

Federal Reserve Bank of New York
Staff Reports

Optimized Taylor Rules for Disinflation When Agents Are Learning

Timothy Cogley
Christian Matthes
Argia M. Sbordone

Staff Report No. 524
November 2011
Revised May 2014



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Optimized Taylor Rules for Disinflation When Agents Are Learning

Timothy Cogley, Christian Matthes, and Argia M. Sbordone

Federal Reserve Bank of New York Staff Reports, no. 524

November 2011; revised May 2014

JEL classification: E31, E52

Abstract

Highly volatile transition dynamics can emerge when a central bank disinflates while operating without full transparency. In our model, a central bank commits to a Taylor rule whose form is known but whose coefficients are not. Private agents learn about policy parameters via Bayesian updating. Under McCallum's (1999) timing protocol, temporarily explosive dynamics can arise, making the transition highly volatile. Locally unstable dynamics emerge when there is substantial disagreement between actual and perceived feedback parameters. The central bank can achieve low average inflation, but its ability to adjust reaction coefficients is more limited.

Key words: inflation, monetary policy, learning, policy reforms, transitions

Cogley: New York University (e-mail: tim.cogley@nyu.edu). Matthes: Federal Reserve Bank of Richmond (e-mail: christian.matthes@rich.frb.org). Sbordone: Federal Reserve Bank of New York (e-mail: argia.sbordone@ny.frb.org). For comments and suggestions, the authors thank Martin Ellison, George Evans, Boyan Jovanovic, Thomas Sargent, Michael Woodford, and seminar participants at the Banque de France, the Centre for Dynamic Macroeconomic Analysis Conference at the University of St. Andrews, the European Central Bank, Duke University, the Federal Reserve Banks of Atlanta, Philadelphia, and Richmond, the Board of Governors of the Federal Reserve System, the Hungarian Central Bank, the London Business School, the National Bank of Poland Conference "DSGE and Beyond," Norges Bank, the Norwegian School of Management, New York University, the University of Oxford, Universitat Pompeu Fabra, Rutgers University, the 2011 Meetings of the Society for Computational Economics, the 2011 Meetings of the Society for Economic Dynamics, the Toulouse School of Economics, and Université du Québec à Montréal. Appendixes to this paper are available at <http://files.nyu.edu/tc60/public/>. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

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 choosing how to implement a vague mandate. We assume the new governor's preferences differ from those of his predecessor and that he wants to disinflate. We seek an optimal Taylor-type rule and study how learning affects the choice of policy parameters.

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. Erceg and Levin (2003) and Goodfriend and King (2005) explain the cost of the Volcker disinflation 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. In their model, learning increases inflation persistence relative to what would occur under full information, thereby raising the sacrifice ratio and producing 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 *normative* questions, such as how learning alters the central bank's choices and what policy is optimal in that case. Our problem is motivated by the Volcker disinflation, and a stylized version of that episode serves as the vehicle for our analysis, but our objective is not to explain the Volcker disinflation. On the contrary, our goal is to illustrate a powerful force that arises when a new policy must be learned and to describe how the bank's choices are affected.

We study this problem in the context of a dynamic new Keynesian model modified in two ways. Following Ascari (2004) and Sbordone (2007), we assume that target inflation need not be zero. We also replace rational expectations with Bayesian learning. We assume the central bank commits to a simple Taylor-type rule whose functional form is known but whose coefficients are not. Private agents learn those coefficients via Bayesian updating. The bank chooses policy-rule parameters by minimizing a discounted quadratic loss function, taking learning into account. Thus our model can be interpreted as representing the consequences of incomplete transparency about policy coefficients when the central bank is committed to a simple rule.

The main quantitative results can be summarized as follows. The optimal simple rule under the full-information benchmark brings inflation down from 4.6% to about zero in four quarters, with reaction coefficients of 1.05 and 0.11, respectively, on inflation and output growth. The sacrifice ratio, defined as the cumulative loss in output divided by the change in target inflation, is approximately 0.5% in this benchmark case. The optimal simple rule under learning reduces target inflation to 1 percent with reaction coefficients of 0.25 on inflation and 0.15 on output growth. The transition takes about 10 quarters with much more oscillation, and the sacrifice ratio is about three times as large than under full information.

The reason why the bank's choice under learning differs substantially from the full-information optimum is that the equilibrium law of motion under learning can be a temporarily explosive process, i.e. one that is asymptotically stationary but which has unstable autoregressive roots during the transition. When locally-unstable dynamics emerge, the transition is highly volatile and dominates expected loss. The central bank's main challenge is to find a way to manage this potential for explosive volatility.

Uncertainty about the inflation target is a lesser evil. In our examples, the bank always achieves low average inflation, though sometimes it stops short of zero – the

optimum under full information¹ – because the transition cost would be too great. Uncertainty about policy feedback parameters is more problematic because this is what creates the potential for temporarily-explosive dynamics. Locally-unstable dynamics emerge when there is substantial disagreement between actual and perceived feedback parameters. It follows that one way for the bank to cope is to adopt a policy that is close to the private sector’s prior. By choosing feedback parameters sufficiently close to the private sector’s prior mode, the bank can ensure that the equilibrium law of motion is nonexplosive throughout the transition, sacrificing better long-term performance for lower transitional volatility. For the model described below, this approximates the optimal strategy. Thus the bank’s choice of feedback parameters is more constrained by the private sector’s initial beliefs.

A lack of transparency can therefore make disinflation very costly even under commitment to a simple rule. Furthermore, although conventional wisdom emphasizes the value of an explicit long-run inflation target, our analysis says that transparency about reaction coefficients is equally important, perhaps even more so.

Our approach to learning differs from much of the macro-learning literature, in particular from the branch emanating from Marcet and Sargent (1989a, 1989b), Cho, Williams, and Sargent (2002), and Evans and Honkapohja (2001, 2003). 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’ theorem. 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, the actual and perceived laws of motion (ALM and PLM, respectively) are tightly linked. In our model, agents know the ALM up to the unknown policy coefficients, and their PLM is the perceived ALM (i.e., the ALM evaluated at their current estimate of the policy coefficients). Because agents know the ALM’s functional form, they can use Bayes’ theorem to update beliefs. Nevertheless, the assumption that agents are Bayesian is not critical. We also examine whether our insights are robust to alternative forms of learning, and we find that they are.

The remainder of the paper is organized as follows. Section 2 describes the model, and section 3 explains how agents update beliefs. Section 4 presents results for our

¹We abstract from the zero lower bound on nominal interest.

baseline specification, and section 5 examines a number of perturbations to that specification. Section 6 explains how our paper relates to a number of others in the literature, and section 7 concludes. A series of online appendices contains additional results.²

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 our behavioral assumptions and the model's structure. 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 treat estimated parameters as if they were known with certainty and formulate plans accordingly. Following McCallum (1999), we assume that the central bank sets the systematic part of its instrument rule at the beginning of the period based on information inherited from $t - 1$. Then period t shocks are realized. Agents observe the central bank's policy action and infer a perceived policy shock $\tilde{\varepsilon}_{it}$. They also observe realizations of the private-sector shocks. Current-period outcomes are then determined in accordance with beginning-of-period plans. After observing those outcomes, private agents update their estimates of policy coefficients and carry them forward to $t + 1$.

