Optimized Taylor Rules for Disinflation When Agents are Learning

Timothy Cogley[a,*], Christian Matthes[b], Argia M. Sbordone[c,†]

[a] Department of Economics, New York University

[b] Research Department, Federal Reserve Bank of Richmond

[c] Research Department, Federal Reserve Bank of New York

Received June 2012; Received in Revised Form January 2015

Abstract

When private agents must learn a new policy rule, an optimal simple Taylor rule for disinflation differs substantially from that under full information. The central bank can reduce target inflation without much difficulty, but adjusting reaction coefficients on lagged inflation and output is more costly. Temporarily explosive dynamics emerge when there is substantial disagreement between perceived and actual feedback parameters, making the transition highly volatile. The bank copes by choosing reaction coefficients close to the private sector's prior mode, thereby sacrificing long-term performance in exchange for achieving lower transitional volatility.

Keywords: Inflation, monetary policy, learning, policy reforms, transitions

JEL classification: E31, E52

*Corresponding author: Timothy Cogley, Department of Economics, New York University, 19 W 4th St., New York, NY 10012. Tel.: 212-992-8679. Email: tim.cogley@nyu.edu.

†For comments and suggestions, we thank Klaus Adam, Martin Ellison, George Evans, Boyan Jovanovic, Ricardo Reis, Thomas Sargent, Michael Woodford, a referee, and participants in numerous seminars and conferences. 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.

1. Introduction

Imagine a newly-appointed central bank governor who inherits high inflation 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. Suppose that the new governor's preferences differ from those of his predecessor and that he wants to disinflate. He seeks an optimal Taylor-type rule and takes learning into account when choosing 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.[1] Goodfriend

[1] Schorfheide (2005) also develops a model in which agents must learn about target inflation. In his model, target inflation follows an exogenous Markov-switching process, and estimates confirm that learning is important for fitting data from the early 1980s. Andolfatto and Gomme (2003) explain the Canadian

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, our question is normative and focuses on how learning alters the central bank's choice of policy. 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 force that arises when a new policy must be learned and to describe how it affects the bank's choices.

The problem is studied in the context of a dynamic new Keynesian model modified in two ways. Following Ascari (2004) and Sbordone (2007), target inflation need not be zero. In addition, Bayesian learning replaces rational expectations. 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.

Our paper contributes to a literature on how to design monetary policy rules when agents are learning. Bullard and Mitra (2002) and Evans and Honkapohja (2003a,b) examine how to specify monetary policy rules so that learning converges to rational expectations and the rational-expectations equilibrium (REE) is determinate. In our model, both conditions are satisfied for the family of simple rules under consideration.[2] We refine the analysis by considering how transition dynamics affect the choice of policy coefficients.

Accounting for transition volatility substantially alters the bank's choice. Compared with the old regime, the optimal simple rule under full information has a lower long-run inflation

experience using a closely related model with high and low money-growth states.

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

target and a higher reaction coefficient on inflation. The optimal simple rule under learning reduces target inflation by almost as much but reacts much less aggressively to inflation. Indeed, the inflation reaction coefficient is only slightly higher than in the old regime.

The reason why the bank's choice differs 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 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 transitional volatility.

As in Eusepi and Preston (2010), uncertainty about policy feedback parameters matters more than uncertainty about target inflation.[3] In our model, the bank always achieves low average inflation. 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.

In this respect, our conclusions differ from those of Orphanides and Williams (2005). They and others examine new Keynesian models with adaptive learning and demonstrate that learning enhances inflation persistence.[4] Orphanides and Williams emphasize that central banks should take steps to counteract this increase in persistence, reacting *more*

[3] For a model with least-squares learning, Eusepi and Preston (2010) study various communications strategies: the central bank credibly communicates (i) target inflation, (ii) the variables on which policy decisions are conditioned, or (iii) the precise details of policy. They demonstrate that the Taylor principle plus strategies (ii) or (iii) guarantees convergence to *REE*, while the Taylor principle plus (i) does not. Our scenario is like case (ii): our agents know the form and arguments of the policy rule, and estimates of policy coefficients converge to the true parameters.

[4] E.g., see Erceg and Levin (2003), Milani (2006, 2007), and Slobodyan and Wouters (2012).

aggressively to inflation than they would under full information. Like us, Orphanides and Williams study optimal simple Taylor rules, but they only consider the consequences of alternative policies once the economy reaches its ergodic distribution.[5] Our conclusions differ because our loss function also penalizes transitional volatility. Concerns about locally-explosive dynamics outweigh other considerations.

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, 2003a,b). Models in that tradition typically assume that agents use reduced-form statistical representations such as vector autoregressions (VARs) for forecasting and that agents update parameter estimates by recursive least squares. In contrast, the agents who inhabit our model build structural models of the economy and update beliefs via Bayes' theorem. Our approach is useful for showing how a bank's policy choice depends on agents' priors, but otherwise it is not critical. Our insights are robust to other forms of learning.

Hagedorn (2011) examines optimal disinflation in a new Keynesian model with perfect credibility and rational expectations. He 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, Smets, and Vestin (2006) 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. The optimal rule shares some features of optimal policy under commitment

[5] They consider a model with constant gain learning, so their agents never fully learn.

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. Although their approach is feasible in models with a low-dimensional state vector, it would run afoul of the curse of dimensionality in ours. We chose to enrich the economic environment at the expense of narrowing the focus to Taylor rules. Scaling their methods to larger models would be another important extension.

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

We work with 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 an intertemporal *IS* curve and an Ascari-Sbordone version of the new Keynesian Phillips curve. A log-linearized version is presented here. Details about how this representation was derived can be found in appendix A.[6]

2.2.1. Monetary policy

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

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

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 rate smoothing (e.g. Woodford (1999)). In addition,

[6] Appendices are posted online in the JME's supplemental material archive.

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. 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

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

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

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

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 central bank minimizes a discounted quadratic loss function,

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

that penalizes variation in inflation and the output gap, and 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_y, taking private-sector learning into account.[7]

2.2.2. Behaviorial 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

[7] The 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.

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

Private agents also behave as anticipated-utility modelers, treating the current estimate ψ_t as if it were known with certainty (Kreps 1998). 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.

Last but not least, our agents adopt the posterior mode as their point estimate. In other words, they do not follow the full Bayesian decision-theoretic route of choosing a point estimate that minimizes an expected loss function implied by their utitility function. The choice of the posterior mode is somewhat arbitrary but seems quite plausible.

