not only has a new governor but also undergoes a fundamental institutional reform. He argues that by suitably changing the rules of the game, the government can persuade the private sector in advance that a low-inflation policy is its best response. In that case, the central bank can engineer a sharp disinflation at low cost. Sargent discusses a number of historical examples that support his theory, emphasizing the institutional changes that establish credibility. Our scenario differs from Sargent's in two ways. We take institutional reform off the table, assuming instead just a change of personnel. We also take away knowledge of the new policy and assume that the private sector must learn about it. This is tantamount to assuming that the private sector does not know the new governor's preferences. Our scenario is more like the Volcker disinflation than the end of interwar hyperinflations.

In contrast to Sargent's case studies, the Volcker disinflation was quite costly. Erceg and Levin (2003) and Goodfriend and King (2005) explain the high cost by pointing to a lack of transparency and credibility. Erceg and Levin contend that Volcker's policy lacked transparency, and they develop a model in which the private sector must learn the central bank's long-run inflation target.¹ In their model, learning makes inflation more persistent relative to what it would be under full information, increases the sacrifice ratio, and produces output losses like those seen in the early 1980s. Goodfriend and King claim that Volcker's disinflation lacked credibility because no important changes were made in the rules of the game. Because the private sector was initially unconvinced that Volcker would disinflate, the new policy collided with expectations inherited from the old regime and brought about a deep recession.

The analysis of Erceg, Levin, Goodfriend, and King is positive and explains why the Volcker disinflation was costly. In contrast, we address a normative question, viz. what policy is optimal when the private sector is not persuaded in advance and must learn the new policy? Are Sargent's recommendations still wise? Or should the disinflation be more gradual, e.g. as recommended by Gordon (1982)? Intuitively, the answer could go either way. Because learning increases inflation persistence and the sacrifice ratio, it might promote gradualism. But gradualism would presumably slow the rate at which the private sector learns, prolonging the transition. Our objective is to assess the competing forces.

We study the question in the context of a dynamic new Keynesian model modified in two ways. Following Ascari (2004) and Sbordone (2007), we assume that target

¹See also Orphanides and Williams (2005) and Milani (2007).

inflation is positive. We also replace rational expectations with Bayesian learning. We assume the central bank follows a simple Taylor-type rule and chooses its coefficients by minimizing a discounted quadratic loss function. The private sector learns the new policy via Bayesian updating, and the central bank takes learning into account when solving its decision problem.

Our approach to learning differs from much of the macro-learning literature, in particular from the branch emanating from Marcet and Sargent (1989a, 1989b) and Evans and Honkapohja (2001). Models in that tradition typically assume that agents use reduced-form statistical representations such as vector autoregressions (VARs) for forecasting. They also commonly assume that agents update parameter estimates by recursive least squares. In contrast, we assume that agents update beliefs via Bayes's theorem. We do this for two reasons, first because we want agents to retain what they know about other aspects of the economy's structure and also because we want the speed of learning to emerge as an endogenous outcome. The agents who inhabit our model utilize VARs for forecasting, but their VARs satisfy cross-equation restrictions analogous to those in rational-expectations models. As a consequence, there is tight link between the actual and perceived laws of motion (ALM and PLM, respectively). In our model, agents know the ALM up to the unknown monetary-policy parameters, and their PLM is the perceived ALM (i.e., the ALM evaluated at their current estimate of the policy coefficients). Because agents know the functional form of the ALM, they can use Bayes theorem to update beliefs, efficiently exploiting information about the new monetary-policy rule.²

For the model described below, the central bank's main challenge is to manage high volatility during the transition. For some choices of Taylor-rule parameters, the equilibrium law of motion is a temporarily explosive process, i.e. one that is asymptotically stationary but which has explosive autoregressive roots during the transition. When this happens, the transition is highly volatile and dominates expected loss. It is up to the central bank to choose a policy that manages this turbulence.

Explosive turbulence emerges when there is a lot of disagreement between the actual and perceived policy. One way for the bank to avoid this is to make an incremental change in policy, i.e. one that represents only a slight perturbation relative to the private sector's prior. Incremental reforms limit initial disagreement and therefore generate a smooth transition at the cost perhaps of improving only slightly on the old regime. More ambitious reforms can succeed if the private sector learns rapidly enough. Because the speed of learning depends on the private sector's prior, optimal policy also depends on the prior. We study a variety of scenarios to explore how the prior affects the nature of the transition and optimal policy.

²Our approach detaches learning from bounded rationality. Although the latter might turn out to be important for understanding disinflation, we think that disentangling the two is helpful for understanding their respective contributions. This paper takes a step in that direction, focusing on the role of learning.

For the examples we study, however, two robust features of policy emerge. First, the potential for explosive arises chiefly from disagreement between actual and perceived feedback parameters for inflation and output, and not so much because of disagreement about the central bank’s long-run inflation target. The bank can therefore move its inflation target far from the private sector’s prior without triggering explosive volatility. Thus it can achieve low average inflation without too volatile a transition. Second, the optimal policy pulls inflation down sharply at impact, overshooting the new long-run inflation target. After that, inflation converges to its new target from below.³ None of our examples rationalize a gradual reduction in inflation.

The remainder of the discussion is organized as follows. Section 2 outlines a dynamic new Keynesian model with non-zero trend inflation, section 3 describes how the private sector learns, and section 4 explains how the optimal-policy problem is solved. Sections 5 and 6 describe optimal policies for various scenarios and examine how they depend on the private sector’s prior. Section 7 summarizes our findings and lists directions for future research.

2 A dynamic new-Keynesian model with positive target inflation

We begin by describing the timing protocol, a critical element in learning models. Then, taking beliefs as given, we describe the model’s structure and our approximation methods. A discussion of how beliefs are updated is deferred to section 3.

2.1 The timing protocol

Private agents enter period t with beliefs about policy coefficients inherited from $t-1$. They form expectations and make current-period plans accordingly. At the same time, the central bank sets the systematic part of its instrument based on information inherited from $t-1$. After that, shocks are realized and current-period outcomes are determined. After observing those outcomes, private agents update their estimates. They treat estimated parameters as if they were known with certainty, and they formulate new plans which they carry forward to period $t+1$.⁴

³Inflation oscillates in some cases, but spends more time below $\bar{\pi}$ than above.

⁴This timing protocol differs slightly from the convention in DSGE models, but it is convenient because it circumvents a simultaneity between the determination of outcomes, the formation of expectations, and the updating of beliefs.

2.2 The model

Our model is a simple dynamic new Keynesian model modified so that target inflation is positive and that agents take expectations with respect to subjective beliefs. Monetary policy is determined in accordance with a Taylor-type rule. Private-sector behavior is characterized by two blocks of equations, a conventional intertemporal *IS* curve and an Ascari-Sbordone version of the aggregate supply curve. The model features habit persistence in consumption and staggered price setting. This section presents a log-linearized version of the model. For details on how we arrived at this representation, see the web appendix.

2.2.1 Monetary policy

We assume that the central bank commits to a Taylor rule in difference form,

$$i_t - i_{t-1} = \psi_\pi(\pi_{t-1} - \bar{\pi}) + \psi_x(y_{t-1} - y_{t-2}) + \varepsilon_{it}, \quad (1)$$

where i_t is the nominal interest rate, π_t is inflation, y_t is log output, and ε_{it} is an *i.i.d.* policy shock. The timing assumption follows McCallum (1999) and fits conveniently within the timing protocol described above.⁵ The policy parameters are $\psi = [\bar{\pi}, \psi_\pi, \psi_x, \sigma_i^2]'$, where $\bar{\pi}$ represents the central bank's long-run inflation target, ψ_π and ψ_x are feedback parameters on the inflation gap and output growth, respectively, and σ_i is the standard deviation of the policy shock.

We adopt this form because others have shown that it is promising for environments like ours. For instance, Coibion and Gorodnichenko (2008) establish that a rule of this form ameliorates indeterminacy problems in Calvo models with positive target inflation. Orphanides and Williams (2007) demonstrate that it performs well under least-squares learning.⁶ More generally, a number of economists have argued that the central bank should engage in a high degree of interest smoothing (e.g. Woodford (1999)). Erceg and Levin (2003) contend that output growth, rather than the output gap, is the appropriate measure to include in an estimated policy reaction function for the U.S.

We assume that private agents know the form of the policy rule but not the policy coefficients. At any given date, the perceived policy rule is

$$i_t - i_{t-1} = \psi_{\pi t}(\pi_{t-1} - \bar{\pi}_t) + \psi_{xt}(y_{t-1} - y_{t-2}) + \varepsilon_{it}, \quad (2)$$

⁵McCallum (1999) contends that monetary policy rules should be specified in terms of lagged variables, on the grounds that the Fed lacks good current-quarter information about inflation, output, and other right-hand variables. This is especially relevant for decisions taken early in the quarter, in accordance with our timing protocol.

⁶Orphanides and Williams (2007) postulate that neither the agents nor the central bank know the true structure of the economy, and they replace rational expectations with least-squares learning. They assume that the central bank cannot observe natural rates and that the bank estimates them via a simple updating rule. They show that an optimized Taylor rule in difference form dominates an optimized standard Taylor rule when learning and time-varying natural rates interact.

where $\psi_t = [\bar{\pi}_t, \psi_{\pi t}, \psi_{xt}, \sigma_{it}^2]$ represents the beginning-of-period estimate of ψ .

The perceived law of motion depends on the perceived policy (2). The actual law of motion depends on actions taken by the central bank and decisions made by the private sector. Hence the actual law of motion involves both the actual policy (1) and the perceived policy (2).

Finally, we assume that the central bank chooses ψ by minimizing a discounted quadratic loss function,

$$L = E_0 \sum_t \beta^t [\pi_t^2 + \lambda_y (y_t - \bar{y})^2 + \lambda_i (i_t - \bar{i})^2], \quad (3)$$

taking private-sector learning into account.⁷ In addition to penalizing variation in inflation and the output gap, the loss function includes a small penalty for deviations of the nominal interest rate from its steady state. The central bank arbitrarily sets σ_i and optimizes with respect to $\bar{\pi}$, ψ_π , and ψ_x .

2.2.2 Approximation methods

We use two approximations when solving the model. As usual, the first-order conditions take the form of non-linear expectational difference equations. We follow the standard practice of log-linearizing around a steady state and solving the resulting system of linear expectational difference equations. However, we expand around the agents' perceived steady state in period t rather than around the true steady state.

