

Appendices for “Optimized Taylor Rules for Disinflation When Agents are Learning”*

Timothy Cogley[†] Christian Matthes[‡] Argia M. Sbordone[§]

January 2015

A The model

The model is composed of a representative household that supplies labor and consumes a final good; monopolistically competitive firms that produce intermediate goods and set prices in a staggered way; a perfectly competitive final good producer, and a central bank that sets monetary policy. Here we describe the problems of household and firms and derive the equilibrium conditions presented in the main text.

A.1 The demand side

The representative household chooses consumption and hours of work to maximize expected discounted utility

$$E_t \sum_{s=0}^{\infty} \beta^s \left(b_{t+s} \log(C_{t+s} - \eta C_{t+s-1}) - \chi_{t+s} \frac{H_{t+s}^{1+\nu}}{1+\nu} \right), \quad (1)$$

subject to a flow budget constraint

$$E_t(Q_{t,t+1}Z_{t+1}) + P_t C_t = Z_t + W_t H_t + \int_0^1 \Psi_t(i) di. \quad (2)$$

*These appendices are not intended for publication. The figures shown below are best viewed in color.

[†]Department of Economics, New York University, 19 W. 4th St., 6FL, New York, NY, 10012, USA. Email: tim.cogley@nyu.edu. Tel. 212-992-8679.

[‡]Research Department, Federal Reserve Bank of Richmond, 701 East Byrd Street, Richmond, VA 23219, USA. Email: christian.matthes@rich.frb.org. Tel. 804-697-4490.

[§]Research Department, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045, USA. Email: argia.sbordone@ny.frb.org. Tel. 212-720-6810.

In the preferences, η measures a degree of internal habit persistence, β is a subjective discount factor and b_t and χ_t are white noise preference shocks. C_t is consumption of the final good, and P_t denotes its price. H_t is an aggregate of the hours supplied by the household to intermediate good-producing firms. Hours are paid an economy-wide nominal wage W_t .

In the budget constraint, $\int_0^1 \Psi_t(i) di$ is profits of intermediate good producers rebated to the household, Z_{t+1} is the state-contingent value of the portfolio of assets held by the household at the beginning of period $t + 1$, and $Q_{t,t+1}$ is a stochastic discount factor.

The first order condition for the choice of consumption is

$$\Xi_t = \beta E_t \left[\Xi_{t+1} \frac{R_t}{\Pi_{t+1}} \right], \quad (3)$$

where Ξ_t is the marginal utility of consumption at t ,

$$\Xi_t = \frac{b_t}{C_t - \eta C_{t-1}} - \beta \eta E_t \frac{b_{t+1}}{C_{t+1} - \eta C_t}, \quad (4)$$

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

$$w_t = \chi_t H_t^\nu \Xi_t^{-1}, \quad (5)$$

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

Growth in this economy is driven by an aggregate technological progress Γ_t (introduced below). We therefore define normalized variables $C_t^a \equiv C_t/\Gamma_t$, $Y_t^a \equiv Y_t/\Gamma_t$, $\gamma_t = \ln \Gamma_t/\Gamma_{t-1}$, and $\Xi_t^a \equiv \Xi_t \Gamma_t$ and, imposing the aggregate resource constraint $C_t^a = Y_t^a$, we express (4) and the equilibrium condition (3) respectively as

$$\Xi_t^a = \frac{b_t}{Y_t^a - \eta Y_{t-1}^a \frac{1}{\gamma_t}} - \beta \eta E_t \frac{b_{t+1}}{Y_{t+1}^a \gamma_{t+1} - \eta Y_t^a} \quad (6)$$

and

$$\Xi_t^a = \beta E_t \left[\Xi_{t+1}^a (\gamma_{t+1})^{-1} \frac{R_t}{\Pi_{t+1}} \right]. \quad (7)$$

Similarly, we re-write the first order condition for labor supply as

$$w_t^a = \chi_t H_t^\nu (\Xi_t^a)^{-1}, \quad (8)$$

where $w_t^a \equiv w_t/\Gamma_t$ is the productivity adjusted real wage.