2.2 The model

Our model is a dynamic new Keynesian model in which agents form expectations using a subjective forecasting model that can differ from the equilibrium law of motion. Monetary policy is determined according to a Taylor-type rule that allows target inflation to differ from zero. Private-sector behavior is characterized by two blocks of equations, an intertemporal *IS* curve and an Ascari-Sbordone version of the new Keynesian Phillips curve. The model features staggered price setting and habit persistence in consumption. A log-linearized version is presented here. For details on

²Appendices are posted at <http://files.nyu.edu/tc60/public/>.

how this representation was derived, see appendix A.

2.2.1 Monetary policy

The baseline model assumes that the central bank commits to a Taylor rule in difference form,

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

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.* normal policy shock with mean zero and variance σ_i^2 . The policy coefficients are collected in a vector $\psi = [\bar{\pi}, \psi_\pi, \psi_y, \sigma_i]'$, where $\bar{\pi}$ represents the central bank's long-run inflation target and ψ_π and ψ_y are feedback parameters on the inflation gap and output growth, respectively.

There are several reasons for specifying a policy rule of this form. Our paper is part of the literature on optimal simple rules, and Taylor-type rules are by far the most influential in this literature. A difference form was adopted because it seems promising for environments like ours. For instance, Coibion and Gorodnichenko (2011) establish that a rule of this form ameliorates indeterminacy problems in Calvo models with positive target inflation, and Orphanides and Williams (2007) demonstrate that it performs well under least-squares learning. 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)). In addition, we agree with McCallum (1999) that monetary policy rules should be specified in terms of lagged variables because the Fed lacks good current-quarter information about inflation and output.³ Last but not least, Erceg and Levin (2003) contend that output growth, rather than the output gap, is more appropriate for estimated policy reaction functions for the U.S.

Private agents know the form of the policy rule but not its coefficients. At any given date, their perceived policy rule is

$$i_t - i_{t-1} = \psi_{\pi t}(\pi_{t-1} - \bar{\pi}_t) + \psi_{yt}(y_{t-1} - y_{t-2}) + \tilde{\varepsilon}_{it}, \quad (2)$$

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

$$\tilde{\varepsilon}_{it} = \varepsilon_{it} + (\psi_\pi - \psi_{\pi t})\pi_{t-1} + (\psi_y - \psi_{yt})\Delta y_{t-1} + \psi_{\pi t}\bar{\pi}_t - \psi_\pi\bar{\pi} \quad (3)$$

³For instance, the Bureau of Economic Analysis released the advance estimate of 2013.Q4 GDP on January 30, 2014, one month after the end of the quarter.

is a perceived policy shock. Private agents believe that $\tilde{\varepsilon}_{it}$ is white noise, but it actually depends on lags of inflation and output growth and errors in estimates of policy coefficients.

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, so it involves both the actual policy (1) and the perceived policy (2).

The central bank minimizes a discounted quadratic loss function,

$$L = E_0 \sum_t \beta^t [\pi_t^2 + \lambda_y (y_t - \bar{y})^2 + \lambda_i (i_t - \bar{i})^2], \quad (4)$$

that penalizes variation in inflation and the output gap, and deviations of the nominal interest rate from its steady state. We assume that the central bank arbitrarily sets σ_i and optimizes with respect to $\bar{\pi}$, ψ_π , and ψ_y , taking private-sector learning into account.⁴

2.2.2 Behavioral assumptions

The agents who inhabit the private sector are boundedly-rational DSGE modelers who know a lot about their environment but not quite as much as agents in a full-information rational-expectations model. They understand the structure of the economy and the form of the monetary-policy rule, but they do not know its coefficients. They build a structural model of the economy and use it for forecasting, decision making, and learning.

Their behavior is boundedly rational in two respects. Their first-order conditions take the form of nonlinear expectational difference equations that they cannot solve. Instead, they log-linearize around a steady state and work with the resulting system of linear expectational difference equations. Not knowing the economy's *true* steady state, however, they expand around the *perceived* steady state in period t . The true steady state \bar{x} is the deterministic steady state associated with the true policy coefficients ψ . The perceived steady state \bar{x}_t is defined 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 $\bar{\pi}$ have level effects on nominal variables and on some real variables (Ascari 2004). Since perceptions of $\bar{\pi}$ change as agents update

⁴The central bank does not experiment because it knows everything. Private agents do not experiment because they are atomistic and cannot unilaterally influence the bank's actions. For both, the marginal cost of experimentation would be positive and the marginal benefit zero.

their beliefs, so do their long-run forecasts. Although nonstandard, expanding around the perceived steady state better reflects the agents' knowledge and state of mind at date t .

Private agents also behave as anticipated-utility modelers (Kreps 1998), treating the current estimate ψ_t as if it were known with certainty. In the context of a single-agent decision problem, Cogley and Sargent (2008) compare the resulting decision rules with exact Bayesian decision rules and demonstrate that the approximation is good as long as precautionary motives are not too strong. Like a log-linear approximation, this imposes a form of certainty equivalence, for it implies that decision rules are the same regardless of the degree of parameter uncertainty. The anticipated-utility approach is standard in the macro-learning literature.

2.2.3 A new-Keynesian IS curve

As usual, a representative household maximizes expected utility subject to a flow budget constraint. The household's period-utility function is

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

where C_t is consumption of a final good, H_t represents hours of work, b_t and χ_t are preference shocks, and η measures the degree of habit persistence in consumption. The first-order condition is a conventional consumption Euler equation. After log-linearizing, agents obtain a version of the new Keynesian *IS* curve,

$$y_t - \bar{y}_t = \xi_t - \xi - E_t^* [\xi_{t+1} - \xi - (y_{t+1} - \bar{y}_t) - (\gamma_{t+1} - \gamma) + i_t - \pi_{t+1} - r], \quad (5)$$

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

$$\xi_t - \xi \equiv \xi_1 (y_t - \bar{y}_t) + \xi_2 [y_{t-1} - \bar{y}_t - (\gamma_t - \gamma) + \beta E_t^* (y_{t+1} - \bar{y}_t + \gamma_{t+1} - \gamma)] + \varepsilon_{yt}. \quad (6)$$

The parameter β is a subjective discount factor, r and γ are steady-state values for the real-interest rate and the growth rate of technological progress, respectively, and \bar{y}_t is the private sector's beginning-of-period long-run forecast for output. The coefficients ξ_1 and ξ_2 are combinations of preference and technology parameters, and γ_t and ε_{yt} are technology and preference shocks, respectively. Further details can be found in appendix A.

This representation differs in three ways from standard *IS* equations. One difference concerns the choice of the expansion point. As mentioned above, agents expand

around the perceived steady state \bar{y}_t instead of the actual steady state \bar{y} . In addition, the 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 (5) and (6). 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 .⁵ Finally, two shocks appear, a persistent shock γ_t to the growth rate of technology,

$$\gamma_t = (1 - \rho_\gamma) \gamma + \rho_\gamma \gamma_{t-1} + \varepsilon_{\gamma t}, \quad (7)$$

and a white-noise shock ε_{yt} .