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

$$U_t = b_t \log\left(C_t - \eta C_{t-1}\right) - \chi_t \frac{H_t^{1+\nu}}{1+\nu}, \tag{5}$$

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,

$$y_t - \overline{y}_t = \xi_t - \xi - E_t^* \left[\xi_{t+1} - \xi - \left(y_{t+1} - \overline{y}_t\right) - \left(\gamma_{t+1} - \gamma\right) + i_t - \pi_{t+1} - r\right], \tag{6}$$

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

$$\xi_t - \xi \equiv \xi_1\left(y_t - \overline{y}_t\right) + \xi_2\left[y_{t-1} - \overline{y}_t - \left(\gamma_t - \gamma\right) + \beta E_t^*\left(y_{t+1} - \overline{y}_t + \gamma_{t+1} - \gamma\right)\right] + \varepsilon_{yt}. \tag{7}$$

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 $\overline{y}_t$ 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 concerns the expansion point. As mentioned above, agents expand around the perceived steady state $\overline{y}_t$ instead of the actual steady state $\overline{y}$. In addition, the anticipated-utility assumption implies that $E_t^*\overline{y}_{t+1} = \overline{y}_t$, explaining the appearance of $\overline{y}_t$ on the right-hand side of equations (6) and (7). 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 is denoted E_t.[8] Finally, two shocks appear, a white-noise shock ε_{yt} and a persistent shock γ_t to the growth rate of technology, $\gamma_t = \left(1 - \rho_\gamma\right)\gamma + \rho_\gamma \gamma_{t-1} + \varepsilon_{\gamma t}$.

2.2.4. A new-Keynesian Phillips curve

A continuum of monopolistically competitive firms produce differentiated intermediate goods that are sold to a final-goods producer. Following Calvo (1983), intermediate-goods producers reset their prices at random intervals. 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,

$$\begin{aligned}
\hat{\pi}_t &= E_t^* \left\{ [\beta_t + \gamma_{1t}(\theta - 1)]\hat{\pi}_{t+1} + \gamma_{1t}\phi_{t+1} \right\} + \kappa_t \hat{y}_t + \varsigma_t \hat{\delta}_t - \widetilde{\kappa}_t \hat{\xi}_t + u_t + \varepsilon_{\pi t}, && (8)\\
\phi_t &= \gamma_{2t} E_t^* [(\theta - 1)\hat{\pi}_{t+1} + \phi_{t+1}], && (9)\\
\hat{\delta}_t &= \lambda_{1t}\hat{\pi}_t + \lambda_{2t}\left(\hat{\delta}_t - \Delta\overline{\delta}_t\right), && (10)
\end{aligned}$$

where gap variables are defined as $\hat{\pi}_t \equiv \pi_t - \bar{\pi}_t$, $\hat{y}_t \equiv y_t - \overline{y}_t$, $\hat{\delta}_t \equiv \delta_t - \overline{\delta}_t$, and $\hat{\xi}_t \equiv \xi_t - \xi$. The NKPC parameters β_t, γ_{1t}, γ_{2t}, κ_t, $\widetilde{\kappa}_t$, ς_t, λ_{1t}, and λ_{2t} are defined in table 1.

Table 1 here

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

[8] We 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 $E_t x_{t+j} \neq E_t^* x_{t+j}$.

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.

Second, a variable

$$\delta_t \equiv \ln\left(\int_0^1 \left(p_t\left(i\right)/P_t\right)^{-\theta} di\right), \tag{11}$$

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, an intermediate variable ϕ_t that has no interesting economic interpretation is added along with equation (9). This is simply a device for obtaining a convenient representation.

Finally, two cost-push shocks are present, a white-noise shock $\varepsilon_{\pi t}$ and a persistent shock u_t that follows an $AR(1)$ process, $u_t = \rho_u u_{t-1} + \varepsilon_{ut}$.

2.2.5. Calibration

Parameters of the pricing model are taken from estimates in Cogley and Sbordone (2008) and are set at $\alpha = 0.6$, $\beta = 0.99$, $\theta = 10$. Preference parameters are calibrated as follows. The parameter $1/\nu$ is 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 compromise between the two, setting the Frisch elasticity equal to 2 ($\nu = 0.5$). This seems 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).

The calibration of loss-function parameters is also standard. 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 $\lambda_y = 1/16$. The parameter λ_i is set 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.

Turning to parameters governing the shocks, γ is set to 0, thereby abstracting from average growth. For the persistent shocks u_t and γ_t, estimates are taken from Cogley, Primiceri, and Sargent (2010), $\rho_u = 0.4$, $\sigma_u = 0.0012$, $\rho_\gamma = 0.27$, $\sigma_\gamma = 0.005$. Last but not least, the standard deviations of the white noise shocks ε_{yt} and $\varepsilon_{\pi t}$ are set equal to $\sigma_\pi = \sigma_y = 0.0025$.

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 explains how this is done. We first describe the 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, the ALM can be used to derive a likelihood function. Agents estimate policy coefficients by combining this likelihood function with their prior.

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 represented as a system of linear expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_t S_{t-1} + D_t \widetilde{\varepsilon}_t, \tag{12}$$

where S_t is the model's state vector, $\widetilde{\varepsilon}_t$ 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). These matrices have time subscripts because they depend on estimates of the policy coefficients ψ_t. The PLM is the reduced-form VAR associated with (12),

$$S_t = F_t S_{t-1} + G_t \widetilde{\varepsilon}_t, \tag{13}$$

where F_t solves $B_t F_t^2 - A_t F_t + C_t = 0$ and $G_t = (A_t - B_t F_t)^{-1} D_t$. As in a conventional rational-expectations model, (13) serves two functions, describing agents' current-quarter plans and how they forecast future outcomes.

3.2. The actual law of motion

To find the ALM, stack the actual policy rule (equation 1) with the equations governing private sector behavior. This results in another system of expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_{at} S_{t-1} + D_t \varepsilon_t. \tag{14}$$

The state vector and the matrices A_t, B_t, and D_t are the same as in (12). 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).

Since outcomes are determined in accordance with agents' plans (equation 13), they depend on the perceived shocks $\tilde{\varepsilon}_t$. A relation between perceived and actual innovations can be found by subtracting (14) from (12),

$$D_t \tilde{\varepsilon}_t = D_t \varepsilon_t + (C_{at} - C_t) S_{t-1}. \tag{15}$$

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

$$S_t = H_t S_{t-1} + G_t \varepsilon_t, \tag{16}$$

where $H_t = F_t + (A_t - B_t F_t)^{-1}(C_{at} - C_t)$. 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.[9]

3.3. The PLM is the perceived ALM

The ALM and PLM are both $VAR(1)$ processes with conditionally Gaussian innovations. Under the ALM, the conditional mean and variance are[10]

$$m_{t|t-1}(\psi_{true}) = H_t(\psi_{true}) S_{t-1}, \quad V_{t|t-1}(\psi_{true}) = G_t V_\varepsilon(\psi_{true}) G_t', \tag{17}$$

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

$$\tilde{m}_{t|t-1}(\psi_t) = F_t S_{t-1}, \quad \tilde{V}_{t|t-1}(\psi_t) = G_t V_\varepsilon(\psi_t) G_t'. \tag{18}$$

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.[11]

[9] 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 would choose one of them. However, this kind of multiplicity never occurs in our simulations.

[10] According 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.