The true steady state $\bar{x}$ is the deterministic steady state associated with the true policy coefficients ψ . We define the perceived steady state $\bar{x}_t$ as the long-horizon forecast associated with the current estimate ψ_t . The private sector's long-run forecast $\bar{x}_t$ varies through time because changes in the central bank's inflation target have level effects on nominal variables and also on some real variables (Ascari 2004). Since perceptions of $\bar{\pi}$ change as agents update their beliefs, so do their long-run forecasts.

We chose to expand around $\bar{x}_t$ instead of $\bar{x}$ because the plans of consumers and firms follow from their first-order conditions, and $\bar{x}_t$ better reflects their state of mind at date t . The perceived steady state $\bar{x}_t$ converges to the full-information steady state $\bar{x}$ if private agents learn the true policy coefficients, but the two differ along the transition path.

Our second approximation involves the assumption that agents treat the current estimate ψ_t as if it were known with certainty. Kreps (1998) calls this as an 'anticipated-utility' model. In the context of a single-agent decision problem, Cogley and Sargent (2008) compare the resulting decision rules with exact Bayesian decision rules, and they demonstrate that the approximation is quite good as long as precautionary motives are not too strong. Like a log-linear approximation, this imposes a

⁷Gaspar et al (2006) distinguish between an unsophisticated central bank - one that accounts for the beliefs of the public but not the dynamic process of learning - and a sophisticated central bank that also takes the learning process into account. Our setting corresponds to the latter assumption.

form of certainty equivalence, for it implies that decision rules are the same regardless of the degree of parameter uncertainty. This approximation is standard in the macro-learning literature.

2.2.3 A new-Keynesian IS curve

As usual, we assume that a representative household maximizes expected utility subject to a flow budget constraint. We assume that the household's period-utility function is

$$U_t = \log(C_t - \eta C_{t-1}) - \chi \frac{H_t^{1+\nu}}{1+\nu},$$

where η measures the degree of habit persistence. The first-order condition is a conventional consumption Euler equation. After log-linearizing, we obtain a version of the new-Keynesian *IS* curve

$$\xi_t - \xi = y_t - \bar{y}_t + E_t^* (\xi_{t+1} - \xi - (y_{t+1} - \bar{y}_t) - (g_{t+1} - g) + i_t - \pi_{t+1} - r) + \varepsilon_{yt}, \quad (4)$$

where ξ_t is a transformation of the marginal utility of consumption,

$$\xi_t - \xi = \xi_1 (y_t - \bar{y}_t) + \xi_2 (y_{t-1} - \bar{y}_t - (g_t - g)) + \xi_3 E_t^* ((y_{t+1} - \bar{y}_t) + g_{t+1} - g). \quad (5)$$

The parameter r is the steady-state real interest rate, $\bar{y}_t$ is the private sector's long-run forecast for output, and g_t and ε_{yt} are shocks.

This equation differs in a number of ways from a standard *IS* equation. One difference concerns the choice of the expansion point. As mentioned above, we expand around the perceived steady state $\bar{y}_t$ instead of the actual steady state $\bar{y}$. In addition, our anticipated-utility assumption implies that $E_t^* \bar{y}_{t+1} = \bar{y}_t$, explaining the appearance of $\bar{y}_t$ on the right-hand side of equations (4) and (5).

A second difference concerns the expectation operator E_t^* , which represents forecasts formed with respect to the private sector's perceived law of motion. In contrast, the central bank takes expectations with respect to the actual law of motion, which we denote by E_t .⁸

Finally, two shocks appear, a persistent shock g_t to the growth rate of technology,

$$g_t = (1 - \rho_g) g + \rho_g g_{t-1} + \varepsilon_{gt}, \quad (6)$$

and a white-noise shock ε_{yt} . We introduce the latter so that the private sector faces a nontrivial signal-extraction problem.

⁸We assume that the central bank knows the private sector's prior over ψ . Since the central bank's information set subsumes that of the private sector, the law of iterated expectations implies $E_t^*(E_t x_{t+j}) = E_t^*(x_{t+j})$ for any random variable x_{t+j} and $j \geq 0$ such that both expectations exist. Because the central bank can reconstruct private forecasts, it also follows that $E_t(E_t^* x_{t+j}) = E_t^*(x_{t+j})$. But $E_t x_{t+j} \neq E_t^* x_{t+j}$.

2.2.4 A new-Keynesian Phillips curve

We adopt a purely-forward looking version of Calvo's (1983) pricing model. A continuum of monopolistically competitive firms produce a variety of differentiated intermediate goods that are sold to a final-goods producer. Firms that produce the intermediate goods reset prices at random intervals. In particular, with probability $1 - \alpha$ an intermediate-goods producer has an opportunity to reset its price, and with probability α its price remains the same. Thus we abstract from indexation, in accordance with the estimates of Cogley and Sbordone (2008). Since pricing and supply decisions depend on the beliefs of private agents, we again log-linearize around perceived steady states, obtaining the following block of equations:

$$\pi_t - \bar{\pi}_{t-1} = \kappa_t(y_t - \bar{y}_t) + \beta_t E_t^*(\pi_{t+1} - \bar{\pi}_t) + \varsigma_{t-1}(\delta_t - \bar{\delta}_t) \quad (7)$$

$$+ \gamma_{1t} E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_t) + \phi_{t+1}] + u_t + \varepsilon_{\pi t},$$

$$\phi_t = \gamma_{2t} E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_t) + \phi_{t+1}], \quad (8)$$

$$\delta_t - \bar{\delta}_t = \lambda_{1t}(\pi_t - \bar{\pi}_t) + \lambda_{2t}(\delta_{t-1} - \bar{\delta}_t). \quad (9)$$

This representation differs in four ways from standard versions of the *NKPC*. First, a variable

$$\delta_t \equiv \log \left(\int_0^1 (p_t(i) / P_t)^{-\theta} di \right), \quad (10)$$

that measures the resource cost induced by cross-sectional price dispersion has first-order effects on inflation and other variables. If target inflation were zero, this variable would drop out of a first-order expansion.

Second, higher-order leads of inflation appear on the right-hand side of (7). To retain a first-order form, we introduce an intermediate variable ϕ_t that has no interesting economic interpretation and add equation (8). This is simply a device for obtaining a convenient representation.

Third, the *NKPC* coefficients depend on both deep parameters and estimates of target inflation $\bar{\pi}_t$,⁹

$$\begin{aligned} \beta_t &= \beta(1 + \bar{\pi}_t), \\ \kappa_t &= (1 + \nu) [1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}] [1 - \alpha\beta(1 + \bar{\pi}_t)^\theta] / \alpha(1 + \bar{\pi}_t)^{\theta-1}, \\ \gamma_{1t} &= \beta\bar{\pi}_t [1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}], \\ \gamma_{2t} &= \alpha\beta(1 + \bar{\pi}_t)^{\theta-1}, \\ \varsigma_t &= \nu [1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}] [1 - \alpha\beta(1 + \bar{\pi}_t)^\theta] / \alpha(1 + \bar{\pi}_t)^{\theta-1}, \\ \lambda_{1t} &= \alpha\theta\bar{\pi}_t(1 + \bar{\pi}_t)^{\theta-1} / (1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}), \\ \lambda_{2t} &= \alpha(1 + \bar{\pi}_t)^\theta. \end{aligned} \quad (11)$$

⁹The *NKPC* parameters collapse to the usual expressions when $\bar{\pi}_t = 0$.

The deep parameters are the subjective discount factor β , the probability $1-\alpha$ that an intermediate-goods producer can reset its price, the elasticity of substitution across varieties θ , and the Frisch elasticity of labor supply $1/\nu$. As Cogley and Sbordone (2008) emphasize, the deep parameters are invariant to changes in policy, but the *NKPC* coefficients are not. The latter change as beliefs about $\bar{\pi}_t$ are updated.

Finally, we assume two cost-push shocks, a persistent shock u_t that follows an *AR*(1) process,

$$u_t = \rho_u u_{t-1} + \varepsilon_{ut}, \quad (12)$$

and a white-noise shock $\varepsilon_{\pi t}$. As before, the latter is included so that agents face a nontrivial signal-extraction problem.

2.3 Calibration

Parameters of the pricing model are taken from estimates in Cogley and Sbordone (2008),

$$\alpha = 0.6, \quad \beta = 0.99, \quad \theta = 10. \quad (13)$$

Notice in particular that we abstract from indexation or other backward-looking pricing influences. We think this is realistic, as it is supported both by the estimates in our earlier paper and by micro data.

We calibrate the labor-disutility parameters as

$$\nu = 0.5, \quad \chi = 1. \quad (14)$$

The parameter ν is the inverse of the Frisch elasticity of labor supply – i.e., the elasticity of hours worked with respect to the real wage, keeping constant the marginal utility of consumption. The literature provides a large range of values for this elasticity, typically high in the macro literature and low in the labor literature. Our calibration implies a Frisch elasticity of 2 and represents a compromise between the two. We think this is reasonable given that the model abstracts from wage rigidities. The parameter η that governs habit formation in consumption is calibrated to 0.7, a value close to those estimated in Smets and Wouters (2007) and Justiniano, Primiceri and Tambalotti (2010).

With respect to parameters governing the shocks, we abstract from average growth, setting $g = 0$. For the persistent shocks u_t and g_t , we take estimates from Cogley, Sargent and Primiceri (2009),

$$\begin{aligned} \rho_u &= 0.4, & 100\sigma_u &= 0.12, \\ \rho_g &= 0.27, & 100\sigma_g &= 0.5. \end{aligned} \quad (15)$$

For the white noise shocks ε_{yt} and $\varepsilon_{\pi t}$ we set

$$\sigma_\pi = \sigma_y = 0.01/4 \quad (16)$$

Finally, we adopt a standard calibration for loss-function parameters. We assume the central bank assigns equal weights to annualized inflation and the output gap. Since the model expresses inflation as a quarterly rate, this corresponds to $\lambda_y = 1/16$. We also set λ_i to 0.1, which implies that the weight on fluctuations of the annualized nominal interest rate is 10% of the weights attached to fluctuations in annualized inflation and the output gap. The results reported below for economies with learning are not sensitive to the choice of λ_i .

