

Cogley, Timothy; Matthes, Christian; Sbordon, Argia M.

Working Paper

Optimal disinflation under learning

Staff Report, No. 524

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Cogley, Timothy; Matthes, Christian; Sbordon, Argia M. (2011) : Optimal disinflation under learning, Staff Report, No. 524, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60829>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



STAFF REPORTS

www.newyorkfed.org/research

Optimal Disinflation under Learning

Timothy Cogley
Christian Matthes
Argia M. Sbordone

Number 524, November 2011

Federal Reserve Bank of New York
Staff Reports

Optimal Disinflation under Learning

Timothy Cogley
Christian Matthes
Argia M. Sbordone

Staff Report No. 524
November 2011

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 are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Optimal Disinflation under Learning

Timothy Cogley, Christian Matthes, and Argia M. Sbordone

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

November 2011

JEL classification: E31, E52

Abstract

We model transitional dynamics that emerge after the adoption of a new monetary policy rule. We assume that private agents learn about the new policy via Bayesian updating, and we study how learning affects the nature of the transition and the choice of a new rule. Temporarily explosive dynamics can emerge when there is substantial disagreement between actual and perceived policies. These dynamics make the transition highly volatile and dominate expected loss. The emergence of temporarily explosive paths depends more on uncertainty about policy-feedback parameters than about the long-run inflation target. For that reason, the central bank can at least achieve low average inflation. Its ability to move feedback parameters away from initial beliefs, however, is more constrained.

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

Cogley: New York University (e-mail: tim.cogley@nyu.edu). Matthes: Universitat Pompeu Fabra and Barcelona GSE (e-mail: christian.matthes@upf.edu). 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.

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.¹

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.² 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, viz. what policy is optimal when the private sector learns the new policy and how learning alter the central bank's choice. We study these questions 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 follows a simple Taylor-type rule and chooses its coefficients by minimizing a discounted quadratic

¹We also assume that the distribution from which preferences are drawn is unknown.

²See also Orphanides and Williams (2005) and Milani (2007).

loss function. The private sector learns the new policy via Bayesian updating, and the central bank takes learning into account when solving its decision problem.

When the private sector learns about the policy rule, 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 potential for explosive volatility.

For this reason, the bank's choice differs substantially from the full-information optimum. Uncertainty about the inflation target is not much of a problem. In our examples, the bank always achieves low average inflation, though sometimes it stops short of zero – the optimum under full information³ – 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.

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 (1998, 2000, 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's 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, there is a tight link between the actual and perceived laws of motion (ALM and PLM, respectively). In our model, agents know the ALM up to the unknown monetary policy parameters, and their PLM is the perceived ALM (i.e., the ALM evaluated at their current estimate of the policy coefficients). Because agents know the functional form of the ALM, they can use Bayes theorem to update beliefs. 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.

³We abstract from the zero lower bound on nominal interest.

2 A dynamic new-Keynesian model with positive target inflation

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

2.1 The timing protocol

Private agents enter period t with beliefs about policy coefficients inherited from $t - 1$. They treat estimated parameters as if they were known with certainty and formulate plans accordingly. 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 from that 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 estimates and carry them forward to $t + 1$.

2.2 The model

Our model is a dynamic new Keynesian model in which agents take expectations with respect to subjective predictive distributions. 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, a conventional intertemporal *IS* curve and an Ascari-Sbordone version of the aggregate supply curve. The model features habit persistence in consumption and staggered price setting. A log-linearized version is presented in this section. For details on how we arrived at this representation, see web appendix A.

2.2.1 Monetary policy

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

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

where i_t is the nominal interest rate, π_t is inflation, y_t is log output, and ε_{it} is an *i.i.d.* policy shock. The timing assumption follows McCallum (1999) and fits conveniently within the timing protocol described above.⁴ The policy coefficients are collected in

⁴McCallum (1999) contends that monetary policy rules should be specified in terms of lagged variables, on the grounds that the Fed lacks good current-quarter information about inflation, output, and other right-hand side variables. This is especially relevant for decisions taken early in the quarter.

a vector $\psi = [\bar{\pi}, \psi_\pi, \psi_y, \sigma_i]'$, where $\bar{\pi}$ represents the central bank's long-run inflation target, ψ_π and ψ_y are feedback parameters on the inflation gap and output growth, respectively, and σ_i is the standard deviation of the policy shock.

We adopt this form because 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 price-setting models with positive target inflation. Orphanides and Williams (2007) demonstrate that it performs well under least-squares learning.⁵ 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, Erceg and Levin (2003) contend that output growth, rather than the output gap, is the appropriate measure to include in an estimated policy reaction function for the U.S.

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

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

is a perceived policy shock that depends on the actual policy shock ε_{it} and on estimation of policy coefficients. Private agents believe that $\tilde{\varepsilon}_{it}$ is white noise, but in fact it involves undetected feedback onto lagged state variables.

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

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

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

taking private-sector learning into account.⁶ In addition to penalizing variation in inflation and the output gap, the loss function includes a small penalty for deviations of the nominal interest rate from its steady state. The central bank arbitrarily sets σ_i and optimizes with respect to $\bar{\pi}$, ψ_π , and ψ_y .

⁵Orphanides and Williams (2007) postulate that neither the agents nor the central bank know the true structure of the economy, and replace rational expectations with least-squares learning. They show that an optimized Taylor rule in difference form dominates an optimized standard Taylor rule when learning and time-varying natural rates interact.

⁶Gaspar et al (2006) distinguish between an unsophisticated central bank - one that accounts for the beliefs of the public but not the dynamic process of learning - and a sophisticated central bank that also takes the learning process into account. Our setting corresponds to the latter assumption.

2.2.2 Approximation methods

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

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

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

Our second approximation involves the assumption that agents treat the current estimate ψ_t as if it were known with certainty. Kreps (1998) calls this an 'anticipated-utility' model. In the context of a single-agent decision problem, Cogley and Sargent (2008) compare the resulting decision rules with exact Bayesian decision rules, and they demonstrate that the approximation is 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. This approximation is standard in the macro-learning literature.

2.2.3 A new-Keynesian IS curve

As usual, we assume that a representative household maximizes expected utility subject to a flow budget constraint. 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 is 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, we obtain a version of the new Keynesian *IS* curve,

$$y_t - \bar{y}_t = \xi_t - \xi - E_t^* [\xi_{t+1} - \xi - (y_{t+1} - \bar{y}_t) - (\gamma_{t+1} - \gamma) + i_t - \pi_{t+1} - r], \quad (4)$$

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

$$\xi_t - \xi \equiv \xi_1 (y_t - \bar{y}_t) + \xi_2 [y_{t-1} - \bar{y}_t - (\gamma_t - \gamma) + \beta E_t^* (y_{t+1} - \bar{y}_t + \gamma_{t+1} - \gamma)] + \varepsilon_{yt}. \quad (5)$$

The parameter β is a subjective discount factor, r is the steady-state real interest rate, γ is the steady-state growth rate for technological progress, 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} reflect technology and preference shocks, respectively. See web appendix A for further details.

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

A second difference concerns the expectation operator E_t^* , which represents forecasts formed with respect to the private sector's perceived law of motion. In contrast, the central bank takes expectations with respect to the actual law of motion, which we denote by E_t .⁷

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

$$\gamma_t = (1 - \rho_\gamma) \gamma + \rho_\gamma \gamma_{t-1} + \varepsilon_{\gamma t}, \quad (6)$$

and a white-noise shock ε_{yt} .