[11] Among 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 $t+j+1$

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

$$p(X^t|\psi) = \prod_{j=1}^{t} p(X_j|X^{j-1}, \psi). \tag{19}$$

Since private agents know the ALM up to the unknown policy parameters, they can use it to evaluate the terms on the right-hand side of (19). According to the ALM, X_t is conditionally normal with mean and variance

$$m^X_{t|t-1}(\psi) = e_X H_t(\psi) S_{t-1}, \quad V^X_{t|t-1}(\psi) = e_X G_t V_\varepsilon(\psi) G_t' e_X', \tag{20}$$

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

$$\ln p(X^t|\psi) = -\frac{1}{2} \sum_{j=1}^{t} \{ \ln |V^X_{j|j-1}(\psi)| + [X_j - m^X_{j|j-1}(\psi)]' \left(V^X_{j|j-1}(\psi)\right)^{-1} [X_j - m^X_{j|j-1}(\psi)]\}. \tag{21}$$

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

hold in expectation at t.

the posterior mode, $\psi_t = \arg\max \left(\ln p(X^t|\psi) + \ln p(\psi) \right)$.

Notice that agents take into account that past outcomes were influenced by past beliefs. 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.

4. Quantitative analysis

A new governor appears at date 0 and formulates a policy 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 $\sigma_i = 0$.

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, the old regime is calibrated to match estimates of the policy rule for the period 1966.Q1-1981.Q1. We assume that the policy rule for that period had the same functional form as in equation (1) and estimate $\bar{\pi}, \psi_\pi, \psi_y$, and σ_i^2 by OLS. Point estimates and standard errors are reported in table 2.

Table 2 here

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 proportion of the 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

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). Learning makes the model nonlinear, however, so expected loss is evaluated numerically. A grid of values is specified for $\bar{\pi}$, ψ_π, and ψ_y. Then, for each node on the grid, 100 sample paths are simulated, with private-sector estimates ψ_t updated 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.[12]

Figure 1 here

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,

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

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 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 in about a year. After 4 quarters, 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, making inflation weakly persistent. The absence of indexation also explains why the bank seeks a substantial rollback in the price level.

Figure 2 here

Under full information, the economy is highly fault tolerant[13] with respect to policies away from the optimum. Figure 2 portrays iso-expected loss contours as a function of $\bar{\pi}, \psi_\pi$, 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 $\bar{\pi} = 0$, 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

[13] Levin and Williams (2003) introduced the term "fault tolerance" to describe the extent to which expected loss increases as policies move away from the optimum.

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

Private agents initially anticipate a continuation of the old regime, and their priors are calibrated using the estimates of policy coefficients for 1966.Q1-1981.Q1 shown in table 2. We assume that their priors are independent across coefficients and that 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 2. 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 table 2. The results are shown in figure 3.

Figure 3 here

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 spread across 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 growth. Overcoming that skepticism will be a major challenge for the central bank.

Figure 4 portrays iso-expected loss contours as a function of $\bar{\pi}, \psi_\pi, \psi_y$. 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.

Figure 4 here

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 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. 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 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 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 $\bar{\pi}$, the figure just shows the $\bar{\pi} = 0$ case. 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.

Figure 5 here

To locate the optimum under learning, we search on a finer grid in the southwest quadrant of the (ψ_π, ψ_y) space. Isoloss 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 in the explosive region, the transition would be initially very turbulent. Furthermore, since the private sector is prejudiced against large values of ψ_π, explosive eigenvalues would remain active for a long time. 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 volatile when this occurs, causing expected loss to rise. For an effective stabilization, the bank must choose a value for ψ_π that 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.

Figure 6 here

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. As shown in appendix C, the sacrifice ratio under learning is comparable to that in a version of our model with adaptive expectations. Under monetary policies optimized for that environment,[14] inflation falls permanently by amounts ranging from 2.6 to 4.6 percent, depending on how the adaptive expectations operator is calibrated, with cumulative output losses of 4.5 to 6.7 percent, implying sacrifice ratios of 1.4 to 1.7. According to Ascari and Ropele (2013), estimates of the sacrifice ratio for a wide range of disinflations lie between 0.5 and 3, so those under learning and adaptive expectations are in the right ballpark.

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

[14] A comparison that holds monetary policy constant across the two models is difficult because policies optimized for one environment work badly in the other. This is why we compare sacrifice ratios under policies optimized for each. The sacrifice ratios therefore differ not only because of the expectations mechanisms but also because of the policy rule. See appendix C for further details.

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.

Figure 7 here

4.5. Intuition about transition dynamics

To develop intuition, we turn to a stripped-down example that can be solved by hand,

$$\pi_t = \beta E_t^* \pi_{t+1} + x_t + u_t, \tag{22}$$

$$x_t = -\psi \pi_{t-1} + \varepsilon_{xt}, \tag{23}$$

$$u_t = \rho_u u_{t-1} + \varepsilon_{ut}. \tag{24}$$

Equation (22) is a stylized version of the NKPC, with x_t representing an abstract policy instrument, equation (23) is a policy rule, and equation (24) is the law of motion for the cost-push shock. The innovations ε_{ut} and ε_{xt} are iid normal with mean zero and variances σ_u^2 and σ_x^2, respectively. The perceived policy is $x_t = -\psi_t \pi_{t-1} + \tilde{\varepsilon}_{xt}$, where ψ_t again represents a beginning-of-period t estimate of ψ, and the perceived policy shock is $\tilde{\varepsilon}_{xt} = \varepsilon_{xt} + (\psi_t - \psi)\pi_{t-1}$. Agents believe that $\tilde{\varepsilon}_{xt}$ is iid normal with mean zero and variance σ_{xt}^2.

As shown in appendix D, the PLM is a $VAR(1)$ for (π_t, u_t),

$$\begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = \begin{bmatrix} \lambda_{1t} & \frac{\rho_u}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & \rho_u \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} \frac{1}{\beta\lambda_{2t}} & \frac{1}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{\varepsilon}_{xt} \\ \varepsilon_{ut} \end{bmatrix}, \tag{25}$$

where λ_{1t} and λ_{2t} are reciprocals of the roots of $(1 - \beta^{-1}L_t - \beta^{-1}\psi_t L^2) = 0$. The eigenvalues of the PLM autoregressive matrix are λ_{1t} and ρ. By construction, both lie inside the unit circle. The ALM is also a $VAR(1)$ for (π_t, u_t),

$$\begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = \begin{bmatrix} \lambda_{1t} + \frac{\psi_t - \psi}{\beta\lambda_{2t}} & \frac{\rho_u}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & \rho_u \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} \frac{1}{\beta\lambda_{2t}} & \frac{1}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_{xt} \\ \varepsilon_{ut} \end{bmatrix}. \quad (26)$$

The eigenvalues of the ALM autoregressive matrix are $\mu_{1t} = \lambda_{1t} + (\psi_t - \psi)/\beta\lambda_{2t}$ and $\mu_{2t} = \rho_u$. The latter always lies inside the unit circle, but the former might lie inside or outside, depending on the difference between ψ_t and ψ.