3 Learning about monetary policy

Private agents and the central bank know the model of the economy and the form of the policy rule, but private agents do not know the policy parameters. Instead, they must solve a signal-extraction problem to learn about them. If ψ entered linearly, this could be done via the Kalman filter. Because ψ enters non-linearly, however, agents must solve a nonlinear filtering problem. This section describes how that is done. We first conjecture a perceived law of motion (PLM) and then derive the actual law of motion (ALM) under the PLM. After that, we verify that the PLM is the perceived ALM. Having verified that private agents know the ALM up to the unknown policy parameters, we use the ALM to derive the likelihood function. Agents combine the likelihood with a prior over policy parameters and search numerically for the posterior mode.

3.1 The perceived law of motion

By stacking the IS equations (4) and (5) along with the aggregate supply block (equations 7-9), exogenous shocks, and perceived monetary policy rule (equation 2), the private sector's beliefs can be represented as a system of linear expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_t S_{t-1} + D_t \varepsilon_t, \quad (17)$$

where S_t is the model's state vector, ε_t is a vector of innovations, and A_t, B_t, C_t , and D_t are matrices that depend on the model's deep parameters (see the web appendix for details). These matrices vary with time because they depend on estimates of the policy coefficients ψ_t . We conjecture that the PLM is the reduced-form VAR associated with (17). Using standard arguments, the reduced form can be expressed as

$$S_t = F_t S_{t-1} + G_t \varepsilon_t, \quad (18)$$

where F_t solves $B_t F_t^2 - A_t F_t + C_t = 0$ and $G_t = A_t^{-1} D_t$.¹⁰

¹⁰As usual, a solution for F_t can be computed by solving a generalized eigenvalue problem (Uhlig 1999).

3.2 The actual law of motion

To find the actual law of motion, we stack the *IS* curve, aggregate supply block, and shocks along with the actual policy rule. This results in another system of expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_{at} S_{t-1} + D_t \varepsilon_t. \quad (19)$$

The state and innovation vectors are the same as in (17), as are the matrices A_t , B_t , and D_t . In addition, all rows of C_{at} agree with those of C_t except for the one corresponding to the monetary-policy rule. In that row, the true policy coefficients ψ replace the estimated coefficients ψ_t . See the web appendix for details.

According to the timing protocol, subjective forecasts are

$$E_t^* S_{t+1} = F_t^2 S_{t-1}. \quad (20)$$

After substituting subjective forecasts into (19), we obtain a system of backward-looking difference equations,

$$A_t S_t = (B_t F_t^2 + C_{at}) S_{t-1} + D_t \varepsilon_t. \quad (21)$$

Premultiplying both sides by A_t^{-1} delivers the ALM under the PLM,

$$S_t = H_t S_{t-1} + J_t \varepsilon_t, \quad (22)$$

where

$$\begin{aligned} H_t &= A_t^{-1} (B_t F_t^2 + C_{at}), \\ J_t &= A_t^{-1} D_t. \end{aligned} \quad (23)$$

The equilibrium law of motion is a *VAR* with time-varying parameters and conditional heteroskedasticity, as in Cogley and Sargent (2005) and Primiceri (2006). Private-sector beliefs about monetary policy matter for price-setting and consumption decisions. Those beliefs are encoded in A_t , B_t , D_t , F_t , and some elements of C_{at} . Outcomes also depend on actions taken by the central bank, which involve the actual policy rule. That is encoded in other elements of C_{at} . An intriguing feature of the equilibrium is that the drifting parameters ψ_t have a lower dimension than the conditional mean parameters $vec(H_t)$. This is qualitatively consistent with a finding of Cogley and Sargent (2005), who reported that drift in an analog to $vec(H_t)$ is confined to a lower dimensional subspace. The form of the conditional variance in (22) differs from their representations, however, so the model disagrees with their identifying restrictions. Another difference is that the model involves temporary drift during a learning transition while their VARs involve perpetual drift.

3.3 The perceived ALM is the PLM

To verify that the perceived ALM is the PLM, first notice that the reduced-form ALM and PLM are both $VAR(1)$ processes. The conditional variance matrices are the same, $G_t = J_t = A_t^{-1}D_t$, and the conditional mean matrices solve

$$\begin{aligned} B_t F_t^2 - A_t F_t + C_t &= 0, \\ B_t F_t^2 - A_t H_t + C_{at} &= 0. \end{aligned} \tag{24}$$

After subtracting one from the other and re-arranging terms, we find

$$H_t - F_t = A_t^{-1}(C_{at} - C_t). \tag{25}$$

By inspection, one can verify that C_{at} and C_t are identical except for the row corresponding to the monetary-policy rule. In that row, C_{at} depends on actual policy coefficients ψ while C_t depends on perceived policy coefficients ψ_t . If we were to interview the agents in the model and ask their view of the ALM, they would answer by replacing ψ in C_{at} with ψ_t , thus obtaining C_t . Since the perceived difference between H_t and F_t is zero, the perceived ALM is the PLM. This is true not only asymptotically but for every date during the transition.

Among other things, this implies that private agents know the ALM up to the unknown policy parameters. Hence they can use the ALM to form a likelihood function.¹¹

3.4 The likelihood function

We collect the observables in a vector $X_t = [\pi_t, u_t, y_t, g_t, i_t]'$. The other elements of the state vector allow us to express the model in first-order form but convey no additional information beyond that contained in the history of X_t . Using the prediction-error decomposition, the likelihood function for data through period t can be expressed as

$$p(X^t|\psi) = \prod_{j=1}^t p(X_j|X^{j-1}, \psi). \tag{26}$$

Since the private sector knows the ALM up to the unknown policy parameters, they can use it to evaluate the terms on the right-hand side of (26).

According to the ALM, X_t is conditionally normal with mean and variance¹²

$$\begin{aligned} m_{t|t-1}(\psi) &= e_X H_t(\psi) S_{t-1}, \\ V_{t|t-1}(\psi) &= e_X J_t V_\varepsilon(\psi) J_t' e_X', \end{aligned} \tag{27}$$

¹¹Another implication is that the private sector's forecasts are consistent with their contingency plans for the future. For instance, for $j > 0$, log-linear consumption Euler equations between periods $t+j$ and $t+j+1$ hold in expectation at t .

¹²According to the timing protocol, H_t and J_t can be regarded either as beginning-of-period t estimates or end-of-period $t-1$ estimates. That is why it is legitimate to use them to calculate the conditional mean and variance.

where e_X is a selection matrix defined in the web appendix, $H_t(\psi)$ is the ALM conditional-mean array evaluated at a particular value of ψ , and $V_\varepsilon(\psi)$ is the variance of the innovation vector ε_t also evaluated at ψ . It follows that the log-likelihood function is

$$\ln p(X^t|\psi) = -\frac{1}{2} \sum_{j=1}^t \left\{ \ln |V_{j|j-1}(\psi)| + [X_j - m_{j|j-1}(\psi)]' V_{j|j-1}^{-1}(\psi) [X_j - m_{j|j-1}(\psi)] \right\}. \quad (28)$$

3.5 The private sector's prior and posterior

Private agents have a prior $p(\psi)$ over the policy coefficients. At each date t , they find the log posterior kernel by summing the log likelihood and log prior. Because of our anticipated-utility assumption, their decisions depend only on a point estimate, not on the entire posterior distribution. Among the various point estimators from which they can choose, we assume they adopt the posterior mode,

$$\psi_t = \arg \max (\ln p(X^t|\psi) + \ln p(\psi)). \quad (29)$$

Agents take into account that past outcomes were influenced by past beliefs. They are not recursively estimating a conventional rational-expectations model. By inspecting the ALM and PLM, one can verify that past values of the conditional mean $m_{j|j-1}(\psi)$ and the conditional variance $V_{j|j-1}(\psi)$ depend on past estimates as well as the current candidate ψ . Past estimates are bygones at t and are held constant when agents update the posterior mode.¹³ The estimates are based not just on the policy rule but also on equations for inflation and output. The agents exploit all information about ψ , taking advantage of cross-equation restrictions implied by the ALM to sharpen estimates.

4 The central bank's decision problem

A new governor appears at date 0 and formulates a new policy rule. After observing the private sector's prior, the new governor chooses the long-run inflation target $\bar{\pi}$ and the reaction coefficients ψ_y, ψ_π to minimize expected loss, with the standard deviation of policy shocks σ_i being set exogenously. The disinflation commences at date 1.

4.1 Initial conditions

The economy is initialized at the steady state under the old regime. Because we are interested in a scenario like the end of the Great Inflation, we calibrate the old

¹³For that reason, whether the estimator can be expressed in a recursive form is unclear.

regime to match estimates of the policy rule for the period 1965-1979. We assume that the policy rule had the same functional form as (1) during that period, and we estimate $\bar{\pi}$, ψ_π , ψ_y , and σ_i^2 by OLS. The point estimates and standard errors are reported in table 1.

Table 1: The Old Regime

$\bar{\pi}$	ψ_π	ψ_y	σ_i
0.0116	0.043	0.12	0.0033
(0.013)	(0.08)	(0.04)	(0.01)

Note: Estimates of policy coefficients, 1965-1979. Standard errors in parentheses.

The point estimate for $\bar{\pi}$ is 0.0116, implying an annualized target-inflation rate of 4.6 percent. The reaction coefficients are both close to zero, with the output coefficient being slightly larger than the inflation coefficient. Policy shocks are large in magnitude and account for a substantial fraction of the total variation in nominal interest. Standard errors are large, especially for $\bar{\pi}$.

We initialize the state vector at the steady state associated with this policy rule. This implies $\pi_0 = 0.0116$, $y_0 = -0.0732$, and $i_0 = 0.0217$, where inflation and nominal interest are expressed as quarterly rates.

4.2 Evaluating expected loss and finding the optimal policy

If the model fell into the linear-quadratic class, the loss function could be evaluated and optimal policy computed using methods developed by Mertens (2009a, 2009b). The central bank has quadratic preferences, and many elements of the transition equation are linear, but learning introduces a nonlinear element. Since that element is essential, we prefer to retain it and use other methods for evaluating expected loss.

We proceed numerically. We start by specifying a grid of values for $\bar{\pi}$, ψ_π , and ψ_y . Then, for each node on the grid, we simulate 100 sample paths, updating private-sector estimates ψ_t by numerical maximization at each date. The sample paths are each 20 years long, and the terminal loss is set equal to zero, representing a decision maker with a long but finite horizon. We calculate realized loss along each sample path and then average realized loss across sample paths to find expected loss. The optimal rule among this family is the node with smallest expected loss.¹⁴

5 A full-information benchmark

To highlight the role of learning, we begin by describing the optimal policy under full information. When private agents know the new policy, the optimal simple rule

¹⁴An alternative procedure would be to write down a dynamic program and solve it numerically, as in Gaspar, Smets, and Vestin (2006, 2009). This is feasible in models with a low-dimensional state vector, but in our model it runs afoul of the curse of dimensionality.