2.2.4 A new-Keynesian Phillips curve

We adopt a purely-forward looking version of Calvo's (1983) price-setting model. A continuum of monopolistically competitive firms produce a variety of differentiated intermediate goods that are sold to a final-goods producer. Firms that produce the intermediate goods reset prices at random intervals. In particular, with probability $1 - \alpha$ an intermediate-goods producer has an opportunity to reset its price, and with probability α its price remains the same. Thus we abstract from indexation 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, we again log-linearize around perceived steady states, obtaining the following

⁷We 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 $E_t x_{t+j} \neq E_t^* x_{t+j}$.

block of equations,

$$\begin{aligned}\pi_t - \bar{\pi}_t &= \beta_t E_t^*(\pi_{t+1} - \bar{\pi}_t) + \kappa_t(y_t - \bar{y}_t) + \varsigma_t(\delta_t - \bar{\delta}_t) - \tilde{\kappa}_t(\xi_t - \xi) \\ &\quad + \gamma_{1t} E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_t) + \phi_{t+1}] + u_t + \varepsilon_{\pi t},\end{aligned}\quad (7)$$

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

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

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

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

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

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

Third, the *NKPC* coefficients depend on deep parameters and estimates of target inflation $\bar{\pi}_t$,⁸

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

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.

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

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

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

⁸The *NKPC* parameters collapse to the usual expressions when $\bar{\pi}_t = 0$.

2.3 Calibration

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

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

We calibrate the preference parameters 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).

Parameters governing the shocks are calibrated as follows. We abstract from average growth, setting $\gamma = 0$. For the persistent shocks u_t and γ_t , we take estimates from Cogley, Sargent and Primiceri (2009),

$$\begin{aligned} \rho_u &= 0.4, & 100\sigma_u &= 0.12, \\ \rho_\gamma &= 0.27, & 100\sigma_\gamma &= 0.5. \end{aligned} \quad (14)$$

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

$$\sigma_\pi = \sigma_y = 0.01/4 \quad (15)$$

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

3 Learning about monetary policy

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, this could be done via the Kalman filter. Because ψ enters non-linearly, however, agents must solve a nonlinear filtering problem. This section describes how that is done. We first conjecture a perceived law of motion (PLM) and then derive the actual law of motion (ALM) under the PLM. After that, we verify that the PLM is the perceived ALM. Having

verified that private agents know the ALM up to the unknown policy 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 the perceived monetary-policy rule, the private sector's model of the economy can be represented as a system of linear expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_t S_{t-1} + D_t \tilde{\varepsilon}_t, \quad (16)$$

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 web appendix B 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 (16). The reduced form can be expressed as

$$S_t = F_t S_{t-1} + G_t \tilde{\varepsilon}_t, \quad (17)$$

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, (17) serves two functions, describing how agents forecast future outcomes and how they make current-quarter decisions.

Because of the anticipated-utility assumption, equation (16) has the same mathematical form as a conventional rational-expectations model. It follows that the conditions for a unique nonexplosive solution are the same. When those conditions are violated, many nonexplosive solutions exist. We use Sims' (2001) Gensys program to solve for (F_t, G_t) . His program delivers the unique nonexplosive solution when outcomes are determinate and one of the many possible nonexplosive solutions when they are indeterminate.

3.2 The actual law of motion

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

$$A_t S_t = B_t E_t^* S_{t+1} + C_{at} S_{t-1} + D_t \varepsilon_t. \quad (18)$$

The state vector is the same as in (16), as are the matrices A_t, B_t , and D_t . In addition, all rows of C_{at} agree with those of C_t except for the one corresponding to the monetary-policy rule. In that row, the true policy coefficients ψ replace the estimated coefficients

ψ_t (see web appendix B). To find the ALM, substitute $E_t^* S_{t+1} = F_t S_t$ from the PLM into (18) and re-arrange terms. After some algebra, we find⁹

$$S_t = H_t S_{t-1} + G_t \varepsilon_t, \quad (19)$$

where

$$H_t = F_t + (A_t - B_t F_t)^{-1} (C_{at} - C_t). \quad (20)$$

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.¹⁰

When the solution for (F_t, G_t) is unique, so is the solution for H_t . When there are multiple solutions for (F_t, G_t) , there are also multiple solutions for H_t . Our program delivers one of the many possible solutions in that case, corresponding to the solution that Gensys selects for (F_t, G_t) .

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 are¹¹

$$\begin{aligned} m_{t|t-1}(\psi_{true}) &= H_t(\psi_{true}) S_{t-1}, \\ V_{t|t-1}(\psi_{true}) &= G_t V_\varepsilon(\psi_{true}) G_t', \end{aligned} \quad (21)$$

⁹The ALM can also be derived as follows. Outcomes are determined in accordance with agents' plans,

$$S_t = F_t S_{t-1} + (A_t - B_t F_t)^{-1} D_t \tilde{\varepsilon}_t.$$

A relation between perceived and actual innovations can be found by subtracting (18) from (16),

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

Substitute this relation into agents' plans to express outcomes in terms of actual shocks,

$$S_t = [F_t + (A_t - B_t F_t)^{-1} (C_{at} - C_t)] S_{t-1} + G_t \varepsilon_t.$$

¹⁰Notice that the ALM is a VAR with time-varying parameters and conditional heteroskedasticity, as in Cogley and Sargent (2005) and Primiceri (2006). An intriguing feature of the equilibrium is that the drifting parameters ψ_t have a lower dimension than the conditional mean parameters $vec(H_t)$. This is qualitatively consistent with a finding of Cogley and Sargent (2005), who reported that drift in an analog to $vec(H_t)$ is confined to a lower dimensional subspace. The form of the conditional variance in (19) differs from their representations, however, so the model disagrees with their identifying restrictions. Another difference is that the model involves temporary drift during a learning transition while their VARs involve perpetual drift.

¹¹According 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. That is why it is legitimate to use them to calculate the conditional mean and variance.

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 we interviewed the agents in the model 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

$$\begin{aligned}\tilde{m}_{t|t-1}(\psi_t) &= F_t S_{t-1}, \\ \tilde{V}_{t|t-1}(\psi_t) &= G_t V_\varepsilon(\psi_t) G_t'.\end{aligned}\tag{22}$$

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. 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 $t + j + 1$ hold in expectation at t .

3.4 The likelihood function

We collect the observables in a vector $X_t = [\pi_t, u_t, y_t, \gamma_t, i_t]'$, which is a subset of the state vector S_t ,

$$X_t = e_X S_t,\tag{23}$$

where e_X is an appropriately defined selection matrix (see web appendix B). The other elements of the state vector allow us to express the model in first-order form but convey no additional information beyond that contained in the history of X_t . Using the prediction-error decomposition, the likelihood function for data through period t can be expressed as

$$p(X^t|\psi) = \prod_{j=1}^t p(X_j|X^{j-1}, \psi).\tag{24}$$

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

$$\begin{aligned}m_{t|t-1}^X(\psi) &= e_X H_t(\psi) S_{t-1}, \\ V_{t|t-1}^X(\psi) &= e_X G_t V_\varepsilon(\psi) G_t' e_X',\end{aligned}\tag{25}$$

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

$$\ln p(X^t|\psi) = -\frac{1}{2} \sum_{j=1}^t \left\{ \ln |V_{j|j-1}^X(\psi)| + [X_j - m_{j|j-1}^X(\psi)]' (V_{j|j-1}^X(\psi))^{-1} [X_j - m_{j|j-1}^X(\psi)] \right\}.\tag{26}$$

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 log likelihood and log prior. Because of our anticipated-utility assumption, their decisions depend only on a point estimate, not on the entire posterior distribution. Among the various point estimators from which they can choose, we assume they adopt the posterior mode,

$$\psi_t = \arg \max (\ln p(X^t|\psi) + \ln p(\psi)) . \quad (27)$$

Agents take into account that past outcomes were influenced by past beliefs. They are not recursively estimating a conventional rational-expectations model. By inspecting the ALM and PLM, one can verify that past values of the conditional mean $m_{j|j-1}^X(\psi)$ and the conditional variance $V_{j|j-1}^X(\psi)$ depend on past estimates as well as the current candidate ψ . Past estimates are bygones at t and are held constant when agents update the posterior mode.¹²

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 is examined below.