To illustrate the model's properties, we plug in some numbers. To represent a scenario in which the old regime involved weak feedback to inflation, ψ_0 is set to 0.1. As before, $\beta = 0.99$ and $\rho_u = 0.4$. These values imply $\lambda_{10} = -0.09$ and $\lambda_{20} = 1.1$.

Figure 8 depicts responses of inflation to a unit cost-push shock for various choices of ψ. Solid curves portray what would happen if ψ were known at date 0. Under full information, the response of inflation becomes more muted as ψ increases, largely because policy guides expectations. When price setters expect a strong countervailing action against future inflation, their incentive to raise prices today is diminished. Thus, the impact effect of a cost-push shock declines as ψ increases. Furthermore, to the extent that an inflation gap does open, a stronger countervailing action would close it more quickly, making inflation less persistent.

Figure 8 here

Dashed and dotted lines depict impulse responses under the PLM and ALM, respectively, for the initial estimate ψ_0. Although both will change shape as ψ_t is updated, we temporarily freeze beliefs in order to build intuition.

Impulse responses under the PLM differ in two respects from those under full information: the impact effect is greater, and the impulse response function decays more slowly. Both follow from the fact that the expectations-management channel that is so important under full information is initially inoperative under learning. The PLM depends on ψ_0, not ψ, and since $\psi_0 < \psi$ agents initially expect a weaker countervailing policy response. It follows that their incentive to raise prices today is greater, amplifying the initial rise in inflation. Beyond lag zero, the impulse response function under the PLM can be interpreted as the revision of the private sector's inflation forecast in response to a cost-push shock. Because price setters systematically underestimate future policy reactions, their inflation forecasts are higher. Hence price setters expect the inflation gap to remain open longer.

A wedge also appears between ALM and PLM. This emerges because a cost-push shock endogenously creates perceived policy shocks, $\tilde{\varepsilon}_{xt} = (\psi_0 - \psi)\pi_{t-1}$.[15] Because price setters do not anticipate future $\tilde{\varepsilon}_{xt}$, perceived policy shocks do not affect the PLM, but they do matter for the ALM. When the difference between ψ_0 and ψ is small, as in the top row of figure 1, the resulting policy shocks are also small and their consequences are minor. Their effects are more apparent when $|\psi_0 - \psi|$ is greater. For instance, in the lower left panel, the negative local feedback between lagged inflation and perceived policy shocks creates a damped oscillation. Inflation rises at impact. At date 1, the central bank reduces x by more than the private sector expects, causing inflation to fall sharply. Because π_1 is negative, the central bank increases x_2 by more than the private sector expects, causing inflation to rise. With positive inflation at date 2, the bank tightens sharply in period 3, causing inflation to fall, and so on. When $|\psi_0 - \psi|$ is sufficiently large, this negative feedback can be strong enough to create an explosive oscillation, as shown in the lower right panel. Thus, inflation is negatively autocorrelated, and explosive dynamics emerge when ψ and ψ_0 are sufficiently far apart.[16]

From the private sector's point of view, the central bank behaves like a madman in a

[15] For these impulse response functions, the true policy shocks are zero.

[16] Since $\psi_0 < \psi$, μ_{10} is farther below zero than λ_{10}. When $|\psi_0 - \psi|$ is small, $|\mu_{10} - \lambda_{10}|$ is also small, and the ALM is close to the PLM. As $|\psi_0 - \psi|$ increases, μ_{10} increases in absolute value and imparts more negative autocorrelation to inflation. The oscillations are damped as long as μ_{10} remains inside the unit circle, but μ_{10} can fall below -1 if ψ is sufficiently far above ψ_0, making the oscillations explosive. Because $|\psi_t - \psi|$ is small in the neighborhood of the *REE*, locally-explosive dynamics vanish in that neighborhood.

shower, dialing the temperature control all the way down when the water is too hot, then all the way up when the temperature is too cold. As beliefs are updated, they learn that the bank is behaving systematically, and the expectations-management channel comes gradually into play. But during the transition, endogenous perceived policy shocks account for high volatility and negative autocorrelation.

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 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, appendix E contrasts the backward-looking Taylor rule in equation (1) with one involving contemporaneous feedback to inflation and output growth,

$$i_t - i_{t-1} = \psi_\pi(\pi_t - \bar{\pi}) + \psi_y(y_t - y_{t-1}) + \varepsilon_{it}. \tag{27}$$

Because actual central banks cannot observe current quarter output or the price level, they could not 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 E, 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 transition is shorter and less volatile than under the backward-looking rule, and the sacrifice ratio is about the same.

Expectations are still sticky, and that is why the contemporaneous rule optimized for learning differs from the full-information optimum. The difference between outcomes under the contemporaneous and backward-looking rules illustrates the quantitative force of locally-explosive dynamics. Sticky expectations and temporarily-explosive dynamics are both important ingredients in the baseline model.

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 σ_i were zero and known with certainty, the signal extraction problem would unravel, with agents perfectly inferring the other three policy coefficients after three periods. This does not happen in our model because agents are uncertain about σ_i, and this is enough to preserve a nontrivial signal-extraction problem. Appendix F confirms that the optimized rule in a $\sigma_i = 0$ economy is similar to that in the benchmark specification. 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 approach

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 G 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.

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 by allowing beliefs about ψ_π and ψ_y to harden around old-regime values during stage 1. 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 about. One of the purposes of a simultaneous reform is to strengthen identification of $\bar{\pi}$ by increasing ψ_π.

As shown in appendix G, target inflation is 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 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 H 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 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. Somewhat to our surprise, single-equation learning is almost as good. Little is to be gained by exploiting cross-equation restrictions.

6. Conclusion

When the private sector must learn about a change in monetary policy, transitional volatility matters for policy design. 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. The potential for locally explosive outcomes dominates expected loss and materially alters the bank's choice of policy coefficients relative to what would be chosen if operating under full information. Locally-unstable dynamics emerge when there is substantial disagreement between actual and perceived feedback parameters. 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.

Although we believe our model has some relevance for understanding the Volcker disinflation, we are reluctant to push it hard as a positive explanation of that episode. The central bank in our model knows the structure of the economy and how agents learn, and we suspect that the Fed under Volcker was not quite so knowledgeable or sophisticated. In addition, our analysis abstracts from features such as Carter's credit controls and political economy factors that were important then. Instead, our contribution is to highlight the importance of accounting for transitional volatility when agents must learn about a new monetary-policy rule.

References

Andolfatto, D., Gomme, P., 2003. Monetary policy regimes and beliefs. International Economic Review 44, 1-30.

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

Ascari, G., Ropele, T., 2013. Disinflation effects in a medium-scale New Keynesian model: money supply rule versus interest rate rule. European Economic Review 61, 77-100.

Bullard, J., Mitra, K., 2002. Learning about monetary policy rules. Journal of Monetary Economics 49, 1105-1129.