2.2.4 A new-Keynesian Phillips curve

Following Calvo (1983), we assume that a continuum of monopolistically competitive firms produce a variety of differentiated intermediate goods that are sold to a final-goods producer. Intermediate-goods producers reset their prices at random intervals, with α representing the probability that their price remains the same. Thus we abstract from indexation or other backward-looking pricing influences, in accordance with the estimates of Cogley and Sbordone (2008). Since pricing and supply decisions depend on the beliefs of private agents, they again log-linearize around perceived steady states, obtaining the following block of equations,

$$\pi_t - \bar{\pi}_t = \beta_t E_t^* (\pi_{t+1} - \bar{\pi}_t) + \kappa_t (y_t - \bar{y}_t) + \varsigma_t (\delta_t - \bar{\delta}_t) - \tilde{\kappa}_t (\xi_t - \xi) \quad (8)$$

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

$$\phi_t = \gamma_{2t} E_t^* [(\theta - 1)(\pi_{t+1} - \bar{\pi}_t) + \phi_{t+1}], \quad (9)$$

$$\delta_t - \bar{\delta}_t = \lambda_{1t} (\pi_t - \bar{\pi}_t) + \lambda_{2t} (\delta_{t-1} - \bar{\delta}_t). \quad (10)$$

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

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

⁵We assume that the central bank knows the private sector's prior over ψ . Because 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 $E_t x_{t+j} \neq E_t^* x_{t+j}$.

depend on deep parameters and estimates of target inflation $\bar{\pi}_t$. 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, even though the deep parameters are invariant to changes in policy, the *NKPC* coefficients are not. The latter change as beliefs about $\bar{\pi}_t$ are updated. Equation (11) collapses to the usual expressions when $\bar{\pi}_t = 0$.

Second, a variable

$$\delta_t \equiv \ln \left(\int_0^1 (p_t(i) / P_t)^{-\theta} di \right), \quad (12)$$

measuring the resource cost of cross-sectional price dispersion, has first-order effects on inflation and other variables. If $\bar{\pi}_t$ were zero, this variable would drop out of a first-order expansion.

Third, higher-order leads of inflation appear on the right-hand side of (8). To retain a first-order form, we introduce an intermediate variable ϕ_t that has no interesting economic interpretation and add equation (9). This is simply a device for obtaining a convenient representation.

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

$$u_t = \rho_u u_{t-1} + \varepsilon_{ut}, \quad (13)$$

and a white-noise shock $\varepsilon_{\pi t}$.

2.2.5 Calibration

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

$$\alpha = 0.6, \quad \beta = 0.99, \quad \theta = 10. \quad (14)$$

Preference parameters are calibrated as follows. The parameter ν is the inverse of the Frisch elasticity of labor supply. The literature provides a large range of values for this elasticity, typically high in the macro literature and low in the labor literature. We set $\nu = 0.5$, which implies a Frisch elasticity of 2 and represents a compromise between the two. We think our calibration 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).

We also 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 0.5, which implies that the weight on fluctuations of the annualized nominal interest rate is half the weights attached to fluctuations in annualized inflation and the output gap.⁶

Turning to parameters governing the shocks, we set $\gamma = 0$, thereby abstracting from average growth. For the persistent shocks u_t and γ_t , estimates are taken from Cogley, Primiceri, and Sargent (2010),

$$\begin{aligned}\rho_u &= 0.4, & 100\sigma_u &= 0.12, \\ \rho_\gamma &= 0.27, & 100\sigma_\gamma &= 0.5.\end{aligned}\tag{15}$$

Last but not least, the standard deviations of the white noise shocks ε_{yt} and $\varepsilon_{\pi t}$ are set equal to

$$\sigma_\pi = \sigma_y = 0.01/4.\tag{16}$$

3 Learning about monetary policy

Everyone knows the model of the economy and the form of the policy rule, but private agents do not know the policy coefficients. Instead, they learn about them by solving a signal-extraction problem. If ψ entered linearly, they could do this with the Kalman filter. Because ψ enters non-linearly, however, agents must solve a nonlinear filtering problem. This section describes how this 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 unknown policy coefficients, we use the ALM to derive the likelihood function. Agents combine the likelihood with a prior over policy parameters and use the posterior mode as their point estimate.

3.1 The perceived law of motion

By stacking the IS equations, the aggregate supply block, exogenous shocks, and perceived monetary-policy rule, the private sector's model of the economy can be

⁶Results for economies with learning are not sensitive to the choice of λ_i .

represented as a system of linear expectational difference equations,

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

where S_t is the model's state vector, $\tilde{\varepsilon}_t$ is a vector of perceived innovations, and A_t, B_t, C_t , and D_t depend on the model's deep parameters (see appendix A.5 for details). These matrices have time subscripts because they depend on estimates of the policy coefficients ψ_t . We conjecture that the PLM is the reduced-form VAR associated with (17),

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

where F_t solves $B_t F_t^2 - A_t F_t + C_t = 0$ and $G_t = (A_t - B_t F_t)^{-1} D_t$.⁷ As in a conventional rational-expectations model, (18) serves two functions, describing agents' current-quarter plans and how they forecast future outcomes.

3.2 The actual law of motion

To find the actual law of motion, we stack the actual policy rule (equation 1) with equations governing private sector behavior (5-7, 8-13). 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 vector and the matrices A_t, B_t , and D_t are the same as in (17). 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 appendix A.5).

A solution for the ALM can be found as follows. Since outcomes are determined in accordance with agents' plans (equation 18), they depend on the perceived shocks $\tilde{\varepsilon}_t$. A relation between perceived and actual innovations can be found by subtracting (19) from (17),

$$D_t \tilde{\varepsilon}_t = D_t \varepsilon_t + (C_{at} - C_t) S_{t-1}. \quad (20)$$

Substituting this relation back into agents' plans expresses outcomes in terms of actual shocks,

$$S_t = H_t S_{t-1} + G_t \varepsilon_t, \quad (21)$$

⁷Following Sims (2001), this matrix quadratic equation is solved using a generalized Schur decomposition.

where

$$H_t = F_t + (A_t - B_t F_t)^{-1}(C_{at} - C_t). \quad (22)$$

The ALM depends on both actual policy coefficients, because that is what governs central bank behavior, and on perceived policy coefficients, because that is what guides private-sector behavior.

When there is a unique nonexplosive solution for (F_t, G_t) , the solution for H_t is also unique but not necessarily nonexplosive. When multiple nonexplosive solutions for (F_t, G_t) exist, there are also multiple solutions for H_t , and our programs choose one of them. However, this kind of multiplicity never occurs in our simulations.

3.3 The PLM is the perceived ALM

The reduced-form ALM and PLM are both $VAR(1)$ processes with conditionally gaussian innovations. Under the ALM, the conditional mean and variance are⁸

$$\begin{aligned} m_{t|t-1}(\psi_{true}) &= H_t(\psi_{true})S_{t-1}, \\ V_{t|t-1}(\psi_{true}) &= G_t V_\varepsilon(\psi_{true})G'_t, \end{aligned} \quad (23)$$

where $H_t(\psi_{true})$ and $V_\varepsilon(\psi_{true})$ are the ALM conditional mean and variance arrays evaluated at the true value ψ_{true} . If the agents in the model were interviewed and asked their view of the ALM, they would answer by replacing ψ_{true} in C_{at} with ψ_t , thus obtaining C_t , implying

$$\begin{aligned} \tilde{m}_{t|t-1}(\psi_t) &= F_t S_{t-1}, \\ \tilde{V}_{t|t-1}(\psi_t) &= G_t V_\varepsilon(\psi_t)G'_t. \end{aligned} \quad (24)$$

These expressions coincide with the conditional mean and variance under the PLM. Hence the PLM is the perceived ALM. This is true not only asymptotically but for every date during the transition.⁹

⁸According to the timing protocol, H_t and G_t can be regarded either as beginning-of-period t estimates or end-of-period $t-1$ estimates, which explains why it is legitimate to use them to calculate the conditional mean and variance.