4 The central bank's decision problem

A new governor appears at date 0 and formulates a new policy rule. After observing the private sector's prior, the 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. The disinflation commences at date 1.

4.1 Initial conditions

The economy is initialized at the steady state under the old regime. Because we are interested in a scenario like the end of the Great Inflation, we calibrate the old regime to match estimates of the policy rule for the period 1965-1979. We assume that the policy rule had the same functional form as (1) during that period, and we estimate $\bar{\pi}, \psi_\pi, \psi_y$, and σ_i^2 by OLS. The point estimates and standard errors are reported in table 1. The point estimate for $\bar{\pi}$ is 0.0116, implying an annualized target-inflation rate of 4.6 percent. The reaction coefficients are both close to zero, with the output coefficient being slightly larger than the inflation coefficient. Policy shocks are large in magnitude and account for a substantial fraction of the total variation in the nominal interest rate. Standard errors are large, especially for $\bar{\pi}$.

¹²For that reason, whether the estimator can be expressed in a recursive form is unclear.

Table 1: The Old Regime

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

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

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

4.2 Evaluating expected loss and finding the optimal policy

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

We proceed numerically. We start by specifying a grid of values for $\bar{\pi}$, ψ_{π} , and ψ_y . Then, for each node on the grid, we simulate 100 sample paths, updating private-sector estimates ψ_t by numerical maximization at each date. The sample paths are each 20 years long, and the terminal loss is set equal to zero, representing a decision maker with a long but finite horizon. We calculate realized loss along each sample path and then average realized loss across sample paths to find expected loss. The optimal rule among this family is the node with smallest expected loss.¹³

5 A full-information benchmark

To highlight the role of learning, we begin by describing the optimal policy under full information. When private agents know the new policy, the optimal simple rule sets $\bar{\pi} = 0$, $\psi_{\pi} = 1.45$, and $\psi_y = 0.2$. Figure 1 portrays a disinflation under this policy. Recall that the economy is initialized in the steady state of the old regime and that the disinflation commences at date 1. The figure depicts responses of inflation, output, and nominal interest gaps, which are defined as deviations from the steady state of the new regime.¹⁴

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

¹³An alternative procedure would be to write down a dynamic program and solve it numerically, as in Gaspar, Smets, and Vestin (2006, 2010). This is feasible in models with a low-dimensional state vector, but in our model it runs afoul of the curse of dimensionality.

¹⁴Values at date 0 represent the difference between steady states of the old and new regimes. Inflation and nominal interest gaps coincide because the steady-state real interest rate is the same.




Figure 1: Disinflation under full information

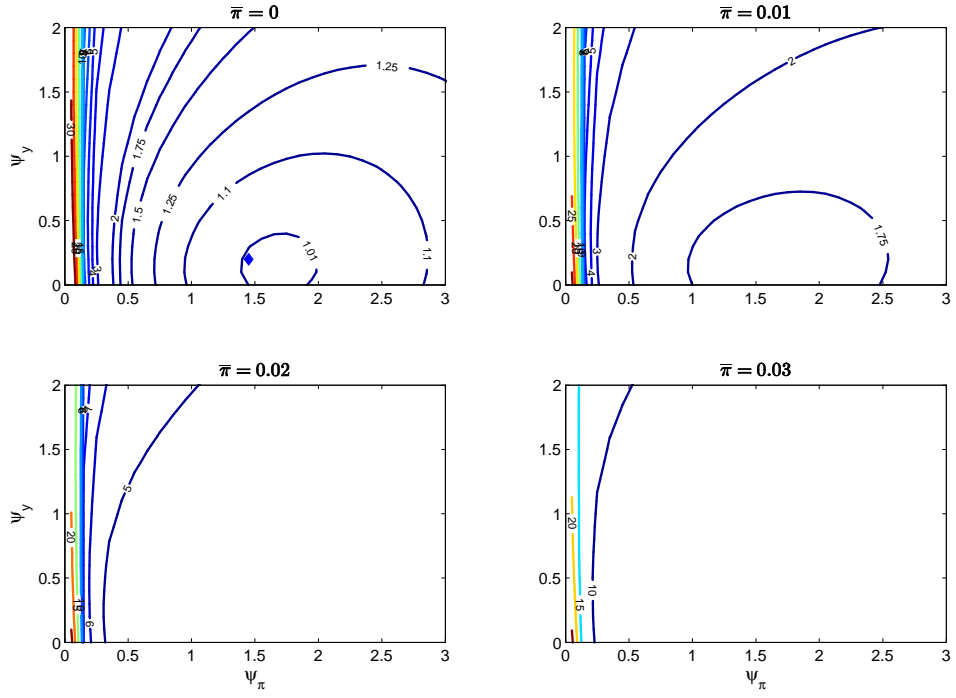


Figure 2: Iso-loss contours under full information

the price level, partially counteracting the effects of high past inflation. As Woodford (2003) explains, a partial rollback of the price level is a feature of optimal monetary policy under commitment. Intuitively, a credible commitment on the part of the central bank to roll back price increases restrains a firm's incentive to increase its price in the first place. The optimal simple rule under full information also has this feature.

The initial increase in the nominal interest rate causes the output gap to fall below zero. Since inflation and output growth are below target at date 1, the central bank cuts the interest rate at date 2, damping the output loss and initiating a recovery. Convergence to the new steady state is rapid, with inflation, output, and interest gaps closing 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 2.4 percent. The sacrifice ratio, defined as the cumulative loss in output divided by the change in target inflation, is approximately 0.5 percent. The reason why the sacrifice ratio is small under full information is that the model has no indexation. Although prices are sticky, the absence of indexation means that inflation is weakly persistent. The absence of indexation also explains why the bank seeks a substantial rollback in the price level.

Another property of the full-information model is that the economy is highly fault tolerant with respect to policies away from the optimum. Figure 2 portrays isoclines for expected loss as a function of $\bar{\pi}$, ψ_π , ψ_y , and σ_i^2 . Each panel involves a different setting for $\bar{\pi}$, ranging from 0 to 3 percent per annum. The feedback parameters ψ_π and ψ_y are shown on the horizontal and vertical axes, respectively. The standard deviation of the policy shock σ_i is held constant at 0.001 in all cases. Expected loss is normalized by dividing by loss under the optimal rule, so that contour lines represent gross deviations from the optimum. The blue diamond in the upper left panel depicts the optimal policy under full information. 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. Later we contrast this with an absence of fault tolerance under learning.

6 An optimal simple rule under learning

Our benchmark model assumes that the private sector initially anticipates a continuation of the old regime. We calibrate their priors using the estimates of policy coefficients for 1965-1979 shown in table 1, thus ensuring that the prior encodes information from the period leading up to the Volcker disinflation. In particular, we assume that agents believe that policy coefficients are independent a priori,

$$p(\psi) = p(\bar{\pi})p(\psi_\pi)p(\psi_y)p(\sigma_i^2). \quad (28)$$




Figure 3: A prior based on the old regime

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

The results are shown in figure 3. Priors for ψ_{π} and ψ_y concentrate slightly to the right of zero, and little prior mass is assigned to values greater than 0.25. On the other hand, priors for $\bar{\pi}$ and σ_i^2 are spread out and assign non-negligible probability to a broad range of values.¹⁵ According to this specification, private agents are skeptical that the central bank will react aggressively to inflation or output, but they are open to persuasion about $\bar{\pi}$ and σ_i^2 .