Calvo, G., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12, 383-398.

Cho, I.K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. Review of Economic Studies 69, 1-40.

Cogley, T., Primiceri, G., Sargent, T.J., 2010. Inflation-gap persistence in the U.S. American Economics Journal: Macroeconomics 2, 43-69.

Cogley, T., Sargent, T.J., 2008. Anticipated utility and rational expectations as approximations of Bayesian decision making. International Economic Review 49, 185-221.

Cogley, T., Sbordone, A.M., 2008. Trend inflation, indexation, and inflation persistence in the new Keynesian Phillips curve. American Economic Review 98, 2101-2126.

Coibion, O., Gorodnichenko, Y., 2011. Monetary policy, trend inflation and the great moderation: an alternative interpretation. American Economic Review 101, 341-370.

Erceg, C., Levin, A., 2003. Imperfect credibility and inflation persistence. Journal of Monetary Economics 50, 915-944.

Eusepi, S., Preston, B., 2010. Central bank communication and expectations stabilization. American Economic Journal: Macroeconomics 2, 235-271.

Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press: Princeton, N.J.

Evans, G.W., Honkapohja, S., 2003a. Expectations and the stability problem for optimal monetary policies. Review of Economic Studies 70, 807-824.

Evans, G.W., Honkapohja, S., 2003b. Adaptive learning and monetary policy design. Journal of Money Credit and Banking 35, 1045-1072.

Gaspar, V., Smets, F., Vestin, D., 2006. Adaptive learning, persistence, and optimal monetary policy. Journal of the European Economic Association 4, 376-385.

Goodfriend, M., King, R.G., 2005. The incredible Volcker disinflation. Journal of Monetary Economics 52, 981-1016.

Hagedorn, M., 2011. Optimal disinflation in new Keynesian models. Journal of Monetary Economics 58, 248-261.

Justiniano, A., Primiceri, G., Tambalotti, A., 2010. Investment shocks and business cycles. Journal of Monetary Economics 57, 132-145.

Kreps, D., 1998. Anticipated utility and dynamic choice. In: Jacobs, D.P., Kalai, E., Kamien, M. (Eds.), Frontiers of Research in Economic Theory, Cambridge University Press, Cambridge, pp. 242–74.

Levin, A., Williams, J.C., 2003. Robust monetary policy with competing reference models. Journal of Monetary Economics 50, 945-975.

Marcet, A., Sargent, T.J., 1989a. Convergence of least-squares learning mechanisms in self-referential linear stochastic models. Journal of Economic Theory 48, 337-368.

Marcet, A., Sargent, T.J., 1989b. Convergence of least-squares learning in environments with hidden state variables and private information. Journal of Political Economy 97, 1306-1322.

McCallum, B.T., 1999. Issues in the design of monetary policy rules. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1C. Elsevier, Amsterdam, pp. 1483-1530.

Mertens, E., 2009a. Managing beliefs about monetary policy under discretion. Unpublished manuscript, Federal Reserve Board.

Mertens, E., 2009b. Discreet commitments and discretion of policymakers with private information. Unpublished manuscript, Federal Reserve Board.

Milani, F., 2006. A Bayesian DSGE model with infinite-horizon learning: do 'mechanical' sources of persistence become superfluous? International Journal of Central Banking 2, 87–106.

Milani, F., 2007. Expectations, learning and macroeconomic persistence. Journal of Monetary Economics 54, 2065-82.

Orphanides, A., Williams, J.C., 2005. Imperfect knowledge, inflation expectations and monetary policy. In: Bernanke, B., Woodford, M. (eds.), The Inflation Targeting Debate. University of Chicago press, Chicago, pp. 201-234.

Orphanides, A., Williams, J.C., 2007. Robust monetary policy with imperfect knowledge. Journal of Monetary Economics 54, 1406-1435.

Preston, B., 2006. Adaptive learning, forecast-based instrument rules and monetary policy. Journal of Monetary Economics 53, 507-535.

Sargent, T.J., 1982. The ends of four big inflations. In: Hall, R. (Ed.), Inflation: Causes and Effects. University of Chicago Press, pp. 41-97.

Sbordone, A.M., 2007. Inflation persistence: Alternative interpretations and policy implications. Journal of Monetary Economics 54, 1311–1339.

Schorfheide, F., 2005. Learning and monetary policy shifts. Review of Economic Dynamics 8, 392-419.

Slobodyan, S., Wouters, R., 2012. Learning in a medium-scale DSGE model with expectations based on small forecasting models. American Economic Journal: Macroeconomics 4, 65-101.

Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles: A Bayesian DSGE approach. American Economic Review 97, 586-606.

Woodford, M., 1999. Optimal Monetary Policy Inertia, NBER working paper 7261.

Woodford, M., 2003. Interest and Prices. Princeton University Press, Princeton NJ.

Table 1: NKPC parameters

$\beta_t = \beta(1+\bar{\pi}_t)$	$\gamma_{1t} = \beta\bar{\pi}_t[1-\alpha(1+\bar{\pi}_t)^{\theta-1}]$	$\gamma_{2t} = \alpha\beta(1+\bar{\pi}_t)^{\theta-1}$	$\varsigma_t = \nu\widetilde{\kappa}_t,$
$\kappa_t = (1+\nu)\,\widetilde{\kappa}_t,$	$\widetilde{\kappa}_t = \frac{[1-\alpha(1+\bar{\pi}_t)^{\theta-1}][1-\alpha\beta(1+\bar{\pi}_t)^{\theta}]}{\alpha(1+\bar{\pi}_t)^{\theta-1}}$	$\lambda_{1t} = \frac{\alpha\theta\bar{\pi}_t(1+\bar{\pi}_t)^{\theta-1}}{(1-\alpha(1+\bar{\pi}_t)^{\theta-1}}$	$\lambda_{2t} = \alpha(1+\bar{\pi}_t)^{\theta}$

Note: This table records the relationship between beliefs about trend inflation $\bar{\pi}_t$ and the parameters of the NKPC (equations 8-10). 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 $1/\nu$.

Table 2: Estimates of Policy Coefficients before the Volcker Disinflation

$\bar{\pi}$	ψ_π	ψ_y	σ_i	R^2
0.0116	0.043	0.12	0.0033	0.12
(0.013)	(0.08)	(0.04)	(0.01)	

Note: This table reports OLS estimates of a backward-looking Taylor rule in difference form (equation 1). $\bar{\pi}$ represents target inflation, ψ_π and ψ_y are feedback parameters on lagged inflation and output growth, and σ_i is the standard deviation of the policy shock. The nominal interest rate is measured by the federal funds rate, and inflation and output growth are measured by the rates of change in the chain-weighted price index for personal consumption expenditures and in the real gross domestic product, respectively. The data are quarterly, and the sample covers the period 1966.Q1-1981.Q1. OLS standard errors are shown in parentheses.