Figure 1: Impulse responses when agents are fully informed

sets $\bar{\pi} = 0$, $\psi_{\pi} = 0.75$, and $\psi_y = 0.5$. Figure 1 portrays a disinflation under this policy. Recall that the economy is initialized in the steady state of the old regime and that the disinflation commences at date 1. The figure depicts responses of inflation, output, and nominal interest gaps, which are defined as deviations from the steady state of the new regime.¹⁵

The nominal interest rate rises at date 1, causing inflation to decline sharply and overshoot the new target. After that, inflation converges from below. This rolls back the price level, partially counteracting the effects of high past inflation. As Woodford (2003) explains, a partial rollback of the price level is a feature of optimal monetary policy under commitment. Intuitively, a credible commitment on the part of the central bank to roll back price increases restrains a firm's incentive to increase its price in the first place. The optimal simple rule under full information also has this feature.

The initial increase in the nominal interest rate causes the output gap to fall below zero. Since inflation and output growth are below target at date 1, the central bank cuts the interest rate at date 2, damping the output loss and initiating a recovery. Convergence to the new steady state is rapid, with inflation, output, and interest gaps closing within 2 years. After 8 periods, inflation is close to its new target, which is 4.6 percentage points below the old target. The cumulative loss in output is approximately 3.25 percent. The sacrifice ratio, defined as the cumulative loss in output divided by the change in target inflation, is approximately 0.7 percent. The reason why the sacrifice ratio is small under full information is that the model has no indexation. Although prices are sticky, the absence of indexation means that inflation

¹⁵Values at date 0 represent the difference between steady states of the old and new regimes. Inflation and nominal interest gaps coincide because the steady-state real interest rate is the same.

is weakly persistent. The absence of indexation also explains why the bank seeks a substantial rollback in the price level.

6 Optimal simple rules under learning

The nature of the transition and the optimal policy depend on the private sector's prior. For that reason, we examine a number of examples involving different priors that are designed to highlight the main features of the central bank's problem.

6.1 The private sector initially anticipates a continuation of the old regime

The first example illustrates what happens when agents initially anticipate a continuation of the old regime. Among our scenarios, this represents a worst case for the central bank. We calibrate the private sector's prior using the estimates of the policy rule for 1965-1979 reported in table 1. In this way, we ensure that it encodes information from the period leading up to the Volcker disinflation.

In particular, we assume that agents believe that policy parameters are independent a priori,

$$p(\psi) = p(\bar{\pi})p(\psi_{\pi})p(\psi_y)p(\sigma_i^2). \quad (30)$$

We also assume that they adopt truncated normal priors for $\bar{\pi}$, ψ_{π} , ψ_y and a gamma prior for σ_i^2 . For $\bar{\pi}$, ψ_{π} , ψ_y , the mean and standard deviation of an *untruncated* normal density are set equal to the numbers shown in table 1. We then truncate at zero to ensure non-negativity and renormalize so that the truncated prior integrates to unity. For σ_i^2 , hyperparameters are chosen so that the implied mode and standard deviation match the numbers in the table.

The results are shown in figure 2. Priors for ψ_{π} and ψ_y concentrate slightly to the right of zero, and little prior mass is assigned to values greater than 0.25. On the other hand, priors for $\bar{\pi}$ and σ_i^2 are spread out and assign non-negligible probability to a broad range of values.¹⁶ According to this specification, private agents are skeptical that the central bank will react aggressively to inflation or output, but they are open to persuasion about $\bar{\pi}$ and σ_i^2 . That private agents are prejudiced against large values of ψ_{π} and ψ_y is important for what follows.

Figure 3 portrays isoclines for expected loss as a function of $\bar{\pi}$, ψ_{π} , ψ_y , and σ_i^2 . Each panel involves a different setting for $\bar{\pi}$, ranging from 0 to 3 percent per annum. The feedback parameters ψ_{π} and ψ_y are shown on the horizontal and vertical axes, respectively. The standard deviation of the policy shock σ_i is held constant at 0.001 in all cases. Expected loss is normalized by dividing by loss under the optimal rule,

¹⁶ $\bar{\pi}$ and σ_i are measured in quarterly rates.




Figure 2: A prior based on the old regime

so that contour lines represent gross deviations from the optimum. The red and blue diamonds in the upper right and left panels depict the optimal policy under learning and full information, respectively.

Regions of low expected loss are concentrated in the southwest quadrant of each panel, near the prior mode for ψ_π and ψ_y . Expected loss increases rapidly as the feedback coefficients move away. Indeed, in the northeast quadrant of each panel, expected loss is 100 times greater than under the optimal policy. For this example, the policy that is optimal under full information lies in the high-loss region and is far from the optimum.

The reason why expected loss is sometimes so large is that the equilibrium law of motion can be a temporarily explosive process, i.e. one that is asymptotically stationary but which has explosive autoregressive roots during the transition. The agents in our model want to be on the stable manifold, but they don't know where it is. Their plans are based on F_t , but outcomes depend on H_t . Under conditions ensuring a unique nonexplosive solution to (17), the eigenvalues of F_t are on or inside the unit circle.¹⁷ The eigenvalues of H_t , however, can be explosive even when those of F_t are not. Thus, actions that would be stable under the PLM can be unstable under the ALM. The matrices H_t and F_t differ because of disagreement between the

¹⁷A unit eigenvalue is associated with the constant in the state vector.




Figure 3: Iso-loss contours for example 1

actual policy ψ and the perceived policy ψ_t (see equation 12). The eigenvalues of H_t are close to those of F_t (and are nonexplosive) when ψ_t is close to ψ . Explosive eigenvalues emerge when there is substantial disagreement between ψ_t and ψ . On almost all simulated paths, the private sector eventually learns enough about ψ to make explosive eigenvalues vanish, but the transition is highly volatile and dominates expected loss when the initial disagreement is large and/or learning is slow.

The gray shaded areas in figure 4 depict regions of the policy-coefficient space for which the eigenvalues of H_1 are nonexplosive.¹⁸ The figure is formatted in the same way as figure 3. The nonexplosive region is similar for all settings of $\bar{\pi}$, but it is sensitive to ψ_π and ψ_y , concentrating near their prior mode. It follows that the emergence of explosive roots depends more on the feedback parameters than the long-run inflation target. The central bank can move its inflation target far from the private sector's prior mode without generating locally explosive dynamics, but moving ψ_π and/or ψ_y far from their prior modes can make the transition turbulent.

The central bank's main challenge is to cope with the potential for explosive volatility. One way to cope is to adopt a policy that is close to the private sector's prior. In particular, by setting ψ_π and ψ_y sufficiently close to their prior modes, the bank could ensure that H_t is nonexplosive throughout the transition. More ambitious changes in policy can succeed provided that the private sector learns quickly enough.

¹⁸The jagged boundary is due to the coarseness of our grid.




Figure 4: Nonexplosive region of H_1

The eigenvalues of H_t might then be explosive for a few periods before becoming nonexplosive. In this way, the central bank could attempt a more ambitious reform that achieves better long-run performance at the cost of higher short-term volatility.

In our examples, the central bank always chooses a policy that is initially outside the nonexplosive region. How far outside depends on the speed of learning. The more rapidly the private sector learns, the farther outside the central bank can reach. It is critical, however, that private beliefs catch up quickly. Otherwise the ALM will remain locally explosive for too long, and transitional volatility will be too high.

The speed of learning depends in part on the private sector's prior. In this example, the private sector is strongly prejudiced against large values of ψ_π and ψ_y . If the bank were to reach far outside the nonexplosive region in figure 4, it would have to fight against that prejudice, and learning would be too slow. For that reason, the optimal policy puts ψ_π and ψ_y only slightly outside. The bank can adjust $\bar{\pi}$ more freely, however, thereby achieving low average inflation.

The optimal simple rule for this example sets $\bar{\pi} = 1$ percent per year, $\psi_\pi = 0.15$, and $\psi_y = 0.3$. Figures 5 and 6 portray outcomes under this policy. Figure 5 plots mean responses of inflation, output, and nominal interest gaps, averaged across 100 sample paths. Figure 6 portrays mean estimates of the policy coefficients, again averaged across 100 sample paths. The true coefficients are shown in red while average estimates are portrayed in blue.




Figure 5: Average responses under the optimal policy




Figure 6: Average estimates under the optimal policy

As shown in figure 5, the transition is longer and more volatile than under full information. Inflation again declines sharply at impact, overshooting $\bar{\pi}$ and partially rolling back past increases in the price level. But the response is greater in amplitude under learning, and inflation oscillates as it converges to its new long-run target. The transition now takes about five years, with inflation remaining below target for most of that time. There is also a shallow but long-lasting decline in output. The output gap reaches a trough of -1.75 percent in quarter 3 and remains negative for four years. The cumulative output gap during this time is -10.25 percent. Since inflation fell permanently by 3.6 percentage points, the sacrifice ratio amounts to 2.8 percent of lost output per percentage point of inflation. The sacrifice ratio is four times larger than under full information, and it is in the ballpark of estimates for the Volcker disinflation.¹⁹

As shown in figure 6, estimates of ψ_π converge to its true value after one year. Rapid convergence of ψ_π is crucial for eliminating locally-explosive dynamics. In this case, beliefs about ψ_π converge rapidly because they don't have far to go. The bank set ψ_π close to its prior mode precisely so that disagreement would not persist. Estimates of $\bar{\pi}$ converge to the true value within 10 quarters, thus centering long-run inflation forecasts near the bank's actual target. Learning about ψ_y and σ_i is slower but also less critical.

To illustrate why a more ambitious reform is suboptimal, we examine an alternative policy that holds $\bar{\pi}$ and ψ_y constant but which reacts more aggressively to inflation, increasing ψ_π from 0.15 to 0.45. This policy is located to the right of the optimum in figures 3 and 4. Figures 7 and 8 depict average outcomes under this rule.