A.2 The supply side

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

$$Y_t = \left[\int_0^1 y_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \quad (9)$$

where θ is the elasticity of substitution across intermediate goods. She chooses intermediate inputs to maximize her profits, taking the price of the final good P_t as given; this determines demand schedules

$$y_t(i) = Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta}. \quad (10)$$

The zero-profit condition then determines the aggregate price level

$$P_t \equiv \left[\int_0^1 p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}}. \quad (11)$$

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

$$y_t(i) = \Gamma_t h_t(i), \quad (12)$$

where Γ_t is an aggregate technological process, whose rate of growth $\gamma_t \equiv \ln \Gamma_t / \Gamma_{t-1}$ evolves as

$$\gamma_t = (1 - \rho_g)\gamma + \rho_g\gamma_{t-1} + \varepsilon_{\gamma t}.$$

We assume staggered Calvo price-setting: intermediate good producers can reset prices at random intervals, and we denote by $1 - \alpha$ the reset probability. The first order condition for the choice of optimal price p_t^* is

$$E_t \sum_{j=0}^{\infty} \alpha^j Q_{t,t+j} Y_{t+j} P_{t+j}^{\theta} \left(p_t^* - \frac{\theta}{\theta-1} MC_{t+j} \right) = 0,$$

where MC_t denotes nominal marginal costs and the index i is suppressed, since all optimizing firms solve the same problem. This condition and the evolution of aggregate prices

$$P_t = \left[(1 - \alpha)p_t^{*1-\theta} + \alpha P_{t-1}^{1-\theta} \right]^{\frac{1}{1-\theta}} \quad (13)$$

jointly determine the dynamics of inflation in the model.

A.2.1 Marginal costs, output and price dispersion

The first order conditions for optimal price setting imply that the optimal price is function of expected future marginal costs MC_{t+j} . These can in turn be expressed as function of aggregate output. Specifically, in equilibrium real marginal cost mc_t ($\equiv MC_t/P_t$) is equal to real wage corrected for productivity

$$mc_t = w_t^a, \quad (14)$$

where the latter is defined by the equilibrium condition (8). Aggregate hours H_t are obtained by aggregating hours worked in each intermediate firm:

$$H_t \equiv \int_0^1 h_t(i) di = \int_0^1 \frac{y_t(i)}{\Gamma_t} di = Y_t^a \int_0^1 \left(\frac{p_t(i)}{P_t} \right)^{-\theta} di = Y_t^a \Delta_t, \quad (15)$$

where we denoted by Δ_t the following measure of price dispersion: $\Delta_t \equiv \int_0^1 \left(\frac{p_t(i)}{P_t} \right)^{-\theta} di$. One can see that aggregate output is equal to the ratio of aggregate hours and the measure of price dispersion:

$$Y_t^a = \frac{H_t}{\Delta_t}, \quad (16)$$

so that in equilibrium higher price dispersion implies that more hours are needed to produce the same amount of output (indeed, labor productivity is the inverse of the price dispersion index.) Substituting expressions (8) and (15) in (14) we obtain a relationship between marginal costs and output, where price dispersion creates a wedge between the two:

$$mc_t = \chi_t H_t^\nu (\Xi_t^a)^{-1} = \chi_t (Y_t^a \Delta_t)^\nu (\Xi_t^a)^{-1}. \quad (17)$$

We will use this expression to derive a Phillips curve in terms of output.

A.3 Steady-state relations

From the definition of Δ_t we can derive that¹

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

where \tilde{p}_t denotes the relative price of the firms that optimizes at t : $\tilde{p}_t \equiv p_t^*(i)/P_t$, whose value can be obtained from the evolution of aggregate prices (13). Price dispersion is therefore the following function of the inflation rate

$$\Delta_t = (1 - \alpha) \left(\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} \right)^{-\frac{\theta}{1-\theta}} + \alpha \Pi_t^\theta \Delta_{t-1}, \quad (18)$$

¹See Schmitt-Grohe and Uribe (2006, 2007).

and in steady state:

$$\bar{\Delta}_t = \frac{1 - \alpha}{1 - \alpha \bar{\Pi}_t^\theta} \left(\frac{1 - \alpha \bar{\Pi}_t^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta}{\theta-1}}. \quad (19)$$

Similarly, from the first order condition of price setting, we can derive the following relation between steady state marginal cost and steady state inflation:

$$\bar{m}c_t = \frac{\theta - 1}{\theta} \frac{\left(1 - \alpha \bar{\Pi}_t^{\theta-1}\right)^{\frac{1}{1-\theta}}}{(1 - \alpha)^{\frac{1}{1-\theta}}} \left[\frac{1 - \alpha \beta (\bar{\Pi}_t)^\theta}{1 - \alpha \beta (\bar{\Pi}_t)^{\theta-1}} \right]. \quad (20)$$