Figure 4 portrays isoclines for expected loss as a function of $\bar{\pi}$, ψ_{π} , ψ_y , and σ_i^2 . The figure is formatted in the same way as figure 2, but we focus on a narrower range of feedback parameters. Expected loss is normalized by dividing by loss under the rule

¹⁵ $\bar{\pi}$ and σ_i are measured in quarterly rates.

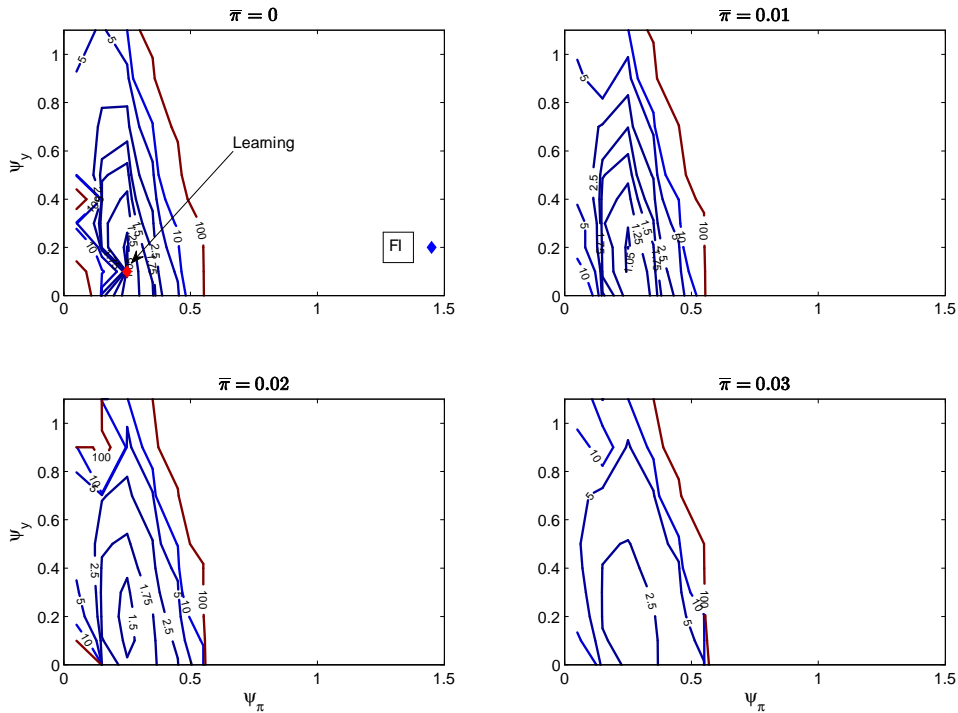


Figure 4: Isoloss contours under learning

optimal under learning. The red and blue diamonds in the upper left panel depict the optimal policy under learning and full information, respectively.

Regions of low expected loss concentrate in the southwest quadrant of each panel, near the prior mode for ψ_π and ψ_y . Expected loss increases rapidly as the feedback coefficients move away. Indeed, in the northeast quadrant of each panel, expected loss is more than 100 times greater than under the optimal policy. The policy that is optimal under full information lies in the high-loss region and performs very badly under learning.

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,¹⁶ 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

¹⁶A unit eigenvalue is associated with the constant in the state vector.

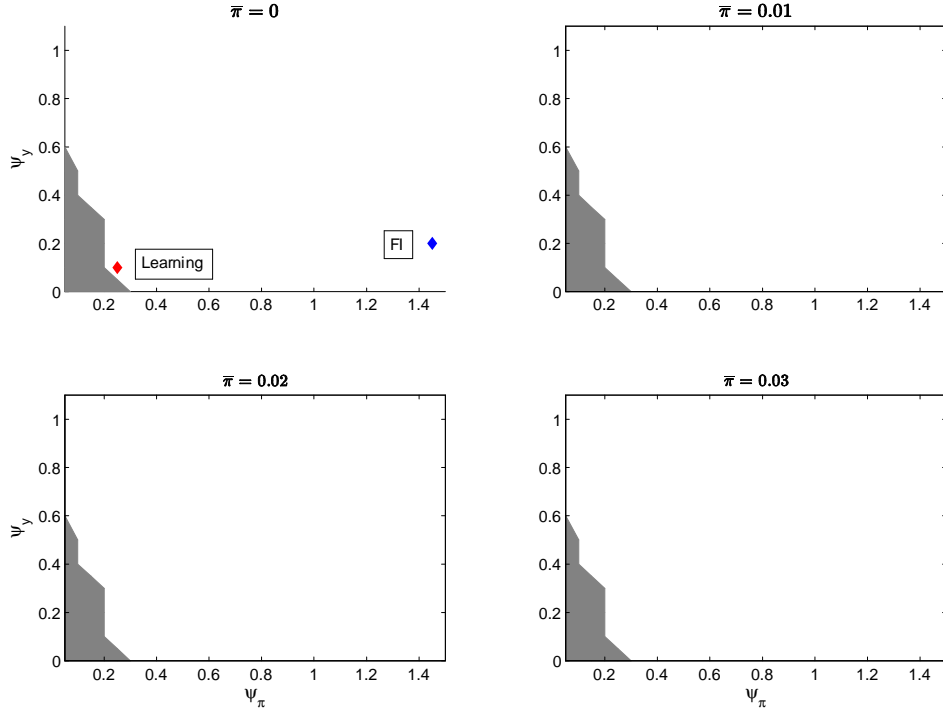


Figure 5: Nonexplosive region for H_1

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 20). 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,¹⁷ but the transition is highly volatile and dominates expected loss when the initial disagreement is large and/or learning is slow.

The gray shaded areas in figure 5 depict regions of the policy-coefficient space for which the eigenvalues of H_1 are nonexplosive.¹⁸ The nonexplosive region is similar

¹⁷Since the private sector uses Bayesian inference and an anticipated utility approach to decision making, standard results for the convergence of estimates formed by Bayesian decision makers (see, among others, El-Gamal and Sundaram (1993)) are not directly applicable. Therefore, we numerically check convergence of the agents' learning algorithm. In particular, we calculated deviations of their parameter estimates from the true values after 40 and 80 periods, both across simulations and across true parameter values. Histograms for those deviations are indeed centered near 0, and the variances of those distributions shrink as the learning horizon grows larger.

¹⁸The jagged boundary reflects the coarseness of our grid. The large losses near the origin in

for all settings of $\bar{\pi}$, but it is sensitive to ψ_π and ψ_y , concentrating near the prior mode. It follows that the emergence of explosive roots depends more on the feedback parameters than on the long-run inflation target. The central bank can move its inflation target 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 can make the transition turbulent.

In this example, the private sector is prejudiced against large values of ψ_π and ψ_y . If the bank were to reach far outside the nonexplosive region in figure 5, it would have to fight that prejudice, and learning would be slow. Explosive eigenvalues would remain active for too long, making the transition too turbulent. For that reason, the optimal policy puts ψ_π and ψ_y only slightly outside. The bank can adjust $\bar{\pi}$ more freely, however, thereby achieving low average inflation.

The optimal simple rule for this example sets $\bar{\pi} = 0$, $\psi_\pi = 0.25$, and $\psi_y = 0.1$. Figures 6 and 7 portray outcomes under this policy. Figure 6 plots mean responses of inflation, output, and nominal interest gaps, averaged across 100 sample paths. Figure 7 portrays mean estimates of the policy coefficients, again averaged across 100 sample paths. The true coefficients are shown as dashed red lines while average estimates are portrayed as solid blue lines.

As shown in figure 6, 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 the response under learning is greater in amplitude, and inflation oscillates as it converges to its new long-run target. The transition now takes about three years, with inflation remaining below target for most of that time. There is also a shallow but long-lasting decline in output. The output gap reaches a trough of -1.1 percent in quarter 5 and remains negative for four years. The cumulative output gap during this time is -10 percent. Since inflation falls permanently by 4.6 percentage points, the sacrifice ratio amounts to 2.15 percent of lost output per percentage point of inflation, four 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.