Under this policy, the central bank is fighting against the private sector's prior, which assigns low probability to neighborhoods of the true values $\psi_\pi = 0.45$ and $\psi_y = 0.3$. It follows that a lot of sample information is needed to overcome the prior. For the sake of intuition, imagine that agents were estimating the policy rule by running a regression. Because the prior assigns low weight to neighborhoods of the true feedback coefficients, the likelihood function would have to concentrate sharply in order to move the posterior there. For that to happen quickly, the right-hand variables in the regression (inflation and output growth) would have to be highly volatile. The bank can create a lot of volatility (see figure 7), and those fluctuations do help the private sector learn (see figure 8). On balance, however, the long-run benefits associated with a stronger reaction to inflation do not justify the higher transitional costs.

¹⁹For instance, Mankiw (2010, p. 398) reports a back-of-the-envelope estimate of 2.8 percent.




Figure 7: Average responses when $\bar{\pi} = 0.01$, $\psi_{\pi} = 0.45$, and $\psi_y = 0.3$




Figure 8: Average estimates when $\bar{\pi} = 0.01$, $\psi_{\pi} = 0.45$, and $\psi_y = 0.3$

6.2 The inflation target is transparent, but other aspects of policy are opaque

A number of central banks take pains to provide clear and credible signals about their inflation targets, but other aspects of policy are often more opaque. The previous results suggest that this will do little to address the bank’s main challenge – coping with high transitional volatility – because transitional volatility depends more on disagreement about the feedback parameters and opacity about ψ_π and ψ_y leaves this in place.

We study this more closely in the web appendix. There we examine a scenario in which the bank makes a transparent and credible announcement about $\bar{\pi}$ at date 0.²⁰ Armed with a credible signal, private agents adopt a prior for $\bar{\pi}$ that is tightly centered on the bank’s announcement. To represent opacity about the other policy coefficients, we assume the bank says nothing about ψ_π , ψ_y , or σ_i , and we adopt the same priors as in figure 2.

As expected, transparency about target inflation accelerates learning about $\bar{\pi}$, but it has little effect on beliefs about the other policy coefficients. As a consequence, regions of low expected loss still concentrate near the prior modes for ψ_π and ψ_y , and temporarily explosive paths still emerge when ψ_π and/or ψ_y deviate too much from prior beliefs. For managing transitional volatility, a clear and credible signal about the feedback parameters would much more helpful.

6.3 Mixture priors

Next we return to example 1 and perturb the prior so that agents are less skeptical about the prospects for change. We imagine that they enter date 0 with beliefs about the old regime $p_{old}(\psi)$ that are the same as in example 1. But instead of assigning probability 1 to that prior, we assume that they expect the central bank to continue the old regime with probability $1 - w$ and to switch to something else with probability w . Their beliefs about a new regime are encoded in a conditional prior $p_{new}(\psi)$. The marginal prior is a mixture of the two conditional priors,

$$p_m(\psi) = (1 - w)p_{old}(\psi) + wp_{new}(\psi), \quad (31)$$

where w measures the public’s skepticism about the prospects for change. A value close to 0 means that the public is highly skeptical and weighs its past experience heavily, while a value close to 1 means that the private sector heavily discounts the past and looks forward to something new. The previous example set $w = 0$. Now we turn to examples in which $w = 0.1$, 0.3 , and 0.5 , respectively.

²⁰We sidestep difficult and important questions about how the bank makes its signal credible and simply assume that it does.




Figure 9: Mixture priors

For $p_{new}(\psi)$, we adopt the same functional forms as for $p_{old}(\psi)$, and we calibrate it so that it is loosely centered on policies that would work well under full information. The details are recorded in table 2. Abstracting from the truncation at zero, the conditional prior mean for $\bar{\pi}$ is 2 percent per year, and a conditional 95 percent confidence band ranges from 0 to 4 percent. Similarly, the conditional prior means for ψ_{π} and ψ_y are 0.5 with conditional confidence bands of plus or minus 0.5. The conditional mode for σ_i is 0.001, and its standard deviation is 0.001.

Table 2: Conditional Prior, $p_{new}(\psi)$

	$\bar{\pi}$	ψ_{π}	ψ_y	σ_i
Mean	0.005	0.5	0.5	0.001
Standard deviation	0.0025	0.25	0.25	0.001

The resulting mixture priors are depicted in figure 9. The components $p_{old}(\psi)$ and $p_{new}(\psi)$ are shown in blue and green, respectively, and mixtures for various values of w are shown in red. For $w = 0.1$, $p_m(\psi)$ is broadly similar to $p_{old}(\psi)$, reflecting that $p_{old}(\psi)$ has a weight of 0.9. However, $p_m(\psi)$ differs in two respects that matter. First, the prior mode for $\bar{\pi}$ is shifted to the left, near the mode for $p_{new}(\bar{\pi})$. Second, the upper tails for the feedback parameters are fatter than those of $p_{old}(\psi)$. For $w = 0.1$, $p_m(\psi)$ still assigns small probabilities to large values of ψ_{π} or ψ_y , but those

probabilities are orders of magnitude larger than under $p_{old}(\psi)$. Thus, although agents remain skeptical that the bank will react aggressively to inflation or output growth, they are less strongly prejudiced against that possibility.

These changes become more pronounced as w increases, but even for $w = 0.5$ the prior modes for ψ_π and ψ_y are the same as in the first example. Hence the initial disagreement about feedback parameters remains the same. Again, what differs is the shape of the upper tails. The priors for ψ_π and ψ_y are less tightly concentrated on old-regime values and assign greater probability to a broad range of possibilities. Because the priors for ψ_π and ψ_y are less tightly concentrated on old-regime values, incoming data are weighed more heavily and agents learn faster, allowing the central bank to move farther away from the prior mode.

Since the initial disagreement about feedback parameters is the same, so is the region of the policy-coefficient space for which the ALM is initially nonexplosive. The gray shaded area in figure 10 depicts this region. Because learning is more rapid, however, locally-explosive dynamics vanish more quickly, enabling the bank to implement more ambitious reforms. Figure 10 also portrays the optimal policy for each prior. The optimum is near the edge of the nonexplosive region when $w = 0$, it moves farther outside as w increases, and it begins to approach the full-information optimum when $w = 0.5$.




Figure 10: Nonexplosive region for H_1

Figures 11-13 depict isoloss contours for the three mixture priors. The chief dif-




Figure 11: Isoloss contours for a mixture prior with $w = 0.1$




Figure 12: Isoloss contours for a mixture prior with $w = 0.3$




Figure 13: Isoloss contours for a mixture prior with $w = 0.5$

ference relative to the case of $w = 0$ is that the expected loss surface is much flatter. Because agents learn faster, the economy is more tolerant of policies far from the prior mode. In each instance there exists a wide range of policies surrounding the optimum that are almost as good. Expected loss still rises as ψ_π moves to the right of the optimum, but it increases by a factor of 2.5 or less, not by a factor of 100. No catastrophes emerge in these simulations.

Figures 14 illustrates the speed of learning under the respective optimal policies. As before, the true policy coefficients are shown in red while ensemble averages of posterior estimates are depicted in blue. The private sector learns quickly, with estimates converging to neighborhoods of the true coefficients within one year. In fact, when $w = 0.5$, the estimates converge after a single quarter. Learning is faster because the private sector's prior assigns greater probability to neighborhoods of the new policies. As a consequence, less sample information is needed to identify the new rule.

Finally, figure 15 portrays average responses of inflation, output, and nominal interest gaps under the respective optimal policies. Although the policy coefficients differ across priors, the impulse response functions are similar. As before, the initial disagreement between actual and perceived policies makes the ALM locally explosive. The initial responses are therefore large in magnitude. The disagreement vanishes quickly, however, and the ALM becomes locally stable, causing volatility to decline.




Figure 14: Average estimates of policy coefficients




Figure 15: Average responses for various mixture priors

Inflation again falls sharply at impact, overshoots the new long-run target, and converges from below. Convergence is rapid, with gaps vanishing within 6-8 quarters. A mild recession occurs, and sacrifice ratios range from 0.7 to 1.6 percent of lost output per percentage point of inflation. As one would expect, the sacrifice ratio depends on the speed of learning and is not an invariant feature of the economy.

Two aspects of these examples are noteworthy. The first is that although fast learning is essential for an ambitious reform, full credibility is not. By ‘full credibility’ we mean that the private sector’s prior is tightly concentrated on the optimum. This would obviously promote fast learning, but it is not necessary. If private agents are sufficiently open minded about potential reforms (i.e. their priors assign nontrivial mass to a range of policies), they will be able to learn quickly, and that is enough for an ambitious reform to succeed. Notice in particular that the optimal policy under learning approaches the full-information optimum when $w = 0.5$. What is interesting is that this happens for a mixture weight well below unity. A weight close to 1 is not needed.

A second and related point is that outcomes depend on subtle features of the private sector’s prior. The prior modes for ψ_π and ψ_y are the same as when $w = 0$. What differs is the shape of the tails. That the upper tails are fatter accelerates learning and allows the central bank to react more aggressively to inflation and output growth. It follows that knowing the prior mode would not be enough for the central bank to choose between the policies recommended here and in example 1. The bank would also have to know the shape of the tails. As w declines toward 0, the economy becomes less fault tolerant, making it more important for the bank to know the entire shape of the prior.

6.4 Comparing the examples

Table 3 summarizes our results, showing the optimal simple rule for each example and reporting expected loss relative to what could be achieved under full information. In the first example, the bank’s concerns about explosive volatility limit its adjustment of ψ_π , and expected loss is more than twice as high as under full information. Worries about explosive volatility are less pressing in the other examples, but expected loss remains 1.6 to 2 times higher than under full information because learning amplifies initial volatility. The more rapidly agents learn, the sooner the initial period of high volatility ends. Thus expected loss declines with w .

Optimal policy coefficients are sensitive to the prior. However, for the examples considered here, a number of robust lessons emerge. First and foremost, the transition dominates expected loss, and what matters most is how well the bank manages transitional volatility. However, the bank can adjust $\bar{\pi}$ freely without triggering explosive volatility. In most cases, it stops short of zero – the optimum under full information²¹

²¹We abstract from the zero lower bound on nominal interest.

– because the transitional cost would be too high, but it always seeks low average inflation.