Substituting (19) and (20) in (17), evaluated in steady state, gives a relationship between inflation and output that should be satisfied in steady state:

$$\bar{Y}_t^a = \left[\frac{\frac{\theta-1}{\theta} \frac{\left(1 - \alpha \bar{\Pi}_t^{\theta-1}\right)^{\frac{1}{1-\theta}}}{(1-\alpha)^{\frac{1}{1-\theta}}} \left[\frac{1 - \alpha \beta (\bar{\Pi}_t)^\theta}{1 - \alpha \beta (\bar{\Pi}_t)^{\theta-1}} \right]}{\frac{e^\gamma - \eta}{e^{\gamma - \beta \eta}} \left(\frac{1 - \alpha}{1 - \alpha \bar{\Pi}_t^\theta} \left(\frac{1 - \alpha \bar{\Pi}_t^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta}{\theta-1}} \right)^\nu} \right]^{\frac{1}{1+\nu}}, \quad (21)$$

where $\frac{e^\gamma - \beta \eta}{e^{\gamma - \eta}}$ is the steady state value of Ξ_t^a . This relationship can be interpreted as a long-run Phillips curve.

A.4 Log-linearizations

For the demand side, the dynamic *IS* block is obtained by log-linearizing equilibrium conditions (6) and (7). It is convenient to define transformed variables $\tilde{\Xi}_t = \Xi_t^a Y_t^a$ and $\tilde{Y}_t^a = Y_t^a / \bar{Y}_t^a$ (with steady state values, respectively, of Ξ and 1). With these transformations, the log-linearized *IS* equation is

$$\hat{\Xi}_t = \hat{Y}_t^a + E_t \left(\hat{\Xi}_{t+1} - \hat{Y}_{t+1}^a - \hat{\gamma}_{t+1} - \hat{\gamma}_{y_{t+1}} + i_t - \pi_{t+1} - r \right) \quad (22)$$

where $\hat{\Xi}_t \equiv \ln \tilde{\Xi}_t / \Xi$ is defined as follows

$$\hat{\Xi}_t = \xi_1 \hat{Y}_t^a + \xi_2 \left[\left(\hat{Y}_{t-1}^a - \hat{\gamma}_t \right) + \beta E_t \left(\hat{Y}_{t+1}^a + \hat{\gamma}_{t+1} + \hat{\gamma}_{y_{t+1}} \right) \right] + \varepsilon_{yt}. \quad (23)$$

The hat variables are, as usual, log deviations from steady state: $\hat{Y}_t^a = \ln Y_t^a - \ln \bar{Y}_t^a$, $i_t = \ln R_t$, $\hat{\gamma}_t = \gamma_t - \gamma$, $\gamma_{yt} \equiv \bar{Y}_t^a / \bar{Y}_{t-1}^a$, r is the steady state real interest rate, and the disturbance ε_{yt} is a transformation of the preference shock b_t .²

²Note the term R_t / Π_{t+1} is stationary, and we denote its (log) steady state (which is equal to the steady state value of the ratio of nominal interest rate to trend inflation) by r . This can be seen by dividing through by $\bar{\Pi}_t$, which gives $\frac{R_t / \bar{\Pi}_t}{(\Pi_{t+1} / \bar{\Pi}_{t+1})(\bar{\Pi}_{t+1} / \bar{\Pi}_t)} = \frac{R_t^{r\pi}}{\bar{\Pi}_t g_{t+1}^\pi}$ whose steady state we denote by $R^{r\pi}$, and $r \equiv \log R^{r\pi}$.

Equations (22) and (23) deliver, respectively, equations (4) and (5) in the main text, by the use of a simplified notation, where $\xi_t \equiv \ln \tilde{\Xi}_t$, $\xi \equiv \ln \Xi$, $y_t \equiv \ln Y_t^a$, $\bar{y}_t \equiv \ln \bar{Y}_t^a$ and $\bar{y}_{t-1} \equiv E_t^* \bar{y}_t$.³

As explained in the text, we also replace rational expectations with subjective expectations and, consistently with this assumption, we take the approximations around the agents' perception of the steady state values. Finally, since $E_t^* \hat{\gamma}_{y_{t+1}} = 0$, that term is suppressed.