As shown in figure 7, estimates of ψ_π converge to the true value after about two years. Rapid convergence of ψ_π is crucial for eliminating locally-explosive dynamics. In this case, beliefs about ψ_π converge quickly because they don't have far to go. The bank sets ψ_π close to its prior mode precisely so that disagreement does not persist. Estimates of ψ_y also converge quickly, again because the optimum is close to the prior. Estimates of $\bar{\pi}$ converge more slowly, approaching the true value after 20 quarters. Learning about σ_i is also slow, but it is also less critical because of our

figure 4 occur in spite of H_1 being non-explosive because for those parameters instability can occur slightly later in the simulations (but within the first 5 periods).




Figure 6: Average responses under the optimal policy




Figure 7: Average estimates of policy coefficients under the optimal rule




Figure 8: Average responses when $\bar{\pi} = 0, \psi_{\pi} = 0.45, \psi_y = 0.1$




Figure 9: Average estimates when $\bar{\pi} = 0, \psi_{\pi} = 0.45, \psi_y = 0.1$

certainty-equivalent approximations.

To illustrate why a more ambitious reform is suboptimal, we examine an alternative policy that holds $\bar{\pi}$ and ψ_y constant but which reacts more aggressively to inflation, increasing ψ_π from 0.25 to 0.45. This policy is located to the right of the optimum in figures 4 and 5. Figures 8 and 9 depict average outcomes under this rule.

Under this policy, the central bank is fighting against the private sector's prior, which assigns low probability to neighborhoods of the true values. A lot of sample information is needed to overcome the prior. For the sake of intuition, imagine that agents were estimating the policy rule by running a regression. Because the prior assigns low weight to neighborhoods of the true value for ψ_π , the likelihood function would have to concentrate sharply in order to move the posterior there. For that to happen quickly, lagged inflation (the right-hand variable in the regression) would have to be highly volatile. The bank can create a lot of volatility (see figure 8), and those fluctuations do help the private sector learn (see figure 9). But that volatility is costly, and the long-run benefits associated with a more aggressive reaction to inflation do not justify the higher transitional costs.

6.1 Uncertainty about target inflation when feedback parameters are known

Our results suggest that uncertainty about feedback parameters is more important than uncertainty about target inflation.¹⁹ To show this more explicitly, we deactivate uncertainty about ψ_π , ψ_y , and σ_i and study a model in which $\bar{\pi}$ is the only unknown policy coefficient.²⁰ For $\bar{\pi}$, we assume the private sector adopts the same prior as in figure 3. The results are depicted in figures 10-13. As before, the gray area in figure 10 portrays the region for which the ALM is initially nonexplosive. Because the feedback parameters are known, this region expands to fill most of the policy-coefficient space. Since the ALM is nonexplosive for most policies, the model has high fault tolerance with respect to rules far from the prior, and the expected-loss surface is flatter. Furthermore, the private sector learns target inflation 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 shown in figure 1.

¹⁹Just to be clear, we are not saying that uncertainty about $\bar{\pi}$ is unimportant, only that it matters less than uncertainty about ψ_π and ψ_y .

²⁰This scenario is analogous to that of Erceg and Levin (2003). Our model differs from theirs in a number of other ways, so this exercise should not be interpreted as an attempt to replicate their analysis.




Figure 10: Nonexplosive region for H_1 when $\bar{\pi}$ is the only unknown policy coefficient

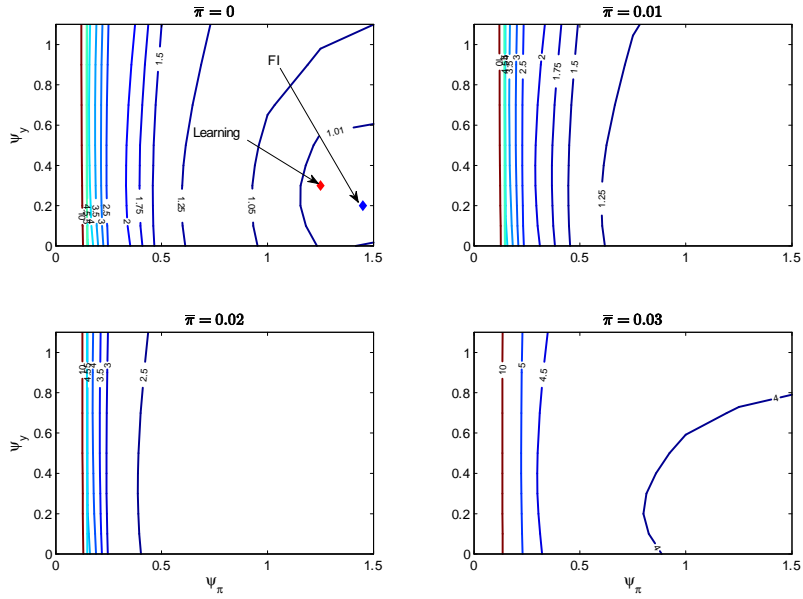


Figure 11: Iso-loss contours when $\bar{\pi}$ is the only unknown policy coefficient




Figure 12: Average responses under the optimal policy when $\bar{\pi}$ is the only unknown policy coefficient




Figure 13: Average estimates under the optimal policy when $\bar{\pi}$ is the only unknown policy coefficient

6.2 Uncertainty about feedback parameters when target inflation is known

Next we deactivate uncertainty about $\bar{\pi}$ and reactivate uncertainty about ψ_π, ψ_y , and σ_i . We assume that the private sector adopts the same priors for the latter coefficients as in figure 3. Results are presented in figures 14-17. At least qualitatively, the outcomes are closer to those for the benchmark learning model than to those under full information. Temporarily explosive paths still emerge when ψ_π and/or ψ_y deviate too much from prior beliefs. In fact, the initial nonexplosive region is larger than in the benchmark learning model. Because of concerns about explosive volatility, the bank chooses a policy close to the prior mode for ψ_π and ψ_y . The transition is volatile, but learning is rapid because the true policy is close to initial beliefs.

7 Alternative models of learning

In this section, we examine the extent to which policy recommendations derived from the baseline model are robust to alternative forms of learning. For the most part, we find that they are. This is good news for a central bank because it means that its choice of policy does not depend sensitively on detailed knowledge about how agents learn.

7.1 Single-equation learning

Agents in the baseline model are highly sophisticated and exploit cross-equation restrictions on the ALM when estimating policy coefficients. Here we step back and consider a less sophisticated form of learning involving single-equation estimation of the policy rule. The priors are the same as in the benchmark model, but we now assume that agents neglect cross-equation restrictions and work with the conditional likelihood function for the policy equation,

$$\ln p(\Delta i^t | \psi, \pi^t, \Delta y^t) = -\frac{1}{2} \sum_{j=1}^t \varphi^j \left\{ \ln \sigma_i^2 + \frac{(\Delta i_j - \psi_\pi(\pi_{j-1} - \bar{\pi}) - \psi_y \Delta y_{j-1})^2}{\sigma_i^2} \right\}. \quad (29)$$

The parameter φ discounts past observations. We consider two forms of single-equation learning, with $\varphi = 1$ and $\varphi < 1$, respectively, to imitate decreasing- and constant-gain learning.²¹ For the discounted case, we set $\varphi = 0.9828$ so that the discount function has a half-life of 40 quarters. We also multiply the log prior by φ^t on the grounds that date-zero beliefs should also be discounted when agents are concerned about structural change.

²¹For background, see Evans and Honkapohja (2001).