Table 3: Optimal Simple Rules

	$\bar{\pi}$	ψ_{π}	ψ_y	Relative Loss	Sacrifice Ratio
Full information	0	0.75	0.5	1	0.7 percent
$w = 0$	0.01	0.15	0.3	2.2	2.8 percent
$w = 0.1$	0.01	0.25	0.5	2.0	1.6 percent
$w = 0.3$	0	0.45	1.1	1.6	1.3 percent
$w = 0.5$	0	0.55	0.5	1.6	0.7 percent

Note: $\sigma_i = 0.001$ in all cases.

Another robust feature is that the central bank always accepts an initial period of high volatility in exchange for better long-run performance. Optimal policies create locally explosive dynamics at the beginning of the transition, which vanish as agents learn. The bank must limit itself to policies that can be learned quickly, however, because the transition would be too costly otherwise.

Third, the optimal policy pulls inflation down sharply at impact, overshooting the new long-run inflation target and partially rolling back past increases in the price level. After that, inflation converges to its new target from below. None of our examples rationalize a gradual reduction in inflation.

Finally, output losses are relatively mild. The sacrifice ratio depends on the private sector’s prior, but in our examples it never exceeds 2.8 percent, and in some cases it is quite a bit less.

7 Conclusion

This paper models the transitional dynamics that emerge after the adoption of a new monetary-policy rule. We assume that private agents must learn about the new policy, and we study how learning affects the nature of the transition and choice of a new rule. Among other things, we derive the relationship between the perceived and actual laws of motion and show that the PLM is the perceived ALM at every date. Our equilibrium therefore lies between that of a conventional rational-expectations model in which the ALM and PLM always coincide and that of a least-squares learning model in which the ALM and PLM might converge, but only asymptotically. In addition, we characterize cross-equation restrictions in a learning equilibrium and show that they are similar but not identical to those arising in a conventional rational-expectations model. We also demonstrate that the ALM can be a temporarily explosive process, i.e. one that is asymptotically stationary but which has autoregressive roots greater

than 1 in magnitude during the transition. The transition is highly volatile when explosive roots emerge.

The central bank's main challenge is to cope with the potential for explosive volatility. Explosive roots emerge when there is too much disagreement between actual and perceived policies. For some priors, the bank's optimal strategy is to adopt an incremental reform that limits the initial disagreement between actual and perceived policies. More ambitious reforms can succeed when priors permit agents to learn fast enough. Examples are given of each. A conventional rational-expectations analysis can be misleading because rules that work well when the new policy is known with certainty sometimes generate a lot of volatility under learning. Because the optimal policy is sensitive to the private sector's prior, we cannot give an unequivocal answer to the policy-design question. For that reason, we view our contribution more as a suggestion about how to analyze the problem than as a definitive description of the optimal strategy.

In some respects, the analysis points toward 'open-mouth' policies for affecting the private sector's prior (e.g. Guthrie and Wright 2000). Our model is capable of analyzing the consequences of alternative priors, but we do not have a theory of how priors are determined or how a central bank might influence them. Our approach would have to be extended in that direction to analyze open-mouth policies. Nevertheless, our results suggest that transparency with respect to feedback parameters is more important than transparency with respect to the inflation target. The bank can achieve low average inflation without too turbulent a transition even when priors on $\bar{\pi}$ are diffuse. Furthermore, transitional volatility depends more on uncertainty about feedback parameters, and transparency with respect to $\bar{\pi}$ matters little if ψ_{π} and ψ_y remain opaque.

Our analysis also leaves open a number of other interesting questions. For instance, incorporating the zero lower bound on nominal interest is important for analyzing policies in which target inflation is close to zero. Since outcomes depend on subtle features of the private sector's prior, thinking about how to design surveys which elicit that information is important. Alternatively, withdrawing the assumption that the central bank can observe the private sector's prior would be interesting. Incorporating real-time data noise is relevant and would affect the private sector's signal-extraction problem. Last but not least, we would eventually like to consider models in which the central bank chooses sequentially, and the private sector learns about a moving target.

References

- [1] Ascari, G., 2004. Staggered prices and trend inflation: some nuisances. *Review of Economic Dynamics* 7, 642–667.

- [2] Calvo, G. 1983. Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics* 12, 383-398.
- [3] Cogley, T., G. Primiceri, and T.J. Sargent. 2010. Inflation-Gap Persistence in the U.S. *American Economics Journal - Macroeconomics* 2, 43-69.
- [4] Cogley, T. and T.J. Sargent. 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262-302.
- [5] Cogley, T. and T.J. Sargent, 2008. Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review* 49, 185-221.
- [6] Cogley, T. and A.M. Sbordone. 2008. Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve. *American Economic Review* 98, 2101-2126.
- [7] Coibion, O. and Y. Gorodnichenko. 2008. Monetary Policy, Trend Inflation and the Great Moderation: An Alternative Interpretation. *NBER wp 14621*.
- [8] Erceg, C. and Levin, A. 2003. Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50 (4): 915-944.
- [9] Evans, G.W. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press: Princeton, N.J.
- [10] Gaspar, Vito, Frank Smets and David Vestin. 2006. Adaptive Learning, Persistence, and Optimal Monetary Policy. *Journal of the European Economic Association* 4, 376-385.
- [11] Gaspar, Vito, Frank Smets and David Vestin. 2009. Inflation Expectations, Adaptive Learning, and Optimal Monetary Policy. Unpublished manuscript, European Central Bank.
- [12] Goodfriend, M. and R.G. King, 2005. The Incredible Volcker Disinflation. *Journal of Monetary Economics*, 52 (5): 981-1016.
- [13] Gordon, R.J. 1982. Why Stopping Inflation May Be Costly: Evidence from Fourteen Historical Episodes, in *Inflation: Causes and Effects*, edited by Robert Hall, University of Chicago press, 11-40.
- [14] Guthrie, G. and J. Wright. 2000. Open mouth operations. *Journal of Monetary Economics*, 46(2): 489-516.
- [15] Justiniano, A. and G. Primiceri and A. Tambalotti. 2010. Investment Shocks and Business Cycles. *Journal of Monetary Economics*, 57(2): 132-145.

- [16] Kreps, D. 1998. Anticipated Utility and Dynamic Choice, in D.P. Jacobs, E. Kalai, and M. Kamien, eds., *Frontiers of Research in Economic Theory*, (Cambridge: Cambridge University Press, 1998), 242–74.
- [17] Mankiw, N.G. 2010. *Macroeconomics*, 7th edition, (New York :Worth Publishers).
- [18] Marcet, A. and T.J. Sargent. 1989a. Convergence of least-squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337-368.
- [19] Marcet, A. and T.J. Sargent. 1989b. Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy* 97, 1306-1322.
- [20] McCallum, B.T., 1999. Issues in the design of monetary policy rules. In: Taylor, J.B., Woodford, M. (Eds.), In: *Handbook of Macroeconomics*, vol. 1C. Elsevier, Amsterdam.
- [21] Mertens, Elmar. 2009a. Managing Beliefs about Monetary Policy under Discretion. Unpublished manuscript, Federal Reserve Board.
- [22] Mertens, Elmar. 2009b. Discreet Commitments and Discretion of Policymakers with Private Information. Unpublished manuscript, Federal Reserve Board.
- [23] Milani, Fabio. 2007. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54 (7): 2065-82
- [24] Okun, A.M., 1978. Efficient disinflationary policies. *American Economic Review Papers and Proceedings* 68, 348–352.
- [25] Orphanides, A. and J.C. Williams, 2002. Robust monetary policy rules with unknown natural rate, *Brookings Papers on Economic Activity*, 2002:2, pp. 63-145.
- [26] Orphanides, A. and J.C. Williams, 2005. Imperfect Knowledge, Inflation Expectations and Monetary Policy, in B. Bernanke and M. Woodford eds, *The Inflation Targeting Debate*
- [27] Orphanides, A. and J.C. Williams, 2007. Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54: 1406-1435.
- [28] Primiceri, G. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72, 821-852.

- [29] Sargent, T.J. 1982. The Ends of Four Big Inflations, in *Inflation: Causes and Effects*, edited by Robert Hall, University of Chicago press, pp. 41-97.
- [30] Sbordone, A.M. 2007. Inflation persistence: Alternative interpretations and policy implications. *Journal of Monetary Economics* 54, 1311–1339
- [31] Schmitt-Grohe, S. and M. Uribe. 2006. Optimal Simple and Implementable Monetary and Fiscal Rules: Expanded version, NBER wp 12402.
- [32] Schmitt-Grohe, S. and M. Uribe. 2007. Optimal Simple and Implementable Monetary and Fiscal Rules. *Journal of Monetary Economics* 54 (6): 1702-1725.
- [33] Smets, F. and R. Wouters. 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3): 586-606.
- [34] Uhlig, H. 1999. A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily. In *Computational Methods for the Study of Dynamic Economies*, edited by R. Marimon and A. Scott, Oxford University Press.
- [35] Woodford, M. 1999. Optimal Monetary Policy Inertia, NBER wp 7261.
- [36] Woodford, M. 2003. Interest and Prices. Princeton University Press: Princeton NJ.
- [37] Yun, Tack, 2005. Optimal Monetary Policy with Relative Price Distortion. *American Economic Review*, vol 95 (1): 89-109.

Appendices

(Not Intended for Publication)

A The model

The demand side of the model consists of equilibrium conditions of a representative household for the optimal choice of consumption and hours of work. The household maximizes expected discounted utility

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\log (C_t - \eta C_{t-1}) - \chi \frac{H_t^{1+\nu}}{1+\nu} \right), \quad (32)$$

subject to a flow budget constraint

$$E_t (Q_{t,t+1} Z_{t+1}) + P_t C_t = Z_t + W_t H_t + \int_0^1 \Psi_t(i) di, \quad (33)$$

where η measures a degree of internal habit persistence, β is a subjective discount factor and C_t is consumption of the final good, with price P_t . The specification of the period utility – separable in consumption and hours and logarithmic in consumption – guarantees the existence of a balanced growth path. The variable

$$H_t = \int h_t(i) di \quad (34)$$

is an aggregate of the number of hours supplied by the household to firms in the intermediate-goods sector, and W_t is the economy-wide nominal wage. Intermediate goods producers earn profits amounting to $\int_0^1 \Psi_t(i) di$, which they rebate to the household. The variable Z_{t+1} is the state-contingent value of the portfolio of assets held by the household at the beginning of period $t+1$, and $Q_{t,t+1}$ is a stochastic discount factor.