⁹Among other things, this implies that private-sector forecasts are consistent with 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 .

3.4 The likelihood function

The observables are stacked in a vector $X_t = [\pi_t, u_t, y_t, \gamma_t, i_t]'$ = $e_X S_t$, where e_X is an appropriately defined selection matrix (see appendix A.5). The other elements of S_t 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

$$p(X^t|\psi) = \prod_{j=1}^t p(X_j|X^{j-1}, \psi). \quad (25)$$

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 (25). According to the ALM, X_t is conditionally normal with mean and variance

$$\begin{aligned} m_{t|t-1}^X(\psi) &= e_X H_t(\psi) S_{t-1}, \\ V_{t|t-1}^X(\psi) &= e_X G_t V_\varepsilon(\psi) G_t' e_X', \end{aligned} \quad (26)$$

where $H_t(\psi)$ and $V_\varepsilon(\psi)$ are the ALM conditional mean and variance, respectively, evaluated at some value of ψ . It follows that the log-likelihood function is

$$\begin{aligned} \ln p(X^t|\psi) &= -\frac{1}{2} \sum_{j=1}^t \{ \ln |V_{j|j-1}^X(\psi)| \\ &\quad + [X_j - m_{j|j-1}^X(\psi)]' (V_{j|j-1}^X(\psi))^{-1} [X_j - m_{j|j-1}^X(\psi)] \}. \end{aligned} \quad (27)$$

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 the 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, they adopt the posterior mode,

$$\psi_t = \arg \max (\ln p(X^t|\psi) + \ln p(\psi)). \quad (28)$$

Notice that agents take into account that past outcomes were influenced by past beliefs. By inspecting the ALM and PLM, one can verify that past values of the conditional mean $m_{j|j-1}^X$ and the conditional variance $V_{j|j-1}^X$ 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.

Notice also that 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. How much the cross-equation restrictions matter in this context is examined below.

4 Quantitative analysis of a backward-looking rule

A new governor appears at date 0 and formulates a new policy rule that becomes operative at date 1. After observing the private sector's prior, the governor chooses the long-run inflation target $\bar{\pi}$ and reaction coefficients ψ_π, ψ_y to minimize expected loss under the new policy, with the standard deviation of policy shocks σ_i being set exogenously. We initially assume that $\sigma_i = 0.001$ (10 basis points per quarter) and later examine what happens when σ_i is zero.

4.1 Initial conditions

The economy is initialized at the steady state under the old regime. To create a scenario like the end of the Great Inflation, we calibrate the old regime to match estimates of the policy rule for the period 1966-1981. We assume that the policy rule for that period had the same functional form as in equation (1), and we estimate $\bar{\pi}, \psi_\pi, \psi_y$, and σ_i^2 by OLS. 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, 1966-1981, with standard errors in parentheses.

The estimate for $\bar{\pi}$ implies an annualized inflation target of 4.6 percent. The reaction coefficients are both close to zero, with ψ_y being slightly larger than ψ_π . Policy shocks are large in magnitude and account for a substantial fraction of the total variation in the nominal interest rate. Standard errors are large, especially for ψ_π . The economy is initialized at the steady state associated with this policy rule, $\pi_0 = 0.0116$, $y_0 = -0.0732$, and $i_0 = 0.0217$, where inflation and nominal interest are expressed as quarterly rates.

4.2 Evaluating expected loss and finding the optimal simple rule

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 this is essential, we 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 25 years long, and the terminal continuation value is set to zero, representing a decision maker with a long but finite horizon. Realized loss is calculated for each sample path, and expected loss is the cross-path average of realized loss. The optimal rule among this family is the node with smallest expected loss.

4.3 A full-information benchmark

To highlight the role of learning, we begin by describing the optimum under full information. When private agents know the new policy coefficients, the optimized Taylor rule sets $\bar{\pi} = 0$, $\psi_\pi = 1.05$, and $\psi_y = 0.11$. Figure 1 depicts average responses of inflation, output, and nominal interest gaps, which are defined as deviations from the steady state of the new regime. Recall that the economy is initialized in the steady state of the old regime and that the disinflation commences at date 1.¹⁰

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 because 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. Under full information, the optimal simple rule shares this property.

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

¹⁰Inflation and nominal interest gaps at date 0 coincide because the steady-state real interest rate is the same under the two regimes.




Figure 1: Response of inflation, output and interest rate under full information

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 in about a year. After 4 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 2.6 percent. The sacrifice ratio, defined as the cumulative loss in output divided by the change in target inflation, is 0.56 percent. The sacrifice ratio is small under full information because the model has no indexation. Although prices are sticky, the absence of indexation means that inflation is weakly persistent. The absence of indexation also explains why the bank seeks a substantial rollback in the price level.

Under full information, the economy is highly fault tolerant¹¹ with respect to policies away from the optimum. Figure 2 portrays iso-expected loss contours as a function of $\bar{\pi}$, ψ_{π} , and ψ_y . Each panel involves a different setting for $\bar{\pi}$, ranging from 0 to 3 percent per annum, and ψ_{π} and ψ_y are shown on the horizontal and vertical axes, respectively. Expected loss is normalized by dividing by the loss under the optimal rule so that contour lines represent gross deviations from the optimum. The diamond in the upper left panel depicts the optimal simple rule. Expected loss increases slowly as policy moves away from the optimum. For instance, when $\bar{\pi} = 0$,

¹¹Levin and Williams (2003) introduced the term “fault tolerance” to describe the extent to which expected loss increases as policies move away from the optimum.

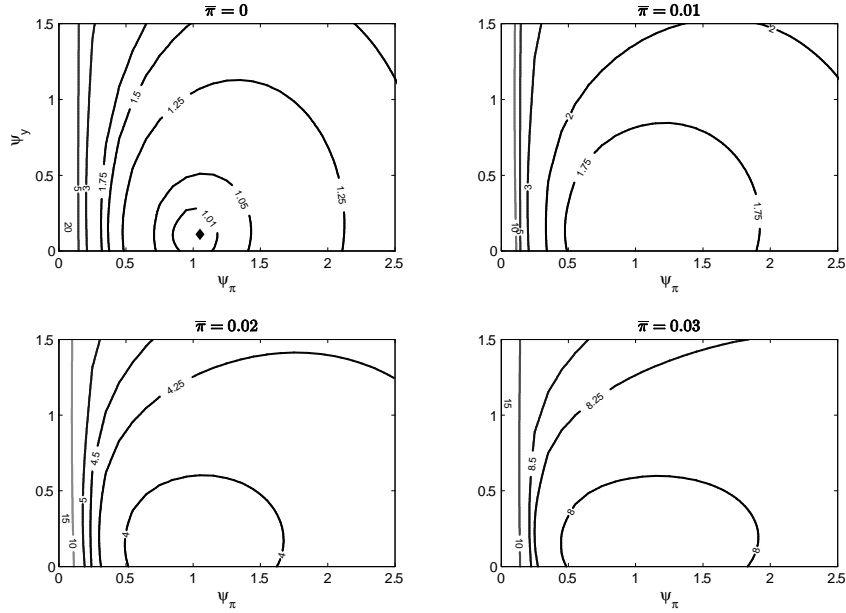


Figure 2: Iso-expected loss contours under full information

relative loss remains below 2 for most combinations of ψ_π and ψ_y and rises above 10 only when ψ_π approaches zero. Although expected loss is higher for higher values of $\bar{\pi}$, the surface remains relatively flat. Later we contrast this with an absence of fault tolerance under learning.