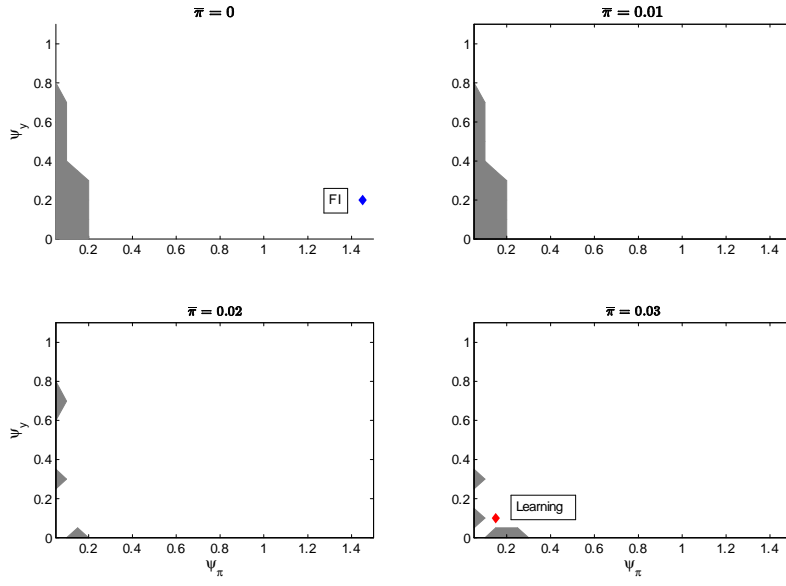


Figure 14: Nonexplosive region for H_1 when $\bar{\pi}$ is known

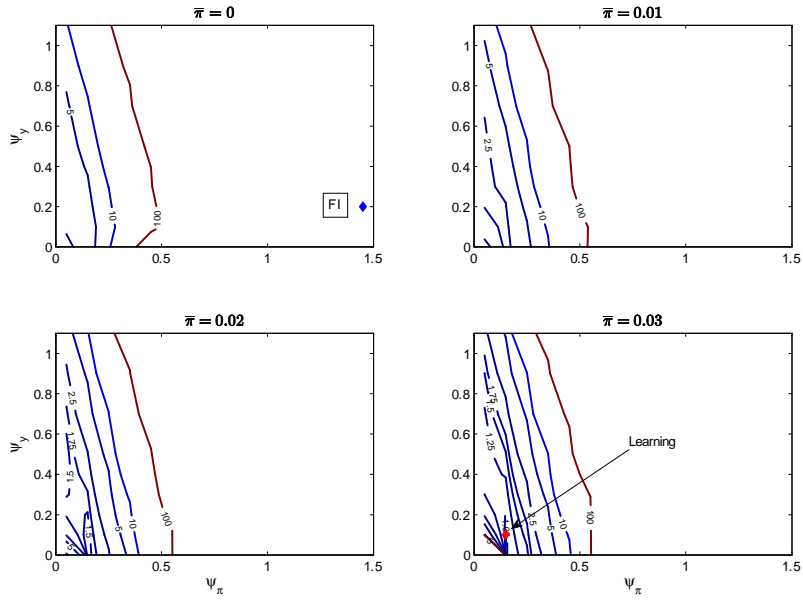


Figure 15: Isoloss contours when $\bar{\pi}$ is known




Figure 16: Average responses under the optimal policy when $\bar{\pi}$ is known

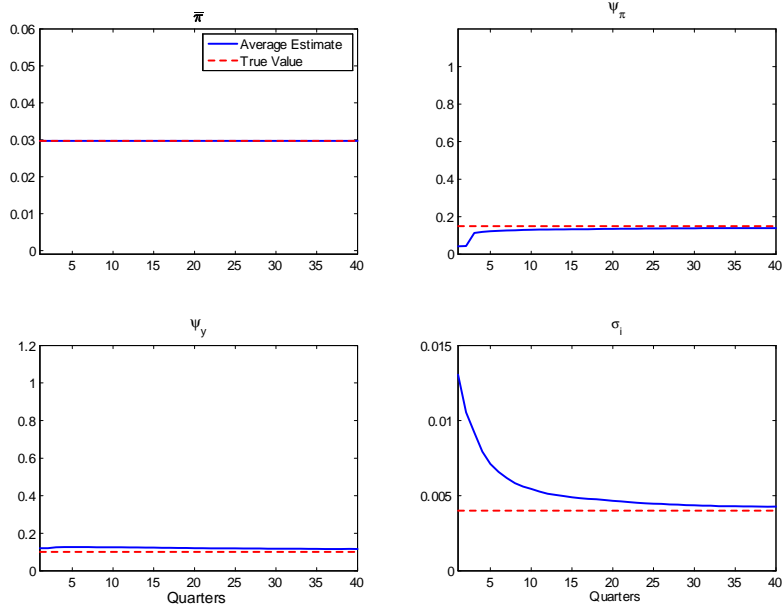


Figure 17: Average estimates under the optimal policy when $\bar{\pi}$ is known

Although policy-coefficient estimates differ, the optimal policies are essentially the same as those in the baseline learning model (see web appendix C). Hence the choice of policy does not depend on whether single-equation or full-system estimators are used.

That results are similar for discounted and undiscounted learning is not surprising because the samples are short and φ is not far from 1 in the discounted case. That the results are similar to those for full-system learning is a statement about the information content of cross-equation restrictions. Evidently those restrictions are less informative under learning 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 information about them. In a learning model, however, private decision rules are predicated on *estimates* of policy coefficients, not on true values. Hence non-policy equations in the ALM encode less information about the true policy. The cross-equation restrictions are informative, especially for policies far from the optimum, but the two estimators are similar for policies near the optimum.

7.2 Alternative priors

The baseline model also assumes that the central bank knows the private sector's prior distribution. While surveys might uncover their prior mode, eliciting information about the tails seems more demanding. Here we examine the robustness of the baseline policy with respect to changes in the private sector's prior. The bank still proceeds as if it knew the private sector's prior, but we now entertain the possibility that the bank's assumptions about the prior are mistaken.

We alter the baseline prior by creating a family of mixture priors. We imagine that agents enter date 0 with beliefs about the old regime $p_{old}(\psi)$ that are the same as figure 3. Instead of assigning probability 1 to that prior, however, we assume they expect the central bank to continue the old regime with probability $1 - w$ and to switch to something else with probability w . Their beliefs about a new regime are encoded in a conditional prior $p_{new}(\psi)$. The marginal prior is a mixture of the two conditional priors,

$$p_m(\psi) = (1 - w)p_{old}(\psi) + wp_{new}(\psi), \quad (30)$$

where w measures the public's beliefs about the prospects for change. A value close to 0 means that the public weighs past experience heavily and is skeptical about the prospects for change, while a value close to 1 means that the private sector discounts the past and looks forward to something new.

For $p_{new}(\psi)$, we adopt the same functional forms as for $p_{old}(\psi)$, and we calibrate it so that it is loosely centered on policies with lower long-run inflation and more aggressive reactions to inflation and output growth. The details are recorded in table 2. Abstracting from the truncation at zero, the conditional prior mean for $\bar{\pi}$ is 2 percent per year, and a conditional 95 percent confidence band ranges from 0 to 4

percent. Similarly, the conditional prior means for ψ_π and ψ_y are 0.5 with conditional confidence bands of plus or minus 0.5, and the conditional mode for σ_i is 0.001 with a standard deviation is 0.001.

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

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

Figure 18 depicts a family of mixture priors. The components $p_{old}(\psi)$ and $p_{new}(\psi)$ are illustrated by solid blue and dashed green lines, respectively. Mixtures are shown as red lines, for weights of $w = 0.3, 0.5$, and 0.7 , respectively. The mixtures differ from $p_{old}(\psi)$ in two respects. The prior mode for $\bar{\pi}$ is shifted to the left, near the mode for $p_{new}(\bar{\pi})$, and the upper tails for ψ_π and ψ_y are fatter. Agents remain skeptical that the bank will react aggressively to inflation or output growth, but they are less strongly prejudiced against that possibility.