To obtain the new-Keynesian Phillips curve, we start from the log-linear NKPC developed in Cogley and Sbordone (2008), where the forcing variable is marginal cost⁴

$$\begin{aligned}\hat{\pi}_t &= \beta_t E_t^* \hat{\pi}_{t+1} + \tilde{\kappa}_{t-1} \widehat{mc}_t + \gamma_{1t} E_t^* [(\theta - 1) \hat{\pi}_{t+1} + \phi_{t+1}] + u_t, \\ \phi_t &= \gamma_{2t} E_t^* [(\theta - 1) \hat{\pi}_{t+1} + \phi_{t+1}].\end{aligned}\quad (24)$$

The parameters are defined in expression (11) in the main text.

Then we transform this equation in an inflation-output dynamic relation, by log-linearizing expression (17) to obtain

$$\widehat{mc}_t = (1 + \nu) \widehat{Y}_t^a + \nu \widehat{\Delta}_t - \widehat{\Xi}_t, \quad (25)$$

which we substitute into (24). In this expression, $\widehat{\Delta}_t \equiv \ln \Delta_t / \bar{\Delta}_t$ is obtained by log-linearizing (18) around steady state, which gives

$$\widehat{\Delta}_t \simeq \lambda_{1t} \widehat{\Pi}_t + \lambda_{2t} (\widehat{\Delta}_{t-1} - \widehat{\gamma}_{\Delta t}). \quad (26)$$

The parameters λ_{1t} and λ_{2t} are defined in the last two rows of (11) in the main text. As most of the other parameters, they are time-varying because they depend on trend inflation. In the main text, for analogy with the other log-linearized equations, instead of notation \widehat{mc}_t and $\widehat{\Delta}_t$ we use the corresponding notation $mc_t - \bar{mc}_t$ and $\delta_t - \bar{\delta}_t$, respectively.

A.5 Structural arrays

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

³The expressions for ξ_1 and ξ_2 are: $\xi_1 = -(\exp(\gamma)\eta(1 + \beta))/((\exp(\gamma) - \eta)(\exp(\gamma) - \beta\eta))$, and $\xi_2 = (\exp(\gamma)\eta)/((\exp(\gamma) - \eta)(\exp(\gamma) - \beta\eta))$.

⁴With minor changes in notation, these equations corresponds to eqs. (46) and (47) in Cogley and Sbordone (2008), simplified to reflect the absence of price indexation and strategic complementarities in the present model. Details on the derivation of the equations can be found in the cited paper.

$$A_t = \begin{bmatrix} 1 & 0 & -\varsigma_t & -1 & -\kappa_t & 0 & 0 & 0 & 0 & \tilde{\kappa}_t \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda_{1t} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -\xi_1 & \xi_2 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (27)$$

$$B_t = \begin{bmatrix} \beta_t + \gamma_{1t}(\theta - 1) & \gamma_{1t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \gamma_{2t}(\theta - 1) & \gamma_{2t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \beta\xi_2 & \beta\xi_2 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (28)$$

$$C_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} & 0 \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} & 0 \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_\gamma & 0 & 0 & \mu_\gamma & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \psi_{\pi t} & 0 & 0 & 0 & \psi_{yt} & 0 & -\psi_{yt} & 1 & -\psi_{\pi t} \bar{\pi}_t & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \xi_2 & 0 & 0 & 0 & \mu_\xi & 0 \end{bmatrix}, \quad (29)$$

$$D_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (30)$$

The expressions for the intercepts in C_t are

$$\begin{aligned}
\mu_{\pi t} &= [1 - \beta_t - \gamma_{1t}(\theta - 1)]\bar{\pi}_t - \kappa_t\bar{y}_t - \varsigma_t\bar{\delta}_t + \tilde{\kappa}\xi, \\
\mu_{\phi_t} &= -\gamma_{2t}(\theta - 1)\bar{\pi}_t, \\
\mu_{\delta t} &= (1 - \lambda_{2t})\bar{\delta}_{t-1} - \lambda_{1t}\bar{\pi}_t, \\
\mu_y &= r - \gamma, \\
\mu_\gamma &= (1 - \rho_\gamma)\gamma, \\
\mu_\xi &= \xi - (\xi_1 + (1 + \beta)\xi_2)\bar{y}_t + \gamma\xi_2(1 - \beta),
\end{aligned} \tag{31}$$

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

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

$$C_{at} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} & 0 \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} & 0 \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_\gamma & 0 & 0 & \mu_\gamma & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \psi_\pi & 0 & 0 & 0 & \psi_y & 0 & -\psi_y & 1 & -\psi_\pi\bar{\pi} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \xi_2 & 0 & 0 & 0 & \mu_\xi & 0 \end{bmatrix}. \tag{32}$$

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

$$e_X = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \tag{33}$$

B The relative importance of uncertainty about feedback parameters and target inflation

To determine the relative importance of uncertainty about feedback parameters and target inflation, we contrast a pair of models that shut down one or the other. The first model deactivates uncertainty about ψ_π , ψ_y , and σ_i while retaining uncertainty about trend inflation. All other aspects of the baseline specification are the same, including the prior for $\bar{\pi}$. In this case, the initial nonexplosive region expands to fill most of the (ψ_π, ψ_y) space (see the top panel in figure B1).