4.4 A Taylor rule optimized for learning

We assume that private agents initially anticipate a continuation of the old regime, and we calibrate their priors using the estimates of policy coefficients for 1966-1981 shown in table 1. In particular, they believe that policy coefficients are independent a priori,

$$p(\psi) = p(\bar{\pi})p(\psi_\pi)p(\psi_y)p(\sigma_i^2), \quad (29)$$

and they adopt truncated normal priors for $\bar{\pi}$, ψ_π , and 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 in table 1. To enforce nonnegativity, the unrestricted priors are truncated at zero and renormalized so that transformed priors integrate 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 3.




Figure 3: A prior based on the old regime

Priors for ψ_{π} and ψ_y concentrate slightly to the right of zero, and little mass is assigned to values greater than 0.25. On the other hand, priors for $\bar{\pi}$ and σ_i assign non-negligible probability to a broad range of values. According to this specification, private agents are open to persuasion about $\bar{\pi}$ and σ_i but are skeptical that the central bank will react aggressively to inflation or output. Overcoming that skepticism will be a major challenge for the central bank.

Figure 4 portrays iso-expected loss contours as a function of $\bar{\pi}$, ψ_{π} , ψ_y . As before, σ_i is held constant at 10 basis points per quarter. The left-hand column depicts the results of a broad search over a coarse grid, while the column on the right portrays calculations based on a finer grid that focuses on the low expected-loss region of the policy-coefficient space. Expected loss is again normalized by dividing by the loss for the rule optimized for learning.

In the left-hand column, regions of low expected loss concentrate in the southwest quadrant of the panels, 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, expected loss is more than 100 times greater than under the optimal simple rule. The optimal simple rule under full information is marked by an asterisk and lies in the

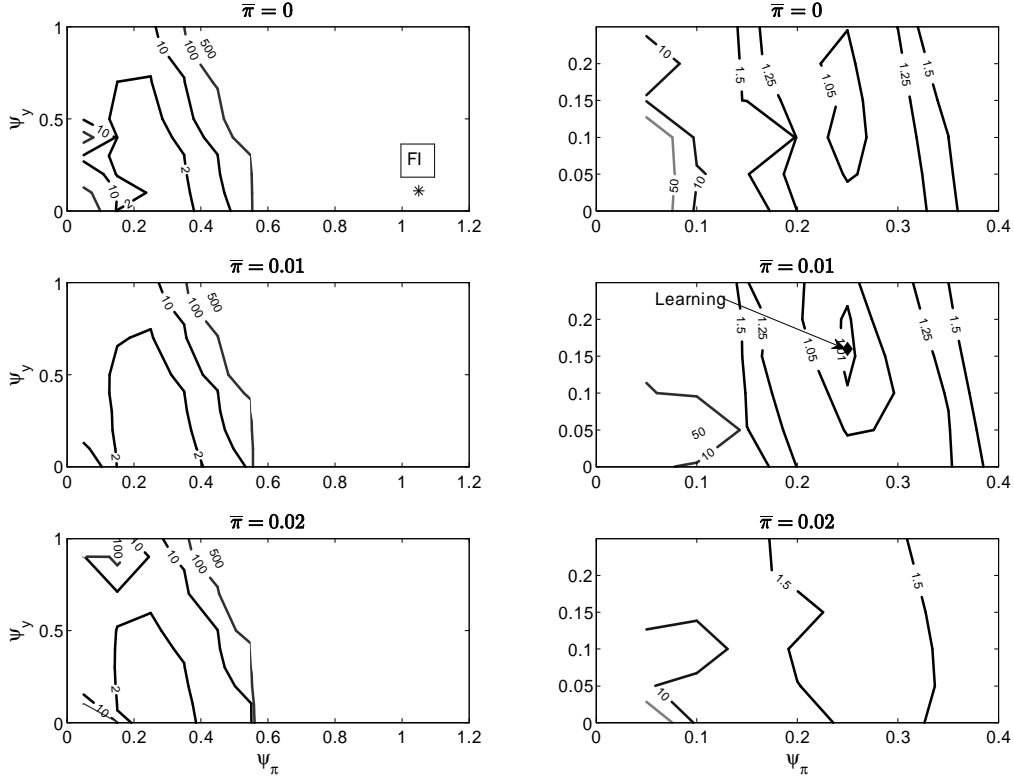


Figure 4: Iso-expected loss contours under learning

high-loss region.

The reason why the economy loses fault tolerance under learning 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 the PLM, which depends on F_t , but outcomes are governed by the ALM, which involves H_t . The eigenvalues of F_t are never outside the unit circle but the eigenvalues of H_t 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 actual policy ψ and the perceived policy ψ_t (see equation 22). The eigenvalues of H_t are close to those of F_t (hence 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 sim-




Figure 5: Nonexplosive region for H_1

ulated 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 shaded area in figure 5 depicts the region of the policy-coefficient space for which the eigenvalues of H_1 are nonexplosive. Since the nonexplosive region is similar for all settings of $\bar{\pi}$, we just show it for $\bar{\pi} = 0$. This region is sensitive to ψ_π and ψ_y , however, and concentrates near the prior mode. The central bank can move $\bar{\pi}$ far from the private sector's prior mode without generating locally-unstable dynamics, but moving ψ_π and/or ψ_y far from their prior modes makes the transition turbulent.

To locate the optimum under learning, we search on a finer grid in the southwest quadrant of the (ψ_π, ψ_y) space. Iso-loss contours are shown in the right column of figure 4, and the optimum is marked by a diamond, $\bar{\pi} = 0.01$, $\psi_\pi = 0.25$, and $\psi_y = 0.15$. Relative to the full-information solution, target inflation is slightly higher, and the reaction to output growth is a bit more aggressive. The main difference, however, is that the central bank responds less aggressively to inflation. Since the full-information optimum $\psi_\pi = 1.05$ lies outside the nonexplosive region, the transition would be initially very turbulent. Furthermore, since the private sector is prejudiced against such large values of ψ_π , explosive eigenvalues would remain active for too long. For these reasons, the optimal policy puts ψ_π and ψ_y only slightly outside the nonexplosive region. The bank can adjust $\bar{\pi}$ more freely, however, thereby achieving low average inflation.