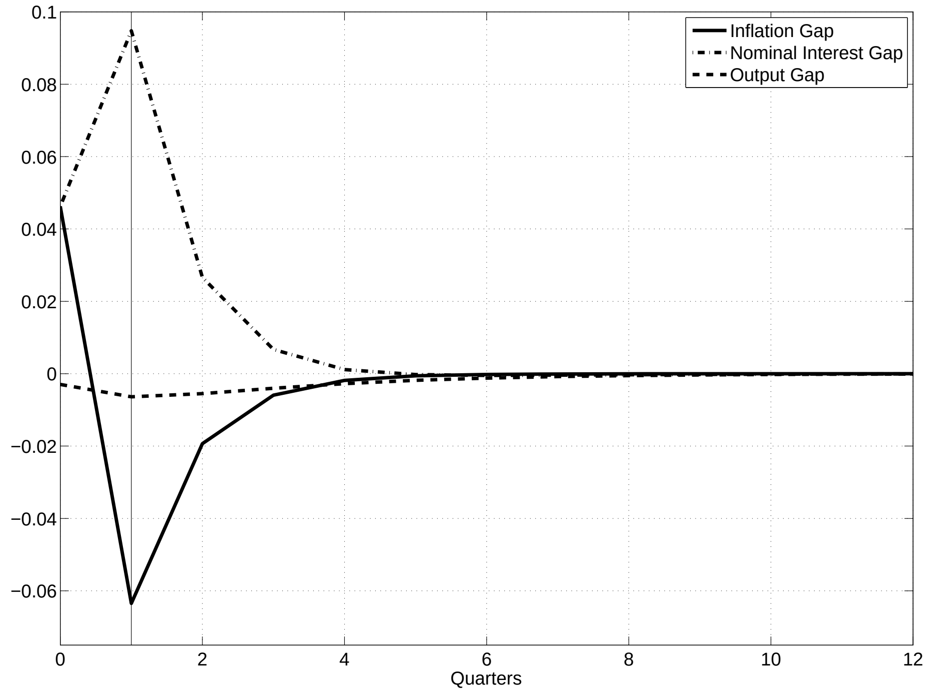

Figure 1: Responses of inflation, output, and interest rate gaps under full information. Values shown at date 0 depict differences between steady states in the old and new regimes. A Taylor rule optimized for full information is introduced at date 1.

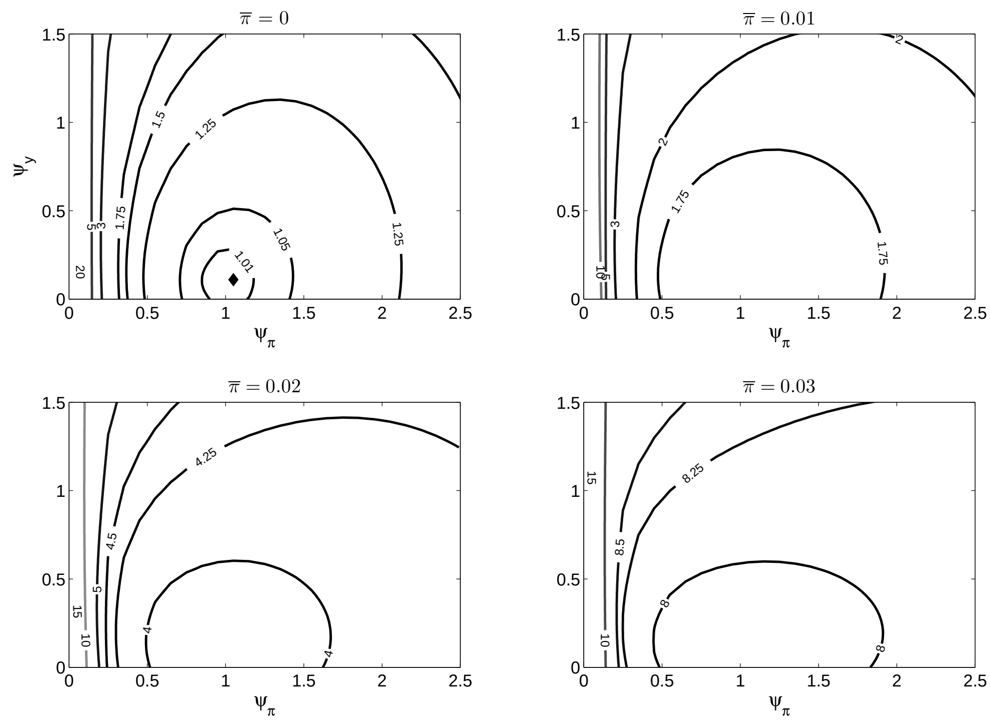

Figure 2: Iso-expected loss contours under full information. $\bar{\pi}$ represents target inflation, and ψ_π and ψ_y are feedback parameters on lagged inflation and output growth. The black diamond in the upper left panel marks the full-information optimum. Contour lines measure the gross increase in expected loss relative to the optimum.

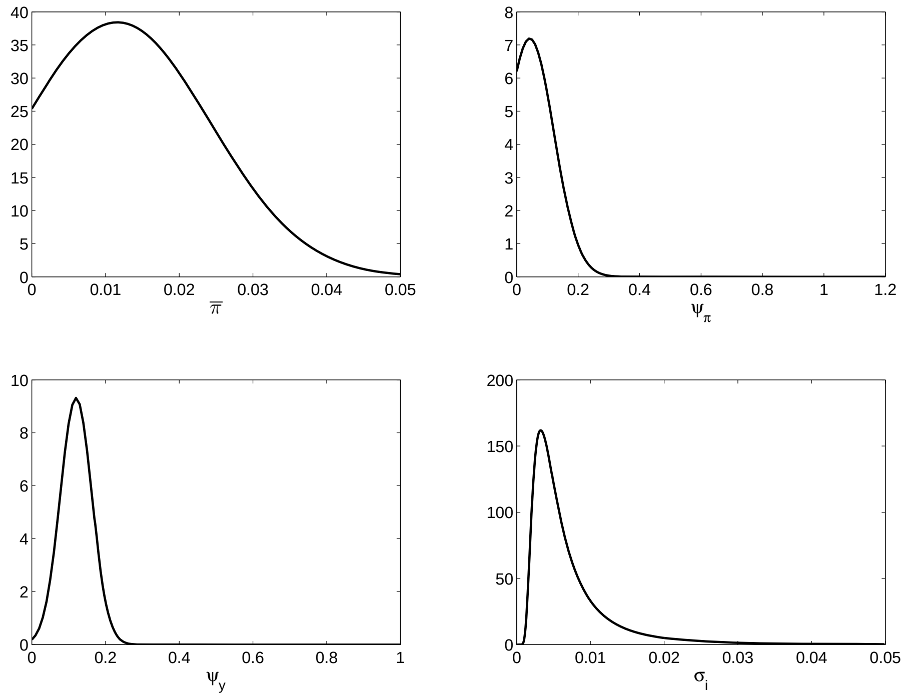

Figure 3: Prior probability distributions on the policy coefficients. $\bar{\pi}$ represents target inflation, ψ_π and ψ_y are feedback parameters on lagged inflation and output growth, and σ_i is the standard deviation of the policy shock. These distributions are calibrated to match aspects of the estimates in table 2.