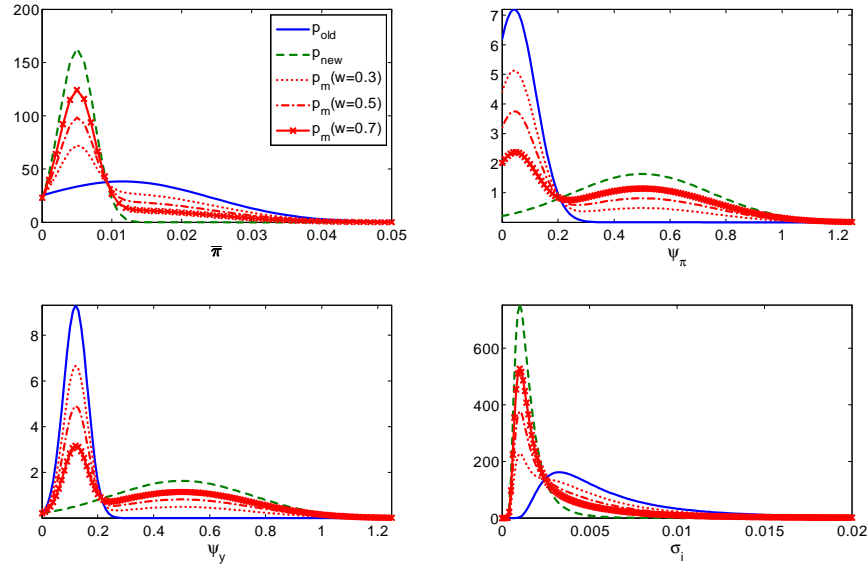


Figure 18: Mixture priors

Tables 3 and 4 report the optimal simple rule and relative expected loss for each model.²² Table 3 shows that the policy recommendations are broadly similar across models. All recommend low average inflation, though not always zero, and feedback

²²Detailed results can be found in web appendix D.

parameters never move far from the private sector’s prior mode. The bank’s main objective is to mitigate turbulence arising from explosive volatility, and all the policies accomplish that.

Table 4 records the results of a McCallum-style robustness analysis, reporting expected loss in the row model under the policy optimized for the column model. Expected loss is normalized by dividing by loss under the rule optimal for the row model. By and large, the policies are robust. The policy optimized for the baseline model ($w = 0$) works fairly well in all the models (see the first column). Policies optimized for priors with $w = 0.5$ or $w = 0.7$ increase expected loss in other models, sometimes by as much as a factor of 7 or 8, but no disasters occur.

Table 3: Optimal Simple Rules

	$\bar{\pi}$	ψ_{π}	ψ_y	Expected Loss
$w = 0$	0	0.25	0.1	1.99e-03
$w = 0.3$	0	0.25	0	1.94e-03
$w = 0.5$	0.01	0.15	0.2	9.27e-03
$w = 0.7$	0.01	0.15	0.3	7.73e-03

Table 4: Relative Expected Loss

	$w = 0$	$w = 0.3$	$w = 0.5$	$w = 0.7$
$w = 0$	1	1.10	1.41	1.28
$w = 0.3$	1.03	1	7.53	8.56
$w = 0.5$	3.14	2.66	1	1.05
$w = 0.7$	1.26	1.4	1.00	1

Relative expected loss in the row model under the policy optimized for the column model

Finally, we also examine models combining mixture priors and single-equation learning. The results are essentially the same as those reported here and in web appendix D.

8 Concluding remarks

Locally-unstable dynamics can emerge in equilibrium when private agents learn a new policy rule. These dynamics make the transition highly volatile and dominate expected loss. The central bank’s main challenge is to guard against this outcome. For the model developed here, uncertainty about policy-feedback parameters is critical because this is what creates the potential for temporarily explosive dynamics. 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.

In ongoing research, we study the properties of other policy rules, including monetary-aggregate rules as well as interest-rate rules with other functional forms. Perhaps we will discover a form of monetary-policy rule for which temporarily-explosive dynamics are less problematic.

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. 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262–302.
- [7] Cogley, T. and T.J. Sargent, 2008. Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review* 49, 185-221.
- [8] 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.
- [9] Coibion, O. and Y. Gorodnichenko. 2011. Monetary policy, trend inflation and the great moderation: an alternative interpretation. *American Economic Review* 101, 341-370.
- [10] Erceg, C. and A. Levin. 2003. Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50 (4): 915-944.
- [11] Evans, G.W. and S. Honkapohja. 1998. Economic dynamics with learning: new stability results. *Review of Economic Studies* 65, 23-44.
- [12] Evans, G.W. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press: Princeton, N.J.

- [13] Evans, G.W. and S. Honkapohja. 2003. Expectations and the stability problem for optimal monetary policies. *Review of Economic Studies*, 70, 807-824.
- [14] Gaspar, Vito, Frank Smets and David Vestin. 2006. Adaptive Learning, Persistence, and Optimal Monetary Policy. *Journal of the European Economic Association* 4, 376-385.
- [15] Gaspar, Vito, Frank Smets and David Vestin. 2010. Inflation Expectations, Adaptive Learning, and Optimal Monetary Policy. In B. Friedman and M. Woodford (eds) *Handbook of Monetary Economics*, vol. 3B, North Holland.
- [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] 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.
- [20] 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.
- [21] 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.
- [22] Mertens, Elmar. 2009a. Managing Beliefs about Monetary Policy under Discretion. Unpublished manuscript, Federal Reserve Board.
- [23] Mertens, Elmar. 2009b. Discreet Commitments and Discretion of Policymakers with Private Information. Unpublished manuscript, Federal Reserve Board.
- [24] Milani, Fabio. 2007. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54 (7): 2065-82
- [25] 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*

- [26] Orphanides, A. and J.C. Williams, 2007. Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54: 1406-1435.
- [27] Primiceri, G. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72, 821-852.
- [28] 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.
- [29] Sbordone, A.M. 2007. Inflation persistence: Alternative interpretations and policy implications. *Journal of Monetary Economics* 54, 1311–1339.
- [30] Sims, C.A. 2001. Solving Linear Rational Expectations Models. *Computational Economics* 20, 1-20.
- [31] Smets, F. and R. Wouters. 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97 (3): 586-606.
- [32] Woodford, M. 1999. Optimal Monetary Policy Inertia, NBER wp 7261.
- [33] Woodford, M. 2003. *Interest and Prices*. Princeton University Press: Princeton NJ.

Staff Reports, 2011

RESPONSES TO THE FINANCIAL CRISIS, TREASURY DEBT, AND THE IMPACT ON SHORT-TERM MONEY MARKETS

Warren B. Hrungr and Jason S. Seligman. Number 481, January 2011.

HOUSEHOLD DEBT AND SAVING DURING THE 2007 RECESSION

Rajashri Chakrabarti, Donghoon Lee, Wilbert van der Klaauw, and Basit Zafar
Number 482, January 2011.

STIGMA IN FINANCIAL MARKETS: EVIDENCE FROM LIQUIDITY AUCTIONS AND DISCOUNT WINDOW BORROWING DURING THE CRISIS

Olivier Armantier, Eric Ghysels, Asani Sarkar, and Jeffrey Shrader
Number 483, January 2011.

COMOVEMENT REVISITED

Maria Kasch and Asani Sarkar. Number 484, February 2011.