Because the location of the nonexplosive regions depends more on ψ_π and ψ_y than on $\bar{\pi}$, uncertainty about reactions coefficients is more problematic than uncertainty about target inflation. As shown in appendix B, when uncertainty about ψ_π , ψ_y , and σ_i is deactivated and $\bar{\pi}$ is the only uncertain policy parameter, the initial nonexplosive region expands to fill most of the (ψ_π, ψ_y) space. Since the ALM becomes nonexplosive for most policies, the economy becomes highly fault tolerant, and private agents learn $\bar{\pi}$ very quickly. For these reasons, the model behaves much as it does under full information. The optimal policy is similar, and impulse response functions resemble those in figure 1. In contrast, when uncertainty about $\bar{\pi}$ is deactivated and ψ_π , ψ_y , and σ_i are uncertain, the results are qualitatively similar to those shown here. Uncertainty about feedback parameters is more costly because it activates locally-explosive dynamics.

A second loss of fault tolerance emerges in the right column of figure 4. For small values of ψ_π , estimates occasionally stray too close to zero, pushing the PLM close to the indeterminacy region. Outcomes are highly volatile when this occurs, causing expected loss to rise. For an effective stabilization, the bank must choose a value for ψ_π that not only achieves determinacy under full information but which also guards against estimates straying too closely to zero during the transition.

Figure 6 portrays impulse response functions for inflation, output, and nominal interest gaps for the optimal simple rule under learning. The transition is longer and more volatile than under full information. Inflation again declines at impact, overshooting $\bar{\pi}$ and partially rolling back past increases in the price level, but now inflation oscillates as it converges to its new long-run target. The transition takes about two and a half 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 -0.9 percent in quarter 5 and remains negative for 3 years. The cumulative output gap during this time is -6.6 percent. Since inflation falls permanently by 3.6 percentage points, the sacrifice ratio amounts to 1.8 percent of lost output per percentage point of inflation, 3 times larger than under full information. According to Ascari and Ropele (2011), most estimates of the sacrifice ratio for the Volcker disinflation lie between 1 and 3, so our model is in the right ballpark.




Figure 6: Average responses under the Taylor rule optimized for learning

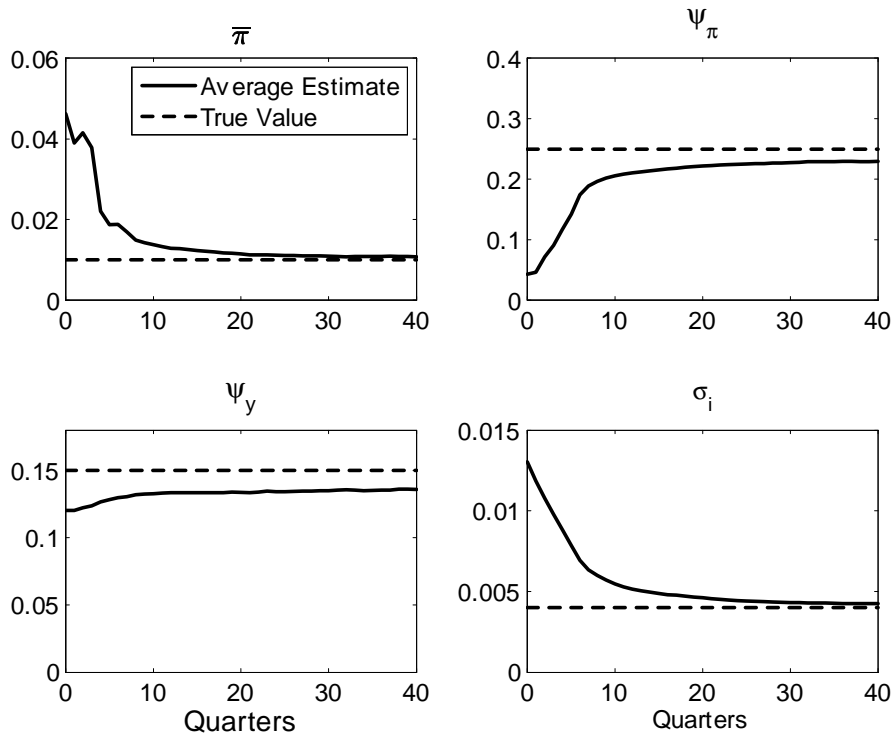


Figure 7: Average estimates of policy coefficients

Figure 7 portrays mean estimates of the policy coefficients, again averaged across 100 sample paths. The true coefficients are shown as dashed lines while average estimates are portrayed as solid lines. The estimates move quickly toward their respective true values and are not far off after 10 quarters. Rapid convergence of ψ_π and ψ_y are crucial for eliminating locally-explosive dynamics. Beliefs about target inflation and the policy shock variance also quickly approach neighborhoods of their respective true values, but this seems secondary for transitional volatility.

5 Perturbations to the baseline learning model

To highlight aspects of the baseline model, we now turn to a number of perturbations. For the sake of brevity, the main points are summarized here, and a full presentation of results is relegated to a series of appendices.

5.1 McCallum’s information constraint

McCallum’s information constraint plays a critical role in our analysis. To highlight its importance, we contrast the backward-looking Taylor rule in equation (1) with one involving contemporaneous feedback to inflation and output growth,

$$i_t - i_{t-1} = \psi_\pi(\pi_t - \bar{\pi}) + \psi_y(y_t - y_{t-1}) + \varepsilon_{it}. \quad (30)$$

There is also a slight change in the timing protocol. Private agents still enter period t with beliefs about policy coefficients inherited from $t - 1$, and they treat estimated parameters as if they were known when updating their decision rules. But now the central bank and private sector simultaneously execute their contingency plans when period t shocks are realized. After observing current-quarter outcomes, private agents update estimates and carry them forward to $t + 1$. All other aspects of the model remain the same, including the prior on $(\bar{\pi}, \psi_\pi, \psi_y, \sigma_i)$.

Because actual central banks cannot observe current quarter output or the price level, they would not be able to implement this policy. We examine it here in order to isolate the consequences of lags in the central bank’s information flow.

As shown in appendix C, locally-explosive dynamics vanish in this case, and the learning economy becomes highly fault tolerant. The model therefore behaves more like its full-information counterpart than did the economy with a backward-looking rule. For instance, while the full-information optimum sets $\bar{\pi} = 0$, $\psi_\pi = 2.4$, and

$\psi_y = 0.1$, the rule optimized for learning sets $\bar{\pi} = 0$, $\psi_\pi = 1.4$, and $\psi_y = 0.1$. The learning rule has the same inflation target and reaction coefficient on output growth as under full information, but it responds to inflation gaps a bit less aggressively. Compared with the baseline model, however, the central bank is less constrained by initial beliefs and freer to adjust its reaction coefficients. The learning transition is also shorter and less volatile than for the backward-looking rule, and the sacrifice ratio is about the same (1.9 percent for the contemporaneous rule as opposed to 1.8 percent for the backward-looking policy). Learning is slower than for the backward-looking rule, but that is because there is less transitional volatility in inflation and output growth.

Many of the difficulties reported for the baseline case follow from the fact that central banks cannot observe current quarter output and inflation. If a rule with contemporaneous feedback to inflation and output were feasible, it would be superior. The difference between contemporaneous and backward-looking Taylor rules is more pronounced under learning than under full information. More important, though, even when allowing contemporaneous feedback in the policy rule, the optimal simple rule under learning responds less aggressively to inflation, and disinflation carries a nontrivial cost. Learning precludes a sharp low-cost disinflation for the contemporaneous rule as well as the backward-looking policy.