The marginal utility of consumption Ξ_t is

$$\Xi_t = \frac{1}{C_t - \eta C_{t-1}} - \beta \eta E_t \frac{1}{C_{t+1} - \eta C_t}. \quad (35)$$

and the first order condition for the choice of consumption is

$$\Xi_t = \beta E_t \left[\Xi_{t+1} \frac{R_t}{\Pi_{t+1}} \right]. \quad (36)$$

where $R_t = [E_t(Q_{t,t+1})]^{-1}$ is the gross nominal interest rate, and Π_t is the gross inflation rate: $\Pi_t = P_t/P_{t-1}$. The first order-conditions for labor supply is

$$w_t = \chi H_t^\nu C_t, \quad (37)$$

were $w_t \equiv W_t/P_t$ is the real wage. Because there is no capital or government, the aggregate resource constraint is simply $C_t = Y_t$.

To eliminate the non stationarity induced by the technological progress Γ_t we express (35) and the equilibrium condition (36) as²² respectively

$$\Xi_t^a = \frac{1}{C_t^a - \eta C_{t-1}^a \frac{\Gamma_{t-1}}{\Gamma_t}} - \beta \eta E_t \frac{1}{C_{t+1}^a \frac{\Gamma_{t+1}}{\Gamma_t} - \eta C_t^a}, \quad (38)$$

and

$$\Xi_t^a = \beta E_t \left[\Xi_{t+1}^a \left(\frac{\Gamma_{t+1}}{\Gamma_t} \right)^{-1} \frac{R_t}{\Pi_{t+1}} \right]. \quad (39)$$

where $C_t^a \equiv C_t/\Gamma_t$, $Y_t^a \equiv Y_t/\Gamma_t$, $\gamma_t = \Gamma_t/\Gamma_{t-1}$, and $\Xi_t^a \equiv \Xi_t \Gamma_t$. Similarly, the first order condition for labor supply can be written as

$$w_t^a = \chi H_t^\nu Y_t^a. \quad (40)$$

where $w_t^a \equiv w_t/\Gamma_t$ is the productivity adjusted real wage, and the aggregate resource constraint becomes $C_t^a = Y_t^a$. Further, imposing market clearing, equation (38) can be rewritten as

$$\Xi_t^a = \frac{1}{Y_t^a - \eta Y_{t-1}^a \frac{\Gamma_{t-1}}{\Gamma_t}} - \beta \eta E_t \frac{1}{Y_{t+1}^a \frac{\Gamma_{t+1}}{\Gamma_t} - \eta Y_t^a}. \quad (41)$$

The log-linearization of the equilibrium condition (41) and (39) gives the dynamic *IS* block. To derive these equations, we need a further transformation of variables to eliminate the non-stationarity induced by trend inflation, so we define $\tilde{\Xi}_t = \Xi_t^a Y_t^a$ and $\tilde{Y}_t^a = Y_t^a/\bar{Y}_t^a$ (with steady state values $\tilde{\Xi}^{ss} = \frac{\gamma - \beta\eta}{\gamma - \eta}$ and $\tilde{Y}^{a,ss} = 1$). With this transformation of variables, we can log-linearize the *IS* equation to obtain

$$\hat{\Xi}_t = \hat{Y}_t^a + E_t \left(\hat{\Xi}_{t+1} - \hat{Y}_{t+1}^a - \hat{\gamma}_{t+1} - \hat{\gamma}_{y,t+1} + i_t - \pi_{t+1} - r \right) \quad (42)$$

where $\hat{\Xi}_t \equiv \log \tilde{\Xi}_t / \tilde{\Xi}^{ss}$ is defined as follows

$$\hat{\Xi}_t = \xi_1 \hat{Y}_t^a + \xi_2 \left[\left(\hat{Y}_{t-1}^a - \hat{\gamma}_t \right) + \beta E_t \left(\hat{Y}_{t+1}^a + \hat{\gamma}_{t+1} + \hat{\gamma}_{y,t+1} \right) \right]. \quad (43)$$

The hat variables are, as usual, log deviations from steady state: $\hat{Y}_t^a = \log Y_t^a - \log \bar{Y}_t^a$, $i_t = \log R_t$, $\hat{\gamma}_t = \log \gamma_t - \log \gamma$, $i_t = \log R_t$, $\gamma_{yt} \equiv \bar{Y}_t^a / \bar{Y}_{t-1}^a$, and r is the steady state real interest rate.²³

²²The nominal interest rate is affected by the non-stationarity of inflation, but its ratio to inflation (and trend inflation) is stationary.

²³Note the term R_t/Π_{t+1} is stationary, and we denote its (log) steady state (which is equal to the steady state value of the ratio of nominal interest rate to trend inflation) by r . This can be seen by dividing through by $\bar{\Pi}_t$, which gives $\frac{R_t/\bar{\Pi}_t}{(\Pi_{t+1}/\bar{\Pi}_{t+1})(\bar{\Pi}_{t+1}/\bar{\Pi}_t)} = \frac{R_t^{\pi}}{\bar{\Pi}_t g_{t+1}^{\pi}}$ whose steady state we denote by $R^{r\pi}$, and $r \equiv \log R^{r\pi}$.

Equations (4) and (5) in the main text are a transformation of (42) and (43), where we adopt a simplified notation, setting $\xi_t \equiv \log \tilde{\Xi}_t$, $\xi \equiv \log \tilde{\Xi}^{ss}$, $y_t \equiv \log Y_t^a$, $\bar{y}_t \equiv \log \bar{Y}_t^a$ and $\bar{y}_{t-1} \equiv E_t^* \bar{y}_t$, $g_t \equiv \ln \gamma_t$ and $g \equiv \ln \gamma$. Furthermore, we replace rational expectations by learning and trend inflation by agents' perception with date $t-1$ information, $\bar{\pi}_{t-1}$. Finally, since $E_t^* \hat{\gamma}_{y_{t+1}} = 0$, the term is suppressed. All steady state variables which are functions of trend inflation are similarly denoted with an overbar and subscript $t-1$.

The equilibrium condition (40) is used in the next section to substitute out the real wage in the marginal cost expression of the supply side of the model.

A.1 The supply side

The supply side of the model consists of equilibrium conditions for a continuum of monopolistically competitive firms that produce intermediate goods and a final-good aggregating firm. These equilibrium conditions determine the dynamics of inflation in the model.

The final-good producer combines $y_t(i)$ units of each intermediate good i to produce Y_t units of the final good with technology

$$Y_t = \left[\int_0^1 y_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}. \quad (44)$$

where θ is the elasticity of substitution across intermediate goods. The final good producer chooses the intermediate inputs to maximize its profits, taking the price of the final good P_t as given, determining demand schedules

$$y_t(i) = Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta}. \quad (45)$$

The zero-profit condition then determines the aggregate price level

$$P_t \equiv \left[\int_0^1 p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}}. \quad (46)$$

Intermediate firm i hires $h_t(i)$ units of labor of type i on an economy-wide competitive market to produce $y_t(i)$ units of intermediate good i with technology

$$y_t(i) = \Gamma_t h_t(i), \quad (47)$$

where Γ_t is an aggregate technological process.

Firms can reset prices at random intervals (we assume a Calvo price-setting mechanism), and we denote by $1 - \alpha$ the probability that an intermediate-goods producer

has an opportunity to reset its price. The first order conditions of the optimal price-setting problem²⁴ and the evolution of aggregate prices jointly determine the dynamics of inflation in the model.

In log-linear form, the supply side can be described by a pair of equations, known as a new Keynesian Phillips curve²⁵

$$\begin{aligned}\pi_t - \bar{\pi}_{t-1} &= \tilde{\kappa}_{t-1}(mc_t - \bar{mc}_{t-1}) + \beta_{t-1}E_t^*(\pi_{t+1} - \bar{\pi}_{t-1}) \\ &\quad + \gamma_{1t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}] + u_t + \varepsilon_{\pi t}, \\ \phi_t &= \gamma_{2t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}],\end{aligned}\tag{48}$$

where $\tilde{\kappa}_{t-1} = \kappa_{t-1}/(1 + \nu)$ and the other parameters are defined in expression (11) in the main text.

A.1.1 Marginal costs, output and price dispersion

In order to write the NKPC as a relation between inflation and output, we solve for marginal cost mc_t as function of aggregate output. Given the production function (47), the real marginal cost for the optimizing firm i (which is setting a new price at time t) is

$$mc_t(i) = \frac{W_t/P_t}{\Gamma_t} = w_t^a = mc_t,\tag{49}$$

where the last equality follows from the fact that the marginal cost doesn't depend on firm's i variables.

From equilibrium condition (40) w_t^a is function of aggregate hours, which are obtained by aggregating hours worked in each firm:

$$H_t \equiv \int_0^1 h_t(i) di = \int_0^1 \frac{y_t(i)}{\Gamma_t} di = Y_t^a \int_0^1 \left(\frac{p_t(i)}{P_t} \right)^{-\theta} di = Y_t^a \Delta_t,\tag{50}$$

where we denoted by Δ_t the following measure of price dispersion: $\Delta_t \equiv \int_0^1 \left(\frac{p_t(i)}{P_t} \right)^{-\theta} di$. Alternatively, we can write aggregate output as function of aggregate hours and a measure of price dispersion:

$$Y_t^a = \frac{H_t}{\Delta_t}.\tag{51}$$

Δ_t measures the resource cost induced by price dispersion in the Calvo model, in equilibrium. One can show that $\Delta_t \geq 1$, which implies that in equilibrium more

²⁴For simplicity we assume away wedges between the individual firm marginal cost and aggregate marginal costs.

²⁵For further detail on the derivation of this Phillips curve, see Cogley and Sbordone (2008). The curve here in that the transformation of inflation into a stationary variable is obtained by dividing current (gross) inflation by perceived (rather than actual) trend inflation. The log-linearization is therefore defined around a point where perceived trend inflation and actual inflation are the same.

hours are needed to produce the same amount of output (indeed, labor productivity is the inverse of the price dispersion index.) Price dispersion is therefore always a costly distortion in this model. By substituting expressions (40) and (50) in (49) we get

$$mc_t = \chi H_t^\nu Y_t^a = \chi \Delta_t^\nu (Y_t^a)^{1+\nu}. \quad (52)$$

This expression shows that price dispersion creates a wedge between marginal costs and output. Substituting out marginal cost, we derive below the log-linear NKPC where both aggregate output and price dispersion are driving variables. Before doing that, we discuss the values of the variables in steady state.