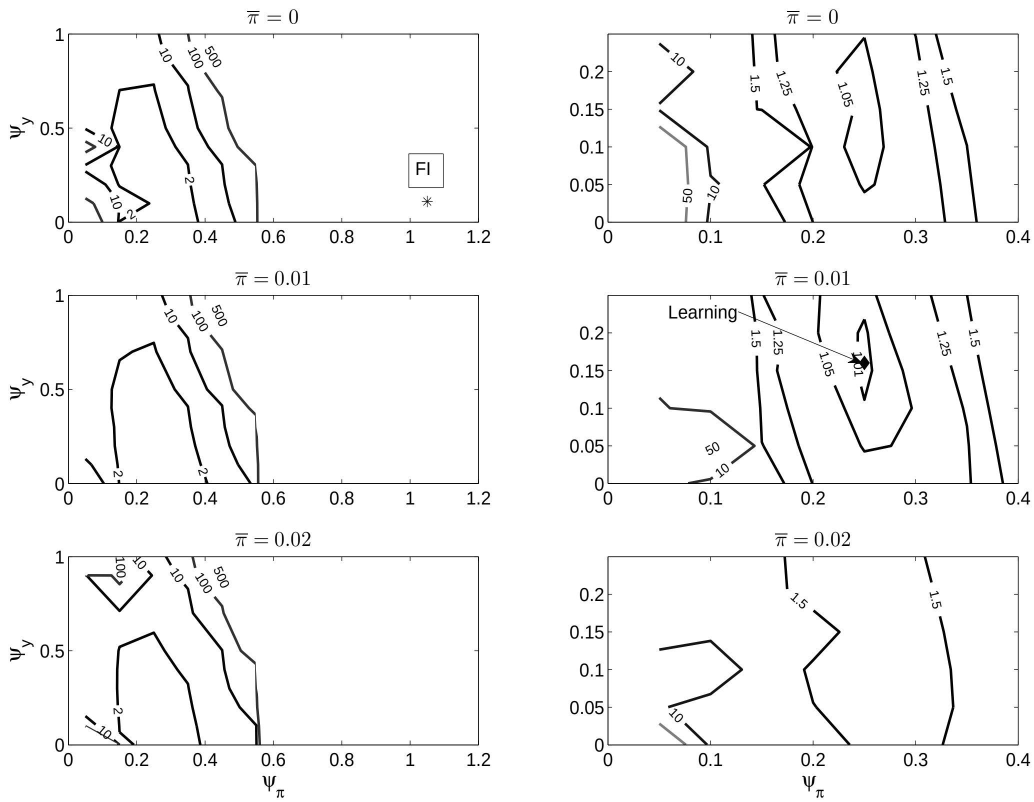

Figure 4: Iso-expected loss contours under learning. $\bar{\pi}$ represents target inflation, and ψ_π and ψ_y are feedback parameters on lagged inflation and output growth. The optimal coefficients under learning and full information are marked, respectively, by a black diamond and an asterisk. Contour lines measure the gross increase in expected loss relative to the optimum under learning.

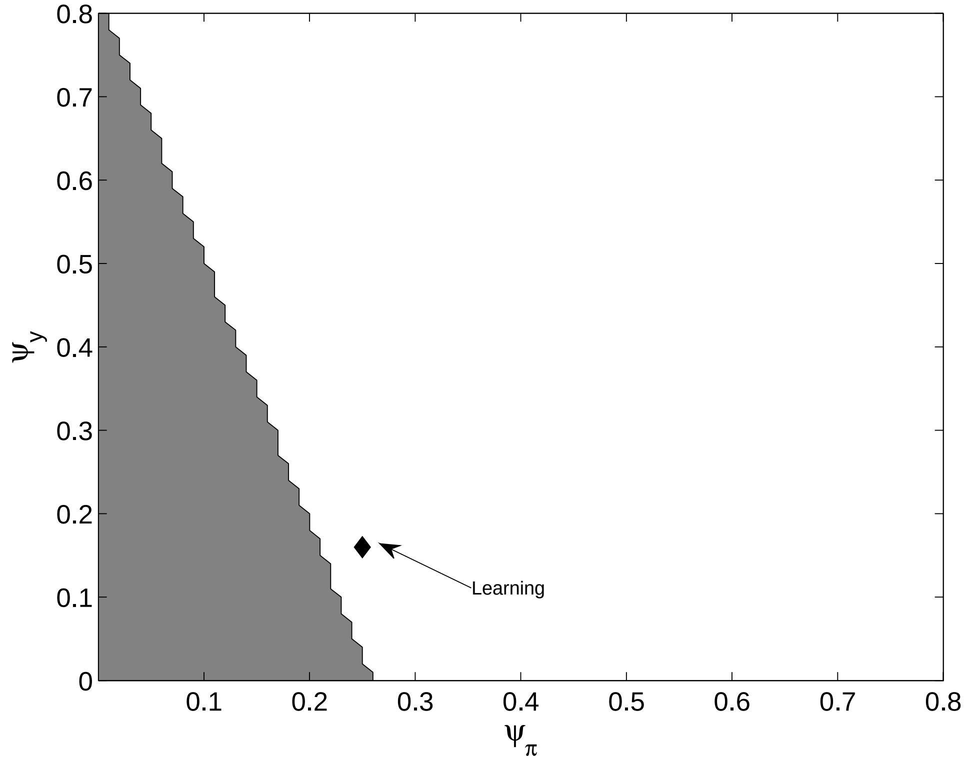

Figure 5: The gray region marks combinations of feedback pararmeters ψ_π and ψ_y on lagged inflation and output growth for which the ALM is nonexplosive at the beginning of the disinflation. Target inflation is set to zero. The black diamond marks the optimal coeffiecients under learning.