5.2 Policy shocks

The baseline calibration for σ_i reflects a tension between two considerations. On the one hand, estimated policy reaction functions never fit exactly, implying $\sigma_i > 0$. On the other, a fully optimal policy would presumably be deterministic, implying $\sigma_i = 0$. The baseline specification compromises with a small positive value ($\sigma_i = 10$ basis points per quarter).

If the true value of σ_i were zero and known with certainty, the signal extraction problem would unravel, with agents perfectly inferring the other three policy coefficients after just three periods. This would not happen in our model even if σ_i were zero because the agents' prior on σ_i encodes a belief that monetary-policy shocks are present. Prior uncertainty about σ_i is enough to preserve a nontrivial signal-extraction problem.

Furthermore, since the initial nonexplosive region depends neither on σ_i nor on prior beliefs about σ_i , the central bank's main challenge in a $\sigma_i = 0$ economy would be

the same. It follows that the optimized rule should be similar. Appendix D confirms this intuition: when $\sigma_i = 0$ and all other aspects of the baseline economy are held constant, the optimized rule sets $\bar{\pi} = 0.01$, $\psi_\pi = 0.15$, and $\psi_y = 0.1$. Thus target inflation and the response to output growth are about the same, and the response to inflation is a bit weaker. The same is true when $\sigma_i = 0$ and the prior standard deviation for σ_i is reduced by half. In both cases, impulse response functions under the optimized rule resemble those in figure 6.

That agents entertain a belief that policy shocks are present is critical. Whether actual policy shocks are small or zero is secondary.

5.3 A two-tier approach¹²

In the baseline model, the central bank introduces two reforms at once, reducing target inflation and strengthening stabilization by responding more aggressively to inflation and output growth. Appendix E contrasts this with a two-tier approach that separates the reforms, with policymakers first switching to a rule designed to bring target inflation down and thereafter changing feedback parameters to stabilize the economy around the new target.

We formulate this approach as follows. We assume that for a certain period whose length is exogenously set, the policymaker reduces $\bar{\pi}$ but continues with response coefficients inherited from the old regime. After this initial period, when beliefs about $\bar{\pi}$ have had a chance to adjust, the policymaker adjusts the reaction coefficients. Once again, all other aspects of the baseline specification remain the same. Appendix E considers models in which the first stage lasts 10 and 20 quarters, respectively.

Alas, the two-tier approach prolongs the transition and makes matters worse. Delaying the second reform postpones but does not circumvent the problem of coping with locally-explosive dynamics. This challenge now emerges at the end stage 1 rather than the beginning of the disinflation, but it does not go away.

A separation of reforms also retards learning. During stage 1, beliefs about ψ_π and ψ_y harden around old-regime values because agents observe more weak responses to inflation and output growth. This hinders learning about ψ_π and ψ_y in stage 2. Less obviously, the separation of reforms also retards learning about target inflation in stage 1. Wherever $\bar{\pi}$ appears in the likelihood function it is multiplied by ψ_π . Since ψ_π remains close to zero during stage 1, $\bar{\pi}$ is weakly identified and hard to learn

¹²We thank Klaus Adam for suggesting this exercise.

about. One of the purposes of a simultaneous reform is to strengthen identification of $\bar{\pi}$ by increasing ψ_{π} . The two-tier approach also postpones this until stage 2.

As shown in appendix E, optimized Taylor rules set $\bar{\pi} = 2$ percent per annum, $\psi_{\pi} = 0.15$, and $\psi_y = 0.15$ or 0.2 . Target inflation is therefore slightly higher than for simultaneous reforms, the inflation response is a bit weaker, and reaction to output growth is about the same. Learning is slower, the transition is longer and more volatile, and expected loss is substantially higher.

5.4 Single-equation learning

Agents in the baseline model exploit cross-equation restrictions on the ALM when estimating policy coefficients. This places a heavy computational burden on decision makers who are supposed to be boundedly rational. Appendix F lightens their burden by assuming that agents estimate equation (1) by recursive least squares with either constant or decreasing gain. All other aspects of the baseline specification remain the same.

Although estimates of policy coefficients sometimes differ from those in the baseline model, optimized Taylor rules are essentially the same. That results are similar for constant and decreasing gain algorithms is not surprising because the samples are small and the rates at which the respective algorithms discount past data are almost the same. That the results are similar to those for full-system learning means that cross-equation restrictions are less informative than in a full-information rational-expectations model. In the latter, private decision rules are predicated on knowledge of the true policy coefficients and therefore convey a lot of information about them. In a learning model, private decision rules are predicated on *estimates* of policy coefficients and encode less information about the true policy. Somewhat to our surprise, little is to be gained by exploiting cross-equation restrictions in the learning economy. Single-equation learning is almost as good.

6 Discussion of related literature

The literature on monetary policy with adaptive learning is vast, and good surveys can be found in Evans and Honkapohja (2009) and Gaspar, et al. (2011). Here we discuss a few papers that are especially relevant to ours.

A number of papers identify attributes of the equilibrium law of motion that

are influential in our analysis. For instance, Erceg and Levin (2003), Orphanides and Williams (2005), Milani (2006, 2007), and Slobodyan and Wouters (2012) examine new Keynesian models with adaptive learning and demonstrate that learning enhances inflation persistence.¹³ Orphanides and Williams emphasize that central banks should take steps to counteract this increase in persistence. In their model, this is done by reacting more aggressively to inflation. In ours, adverse initial beliefs that reaction coefficients are close to zero prevent the bank from responding aggressively, and inflation persistence is restrained – at least during the transition – by keeping reaction coefficients close to prior beliefs.

The conclusion that knowledge of the inflation target affords little benefit for stabilization when other aspects of monetary policy are uncertain also appears in Eusepi and Preston (2010), but the reasons supporting this conclusion differ. First and foremost, our notions of stability differ. Eusepi and Preston say that learning dynamics are stable if they converge to a rational-expectations equilibrium (*REE*), whereas we examine volatility during the transition. For a model with least-squares learning, they demonstrate that learning dynamics can fail to converge to *REE* when the central bank’s inflation target is known but other aspects of monetary policy are not, and this is what they mean when they write that knowledge of the inflation target affords little benefit for stabilization.

Eusepi and Preston also show that convergence to a rational-expectations equilibrium is restored if the central bank can credibly communicate the variables upon which nominal interest decisions are conditioned. This is the case that we examine, albeit with Bayesian learning. Our agents know the arguments and form of the policy rule, and their estimates eventually converge to the true policy coefficients.¹⁴ Even when policy estimates converge to the true parameters, however, uncertainty about reaction coefficients is more costly in our environment because this is what activates temporarily explosive transition dynamics. Our conclusion depends not on limiting beliefs but on the nature of the transition.

Among the above cited papers, only Erceg and Levin consider the implications for disinflation, and their analysis focuses on uncertainty about target inflation. None analyze how the potential for explosive transitional dynamics influences the central

¹³Milani and Slobodyan and Wouters estimate DSGE models with learning and show that structural sources of inflation persistence such as indexation lose empirical support when learning is introduced and that exogenous shocks become less persistent.

¹⁴We have no theorem to this effect, but this is what happens in the simulations.

bank's strategy.