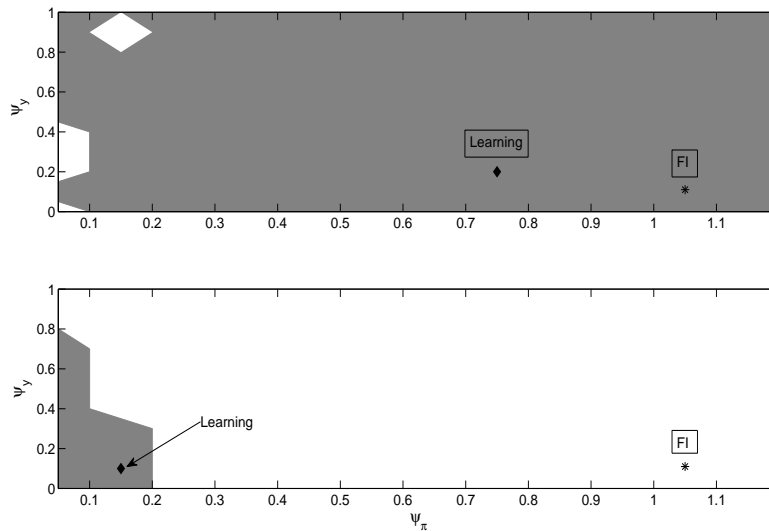


Figure B1: Gray areas mark the nonexplosive region for H_1 . In the top row, the feedback parameters are known and $\bar{\pi}$ is unknown. In the bottom row, $\bar{\pi}$ is known and the feedback parameters are unknown.

Since the ALM is nonexplosive for most policies, the model has high fault tolerance (see the left column of figure B2). Furthermore, private agents learn $\bar{\pi}$ very quickly (see the top left panel of figure B4). 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 the main text (see the left panel of figure B3).

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 in the main text. 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 (see the bottom panel in figure B1). Because of concerns about explosive volatility, the bank chooses a policy close to the prior mode for ψ_π and ψ_y (see the bottom right panel of figure B2). The transition is volatile (see the right panel of figure B3), but learning is rapid because the true policy is close to initial beliefs (see the right column of figure B4).

From this we conclude that uncertainty about feedback parameters is more costly. Uncertainty about target inflation is a lesser evil.