A.1.2 Steady-state relations

From the definition of Δ_t we can derive that²⁶

$$\Delta_t = (1 - \alpha) (\tilde{p}_t)^{-\theta} + \alpha \Pi_t^\theta \Delta_{t-1},$$

where for ease of notation we indicate by $\tilde{p}_t$ the relative price of the firms that optimizes at t : $\tilde{p}_t \equiv p_t^*(i)/P_t$. Then substituting the value of $\tilde{p}_t$ from the evolution of aggregate prices

$$\tilde{p}_t = \left[\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} \right]^{\frac{1}{1-\theta}}, \quad (53)$$

we get²⁷

$$\Delta_t = (1 - \alpha) \left(\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} \right)^{-\frac{\theta}{1-\theta}} + \alpha \Pi_t^\theta \Delta_{t-1}. \quad (54)$$

From this expression we then obtain a relationship between price dispersion and trend inflation in steady state:

$$\bar{\Delta}_t = \frac{1 - \alpha}{1 - \alpha \bar{\Pi}_t^\theta} \left(\frac{1 - \alpha \bar{\Pi}_t^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta}{\theta-1}}. \quad (55)$$

We can now use the relation between steady state marginal cost and steady state inflation, namely:

$$\bar{mc}_t = \frac{\theta - 1}{\theta} \frac{\left(1 - \alpha \bar{\Pi}_t^{\theta-1} \right)^{\frac{1}{1-\theta}}}{(1 - \alpha)^{\frac{1}{1-\theta}}} \left[\frac{1 - \alpha \beta (\bar{\Pi}_t)^\theta}{1 - \alpha \beta (\bar{\Pi}_t)^{\theta-1}} \right], \quad (56)$$

together with (52) evaluated in steady state,

$$\bar{mc}_t = \chi (\bar{Y}_t^a)^{1+\nu} \bar{\Delta}_t^\nu, \quad (57)$$

²⁶This derivation follows Schmitt-Grohe and Uribe (2006, 2007).

²⁷This expression corresponds to the derivation of Yun (2005).

to obtain a relationship between inflation and output that should be satisfied in steady state. Equating (56) and (57), substituting $\bar{\Delta}_t$ from (55) and rearranging, we get

$$\bar{Y}_t^a = \left[\frac{\frac{\theta-1}{\theta} \frac{(1-\alpha\bar{\Pi}_t^{\theta-1})^{\frac{1}{1-\theta}}}{(1-\alpha)^{\frac{1}{1-\theta}}} \left[\frac{1-\alpha}{1-\alpha} \frac{\beta(\bar{\Pi}_t)^\theta}{\beta(\bar{\Pi}_t)^{\theta-1}} \right]}{\chi \left(\frac{1-\alpha}{1-\alpha\bar{\Pi}_t^\theta} \left(\frac{1-\alpha\bar{\Pi}_t^{\theta-1}}{1-\alpha} \right)^{\frac{\theta}{\theta-1}} \right)^\nu} \right]^{\frac{1}{1+\nu}} \quad (58)$$

which can be interpreted as a long-run Phillips curve relationship between inflation and output.

A.1.3 Log-linearizations

We start from a log-linear NKPC with marginal cost as forcing variable, and want to transform it into a log-linear Phillips curve in output deviations from steady state: $\tilde{Y}_t^a \equiv Y_t^a / \bar{Y}_t^a$.

To do that we only need to obtain the log-linearization of (52) around $\bar{m}\bar{c}_t$ (as defined in (57)), and substitute it into the marginal cost *NKPC*. From (52) we get

$$\widehat{mc}_t = (1 + \nu) \widehat{Y}_t^a + \nu \widehat{\Delta}_t. \quad (59)$$

For the log-linearization of (54), we first let $\widehat{\Delta}_t = \log \Delta_t / \bar{\Delta}_t$ (the non-stationarity of Π_t implies that Δ_t is also non-stationary, but its ratio to trend is by definition stationary).²⁸ Then we log-linearize the resulting expression around a steady state where $\tilde{\Delta} = \tilde{\Pi} = 1$, obtaining

$$\widehat{\Delta}_t \simeq \lambda_{1t} \widehat{\Pi}_t + \lambda_{2t} \left(\widehat{\Delta}_{t-1} - \widehat{\gamma}_{\Delta t} \right), \quad (60)$$

where the parameters λ_{1t} and λ_{2t} are defined in the last two rows of (11) in the main text. They are time-varying because they depend on trend inflation. In the main text, for analogy with the other equations, for $\widehat{\Delta}_t$ we use the notation $\delta_t - \bar{\delta}_{t-1}$.

²⁸In first order approximations around a steady state with zero inflation, the variable Δ_t can be ignored (the log deviation $\widehat{\Delta}_t$ would be a first order process with no real consequences for the stationary distribution of the other endogenous variables). But price dispersion must be taken into account if one analyzes economies with trend inflation and imperfect price indexation, as the one in this paper. (see Schmitt-Grohe and Uribe (2007)).

B Arrays for structural representations

The state vector is $S_t = [\pi_t \ \phi_t \ \delta_t \ u_t \ y_t \ g_t \ y_{t-1} \ i_t \ 1 \ \xi_t]'$. The matrices entering the PLM are defined as:

$$A_t = \begin{bmatrix} 1 & 0 & -\varsigma_t & -1 & -\kappa_t & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda_{1t} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -\xi_1 & \xi_2 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (61)$$

$$B_t = \begin{bmatrix} \beta_t + \gamma_{1t}(\theta - 1) & \gamma_{1t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \gamma_{2t}(\theta - 1) & \gamma_{2t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \beta\xi_2 & \beta\xi_2 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (62)$$

$$C_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} & 0 \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} & 0 \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_g & 0 & 0 & \mu_g & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \psi_{\pi t} & 0 & 0 & 0 & \psi_{yt} & 0 & -\psi_{yt} & 1 & -\psi_{\pi t}\bar{\pi}_t & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \xi_2 & 0 & 0 & 0 & \mu_\xi & 0 \end{bmatrix}, \quad (63)$$

$$D_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (64)$$

The expressions for the intercepts in C_t are

$$\begin{aligned} \mu_{\pi t} &= [1 - \beta_t - \gamma_{1t}(\theta - 1)]\bar{\pi}_t - \kappa_t \bar{y}_t - \varsigma_t \bar{\delta}_t, \\ \mu_{\phi_t} &= -\gamma_{2t}(\theta - 1)\bar{\pi}_t, \\ \mu_{\delta t} &= (1 - \lambda_{2t})\bar{\delta}_{t-1} - \lambda_{1t}\bar{\pi}_t, \\ \mu_y &= r - g, \\ \mu_g &= (1 - \rho_g)g, \\ \mu_\xi &= \xi - (\xi_1 + (1 + \beta)\xi_2)\bar{y}_t + g\xi_2(1 - \beta), \end{aligned} \quad (65)$$

where $\bar{y}_t$ and $\bar{\pi}_t$ are private-sector estimates respectively of steady-state output and trend inflation, and r and g are the steady-state real-interest rate and real-growth rate, respectively.

The matrices A_t , B_t , and D_t also appear in the ALM. However, C_t is replaced by

$$C_{at} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} & 0 \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} & 0 \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_g & 0 & 0 & \mu_g & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \psi_\pi & 0 & 0 & 0 & \psi_y & 0 & -\psi_y & 1 & -\psi_\pi \bar{\pi} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \xi_2 & 0 & 0 & 0 & \mu_\xi & 0 \end{bmatrix}. \quad (66)$$

The selection matrix e_X used to evaluate the likelihood function is defined as

$$e_X = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (67)$$

C The inflation target is transparent, but other aspects of policy are opaque

In a number of countries, the central bank works hard to provide a clear and credible signal of its inflation target, but frequently other elements of the policy rule are more opaque. Here we examine the extent to which a signal about $\bar{\pi}$ mitigates the problems described in the text.

We study the following scenario. We assume the central bank makes an announcement about $\bar{\pi}$ at date 0. We sidestep difficult and important questions about how the bank makes its signal credible and simply assume that it does. The private sector adopts a prior for $\bar{\pi}$ that is tightly centered on the Bank's announcement. In particular, we assume that $p(\bar{\pi})$ is truncated normal, that the untruncated density is centered on the true value, and that the standard deviation of the untruncated prior is 25 basis points per annum. Figure A1 depicts an example. The blue line reproduces the prior described above, and the red line depicts a normal prior centered on 3 percent per year with a 95 percent confidence band ranging from 2.5 to 3.5 percent per year.




Figure A1: An inflation-targeting prior

To represent opacity about other aspects of the policy rule, we assume that the bank says nothing at date 0 about ψ_π , ψ_y , or σ_i . Priors for these parameters are therefore assumed to be the same as in figure 2 in the paper.

The good news is that the private sector learns quickly about $\bar{\pi}$ and this centers long-run forecasts of inflation on the bank's target. The not-so-good news is that opacity with respect to ψ_π and ψ_y means that the potential for explosive volatility remains.

The optimal policy in this case sets $\bar{\pi} = 0.03$, $\psi_{\pi} = 0.05$, and $\psi_y = 0.7$. Average parameter estimates and responses of inflation, output, and, nominal interest are shown in figures A2 and A3, respectively. Because $p(\bar{\pi})$ is tightly concentrated on the inflation target, the private sector learns $\bar{\pi}$ right away. In figure A2, the prior mode for $\bar{\pi}$ centers on the true value, and the posterior mode remains close thereafter. As shown in figure A3, inflation again declines sharply at impact and then oscillates as it converges. After a decade, inflation has declined permanently by 1.6 percentage points at the cost of a cumulative loss in output of 7.6 percent of a year's GDP. Thus, the sacrifice ratio is 4.75 percentage points of lost output per point of inflation.

Nonexplosive regions of H_1 , shown in figure A4, are essentially the same as those in the text. Isoloss contours for this example, shown in figure A5, are also similar to those for example 1. Regions of low expected loss still concentrate in the vicinity of the prior modes for ψ_{π} and ψ_y , and temporarily explosive paths still emerge when ψ_{π} and/or ψ_y deviate too much from prior beliefs. Since the volatility of the transition depends more on the feedback parameters, what is needed is a credible signal about ψ_{π} and ψ_y .




Figure A2: Average estimates under the optimal policy




Figure A3: Average responses under the optimal policy




Figure A4: Nonexplosive region of H_1




Figure A5: Iso-loss contours