Hagedorn (2011) examines optimal disinflation in a new Keynesian model with perfect credibility and rational expectations. He demonstrates that the transition path for the nominal interest rate is uniformly lower than it would be under the original inflation target. Hagedorn stops short of characterizing optimal policy under learning, however, commenting that this would require solving a challenging signal-extraction problem. His notion of optimality is broader than ours, but we tackle the signal-extraction problem. The price of extending the model in this direction was narrowing the family of policies to Taylor rules. Embracing a broader notion of optimality would be an important extension.

For a stylized, small-scale new Keynesian model, Gaspar, et al. (2006, 2011) show how to do this. They study optimal monetary policy in an environment where agents learn adaptively and the central bank takes the learning process into account when formulating its policy.¹⁵ The optimal rule shares some features of optimal policy under commitment and rational expectations, but commitment plays no role and the bank relies instead on its ability to influence estimated inflation persistence. Like Hagedorn, their notion of optimality is broader than ours, and they characterize the optimal policy by numerically solving a dynamic program. Their approach is feasible in models with a low-dimensional state vector but would run afoul of the curse of dimensionality in ours. We chose to enrich the economic environment at the expense of narrowing our focus to Taylor rules. Scaling their methods to a larger model would be another important extension.

7 Conclusion

We model transitional dynamics that emerge when a central bank tries to disinflate when operating without full transparency. The bank commits to a simple Taylor rule whose form is known but whose coefficients are not. Private agents learn about those coefficients via Bayesian updating. Under a McCallum timing protocol, locally-explosive dynamics can emerge when the new policy lacks transparency, making the transition highly volatile. The potential for locally explosive outcomes dominates expected loss and materially alters the bank's choice of feedback parameters relative to what would be chosen if operating under complete transparency and credibility.

¹⁵They do not study disinflation, however.

The bank copes by choosing feedback parameters close to the private sector's initial beliefs. Uncertainty about target inflation is secondary, and the bank can reduce average inflation substantially without generating much turbulence. Its ability to achieve greater stability by adjusting reaction coefficients is more limited.

Acknowledgement: *For comments and suggestions, we thank Klaus Adam, Martin Ellison, George Evans, Boyan Jovanovic, Thomas Sargent, Michael Woodford, a referee, and seminar participants at the Banque de France, the Centre for Dynamic Macroeconomics conference at the University of St Andrews, Duke, the European Central Bank, FRB Atlanta, FRB Philadelphia, FRB Richmond, the Federal Reserve Board, the Hungarian Central Bank, London Business School, the National Bank of Poland conference "DSGE and Beyond," the Norges Bank, the Norwegian School of Management, NYU, Oxford, Pompeu Fabra, Rutgers, the 2011 SCE and SED meetings, the Toulouse School of Economics, and UQAM. The views expressed here do not necessarily reflect the position of the Federal Reserve Banks of New York or Richmond or the Federal Reserve System.*

References

- [1] Ascari, G., 2004. Staggered prices and trend inflation: some nuisances. *Review of Economic Dynamics* 7, 642–667.
- [2] Ascari, G. and T. Ropele. 2011. Disinflations in a Medium-Scale DSGE Model: Money Supply versus Interest Rate Rules, unpublished manuscript.
- [3] Calvo, G. 1983. Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics* 12, 383-398.
- [4] Cho, I.K., N. Williams, and T.J. Sargent. 2002. Escaping Nash Inflation. *Review of Economic Studies* 69, 1-40.
- [5] Cogley, T., G. Primiceri, and T.J. Sargent. 2010. Inflation-Gap Persistence in the U.S. *American Economics Journal - Macroeconomics* 2, 43-69.
- [6] Cogley, T. and T.J Sargent, 2008. Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review* 49, 185-221.

- [7] 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.
- [8] Coibion, O. and Y. Gorodnichenko. 2011. Monetary policy, trend inflation and the great moderation: an alternative interpretation. *American Economic Review* 101, 341-370.
- [9] Erceg, C. and A. Levin. 2003. Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50 (4): 915-944.
- [10] Eusepi, S. and B. Preston. 2010. Central bank communication and expectations stabilization. *American Economic Journal: Macroeconomics* 2: 235-271.
- [11] Evans, G.W. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press: Princeton, N.J.
- [12] Evans, G.W. and S. Honkapohja. 2003. Expectations and the stability problem for optimal monetary policies. *Review of Economic Studies*, 70, 807-824.
- [13] Evans, G.W. and S. Honkapohja. 2009. Expectations, Learning and Monetary Policy: An Overview of Recent Research. in K. Schmidt-Hebbel and C. Walsh (eds.) *Monetary Policy under Uncertainty and Learning*, Central Bank of Chile: 27-76.
- [14] Gaspar, Vitor, Frank Smets and David Vestin. 2006. Adaptive Learning, Persistence, and Optimal Monetary Policy. *Journal of the European Economic Association* 4, 376-385.
- [15] Gaspar, Vitor, Frank Smets and David Vestin. 2011. Inflation Expectations, Adaptive Learning and Optimal Monetary Policy, in B. Friedman and M. Woodford (eds) *Handbook of Monetary Economics*, vol 3B, North-Holland: 1055-1095.
- [16] Goodfriend, M. and R.G. King, 2005. The Incredible Volcker Disinflation. *Journal of Monetary Economics*, 52 (5): 981-1016.
- [17] Justiniano, A., G. Primiceri, and A. Tambalotti. 2010. Investment Shocks and Business Cycles. *Journal of Monetary Economics*, 57(2): 132-145.

- [18] 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.
- [19] Levin, A.T. and J.C. Williams. 2003. Robust monetary policy with competing reference models. *Journal of Monetary Economics*, 50 (5): 945-975.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] Mertens, Elmar. 2009a. Managing Beliefs about Monetary Policy under Discretion. Unpublished manuscript, Federal Reserve Board.
- [24] Mertens, Elmar. 2009b. Discreet Commitments and Discretion of Policymakers with Private Information. Unpublished manuscript, Federal Reserve Board.
- [25] Milani, Fabio. 2006. A Bayesian DSGE Model with Infinite-Horizon Learning: Do ‘Mechanical’ Sources of Persistence Become Superfluous? *International Journal of Central Banking* 2 (3): 87–106.
- [26] Milani, Fabio. 2007. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54 (7): 2065-82
- [27] Hagedorn, Marcus, 2011. Optimal disinflation in new Keynesian models. *Journal of Monetary Economics* 58: 248-261.
- [28] 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*

- [29] Orphanides, A. and J.C. Williams, 2007. Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54: 1406-1435.
- [30] 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.
- [31] Sbordone, A.M. 2007. Inflation persistence: Alternative interpretations and policy implications. *Journal of Monetary Economics* 54, 1311–1339.
- [32] Sims, C.A. 2001. Solving Linear Rational Expectations Models. *Computational Economics* 20, 1-20.
- [33] Slobodyan, Sergey and Raf Wouters, 2012. Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models. *American Economic Journal: Macroeconomics* 4(2): 65-101
- [34] Smets, F. and R. Wouters. 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97 (3): 586-606.
- [35] Woodford, M. 1999. Optimal Monetary Policy Inertia, NBER wp 7261.
- [36] Woodford, M. 2003. *Interest and Prices*. Princeton University Press: Princeton NJ.