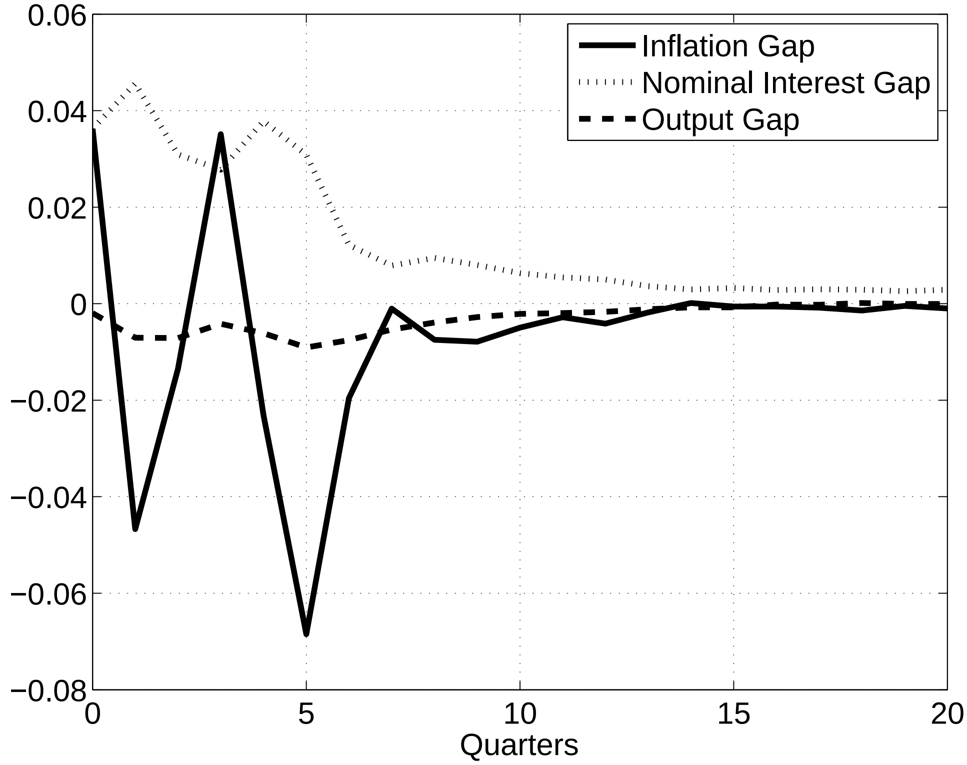

Figure 6: Responses of inflation, output, and interest rate gaps under learning. Values shown at date 0 depict differences between steady states in the old and new regimes. A Taylor rule optimized for learning is introduced at date 1.

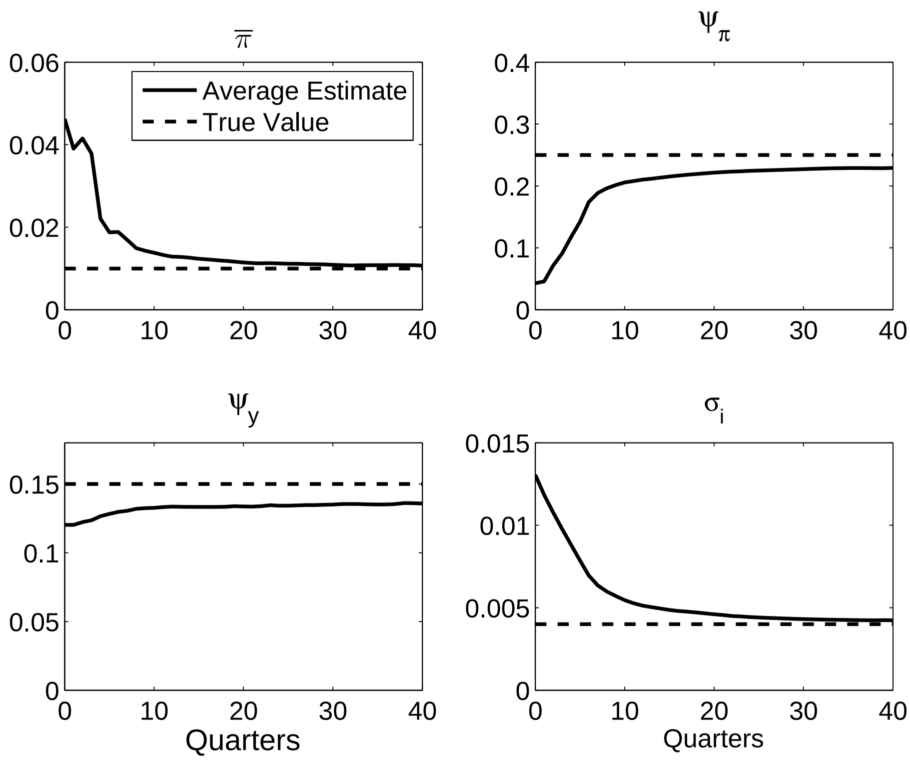

Figure 7: Average estimates of policy coefficients. $\bar{\pi}$ represents target inflation, ψ_π and ψ_y are feedback parameters on lagged inflation and output growth, and σ_i is the standard deviation of the policy shock. Values for date 0 are prior modes, while those shown at date 1 and after are cross-sample-path averages of posterior modes.

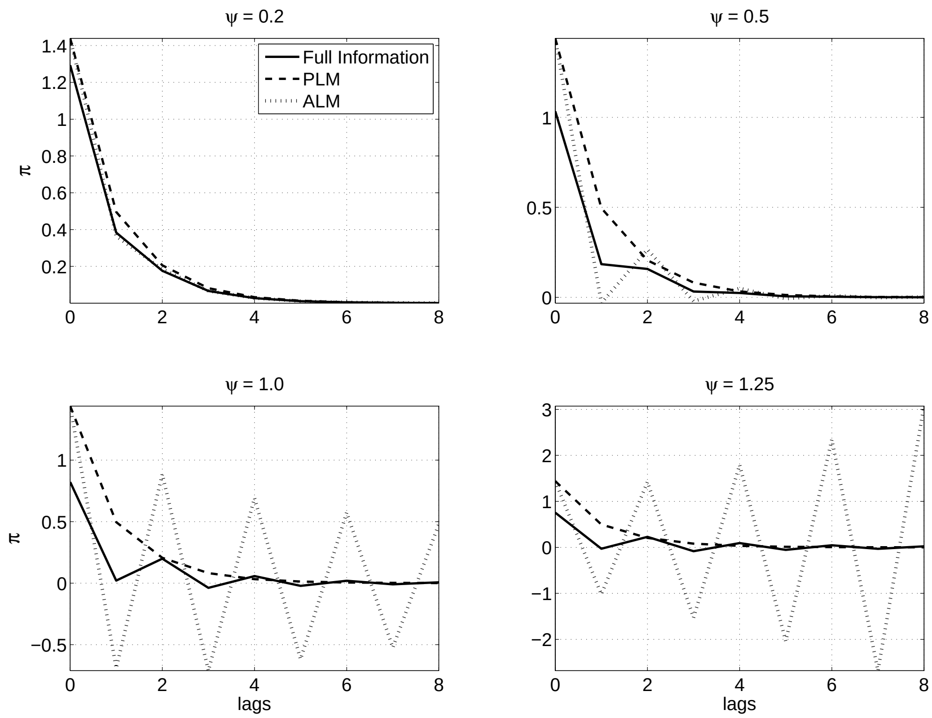

Figure 8: Response of inflation in the stripped-down example to a unit cost-push shock. ψ is the reaction coefficient on lagged inflation in the abstract policy rule (equation 24). The prior mode for ψ is $\psi_0 = 0.1$, the discount rate β is set to 0.99, and the autoregressive parameter for cost-push shocks is $\rho_u = 0.4$. The ALM and PLM are both based on initial beliefs.