BASEL III: LONG-TERM IMPACT ON ECONOMIC PERFORMANCE AND FLUCTUATIONS

Paolo Angelini, Laurent Clerc, Vasco Cúrdia, Leonardo Gambacorta, Andrea Gerali, Alberto Locarno, Roberto Motto, Werner Roeger, Skander Van den Heuvel, and Jan Vlček.
Number 485, February 2011.

VOUCHERS, RESPONSES, AND THE TEST-TAKING POPULATION: REGRESSION DISCONTINUITY EVIDENCE FROM FLORIDA

Rajashri Chakrabarti. Number 486, March 2011.

CENTRAL BANK TRANSPARENCY AND THE CROWDING OUT OF PRIVATE INFORMATION IN AN EXPERIMENTAL ASSET MARKET

Menno Middelorp and Stephanie Rosenkranz. Number 487, March 2011.

LIQUIDITY HOARDING

Douglas Gale and Tanju Yorulmazer. Number 488, March 2011.

EXPECTATIONS OF INFLATION: THE BIASING EFFECT OF THOUGHTS ABOUT SPECIFIC PRICES

Wändi Bruine de Bruin, Wilbert van der Klaauw, and Giorgio Topa. Number 489, April 2011.

ROBUST CAPITAL REGULATION

Viral Acharya, Hamid Mehran, Til Schuermann, and Anjan Thakor. Number 490, April 2011;
revised June 2011.

FOMC COMMUNICATION POLICY AND THE ACCURACY OF FED FUNDS FUTURES

Menno Middelorp. Number 491, April 2011.

BANK CAPITAL REGULATION AND STRUCTURED FINANCE

Antoine Martin and Bruno M. Parigi. Number 492, May 2011; revised July 2011.

EFFICIENT, REGRESSION-BASED ESTIMATION OF DYNAMIC ASSET PRICING MODELS
Tobias Adrian, Richard K. Crump, and Emanuel Moench. Number 493, May 2011.

ARE CREDIT DEFAULT SWAPS ASSOCIATED WITH HIGHER CORPORATE DEFAULTS?
Stavros Peristiani and Vanessa Savino. Number 494, May 2011.

SECTORAL PRICE FACTS IN A STICKY-PRICE MODEL
Carlos Carvalho and Jae Won Lee. Number 495, May 2011.

CENTRAL BANK TRANSPARENCY, THE ACCURACY OF PROFESSIONAL FORECASTS,
AND INTEREST RATE VOLATILITY
Menno Middeldorp. Number 496, May 2011.

A NOTE ON BANK LENDING IN TIMES OF LARGE BANK RESERVES
Antoine Martin, James McAndrews, and David Skeie. Number 497, May 2011.

A MODEL OF LIQUIDITY HOARDING AND TERM PREMIA IN INTER-BANK MARKETS
Viral V. Acharya and David Skeie. Number 498, May 2011.

GLOBAL BOND RISK PREMIUMS
Rebecca Hellerstein. Number 499, June 2011.

DETERMINANTS OF COLLEGE MAJOR CHOICE: IDENTIFICATION USING AN INFORMATION
EXPERIMENT
Matthew Wiswall and Basit Zafar. Number 500, June 2011.

STEREOTYPES AND MADRASSAS: EXPERIMENTAL EVIDENCE FROM PAKISTAN
Adeline Delavande and Basit Zafar. Number 501, June 2011.

CORPORATE GOVERNANCE AND BANKS: WHAT HAVE WE LEARNED FROM THE FINANCIAL CRISIS?
Hamid Mehran, Alan Morrison, and Joel Shapiro. Number 502, June 2011.

THE PRODUCTION IMPACT OF “CASH-FOR-CLUNKERS”: IMPLICATIONS FOR STABILIZATION POLICY
Adam Copeland and James Kahn. Number 503, July 2011.

THE EMPIRICAL CONTENT OF MODELS WITH MULTIPLE EQUILIBRIA IN ECONOMIES
WITH SOCIAL INTERACTIONS
Alberto Bisin, Andrea Moro, and Giorgio Topa. Number 504, July 2011.

SUSTAINABLE SOCIAL SECURITY: FOUR OPTIONS
Sagiri Kitao. Number 505, July 2011.

REPO RUNS: EVIDENCE FROM THE TRI-PARTY REPO MARKET
Adam Copeland, Antoine Martin, and Michael Walker. Number 506, July 2011.

MAPPING CHANGE IN THE FEDERAL FUNDS MARKET
Morten L. Bech, Carl T. Bergstrom, Rodney J. Garratt, and Martin Rosvall.
Number 507, August 2011.

THE DYNAMICS AND DIFFERENTIATION OF LATIN AMERICAN METAL EXPORTS
Benjamin Mandel. Number 508, August 2011.

INFLATION EXPECTATIONS AND BEHAVIOR: DO SURVEY RESPONDENTS ACT ON THEIR BELIEFS?
Olivier Armantier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. Number 509, August 2011.

EVALUATING INTEREST RATE RULES IN AN ESTIMATED DSGE MODEL
Vasco Cúrdia, Andrea Ferrero, Ging Cee Ng, and Andrea Tambalotti. Number 510, August 2011.

LIQUIDITY MANAGEMENT OF U.S. GLOBAL BANKS: INTERNAL CAPITAL MARKETS IN THE GREAT RECESSION
Nicola Cetorelli and Linda Goldberg. Number 511, August 2011.

THE PRE-FOMC ANNOUNCEMENT DRIFT
David O. Lucca and Emanuel Moench. Number 512, September 2011.

DECOMPOSING SHORT-TERM RETURN REVERSAL
Zhi Da, Qianqiu Liu, and Ernst Schaumburg. Number 513, September 2011.

REAL ESTATE INVESTORS, THE LEVERAGE CYCLE, AND THE HOUSING MARKET CRISIS
Andrew Haughwout, Donghoon Lee, Joseph Tracy, and Wilbert van der Klaauw. Number 514, September 2011.

LEARNING THE FISCAL THEORY OF THE PRICE LEVEL: SOME CONSEQUENCES OF DEBT MANAGEMENT POLICY
Stefano Eusepi and Bruce Preston. Number 515, September 2011.

BELIEF UPDATING AMONG COLLEGE STUDENTS: EVIDENCE FROM EXPERIMENTAL VARIATION IN INFORMATION
Matthew Wiswall and Basit Zafar. Number 516, September 2011.

AN ANALYSIS OF CDS TRANSACTIONS: IMPLICATIONS FOR PUBLIC REPORTING
Kathryn Chen, Michael Fleming, John Jackson, Ada Li, and Asani Sarkar. Number 517, September 2011.

MARKET DECLINES: IS BANNING SHORT SELLING THE SOLUTION?
Robert Battalio, Hamid Mehran, and Paul Schultz. Number 518, September 2011.

EXPECTATIONS VERSUS FUNDAMENTALS: DOES THE CAUSE OF BANKING PANICS MATTER FOR PRUDENTIAL POLICY?
Todd Keister and Vijay Narasiman. Number 519, October 2011.

THE GREAT ESCAPE? A QUANTITATIVE EVALUATION OF THE FED'S LIQUIDITY FACILITIES
Marco Del Negro, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki. Number 520, October 2011.

EARLY CONTRACT RENEGOTIATION: AN ANALYSIS OF U.S. LABOR CONTRACTS FROM 1970 TO 1995

Robert Rich and Joseph Tracy. Number 521, October 2011.

THE INTERNATIONAL ROLE OF THE DOLLAR: DOES IT MATTER IF THIS CHANGES?

Linda Goldberg. Number 522, October 2011.

DO WE KNOW WHAT WE OWE? A COMPARISON OF BORROWER- AND LENDER-REPORTED CONSUMER DEBT

Meta Brown, Andrew Haughwout, Donghoon Lee, and Wilbert van der Klaauw. Number 523, October 2011.