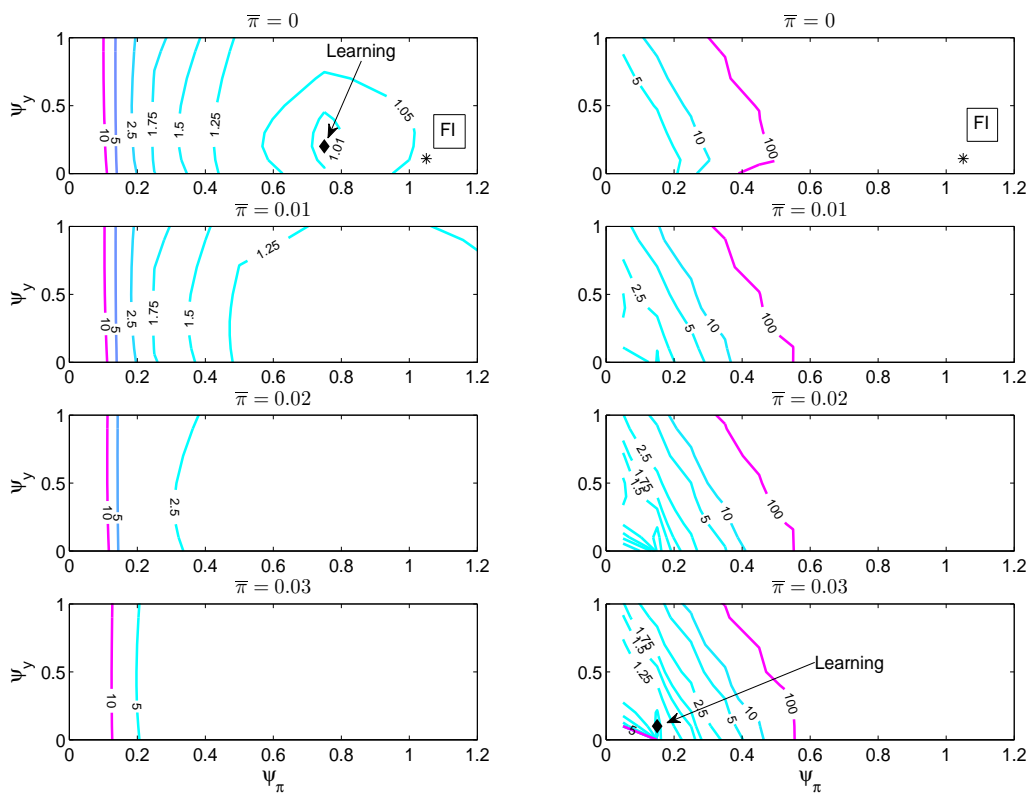


Figure B2: Iso-expected-loss contours. In the left column, the feedback parameters are known and $\bar{\pi}$ is unknown. In the right column, $\bar{\pi}$ is known and the feedback parameters are unknown.




Figure B3: Average responses under the optimized rule. In the left column, the feedback parameters are known and $\bar{\pi}$ is unknown. In the right column, $\bar{\pi}$ is known and the feedback parameters are unknown.

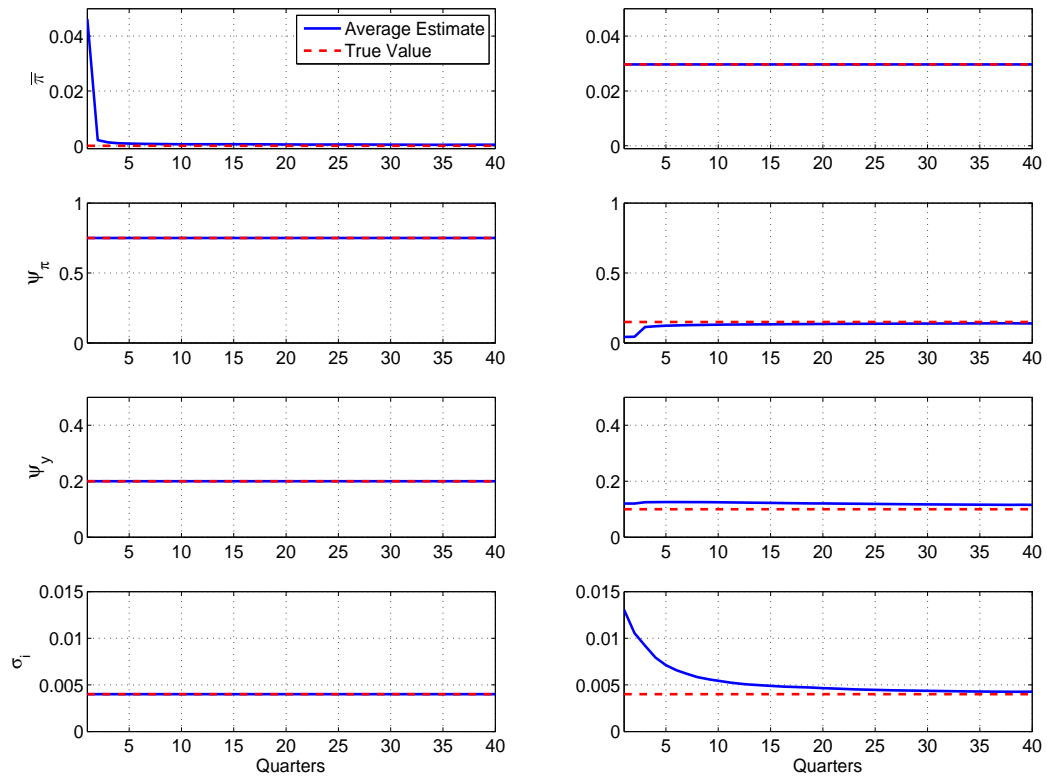


Figure B4: Average estimates of policy coefficients. In the left column, the feedback parameters are known and $\bar{\pi}$ is unknown. In the right column, $\bar{\pi}$ is known and the feedback parameters are unknown.

C Adaptive expectations

Here we examine how our model behaves when learning is deactivated and replaced by adaptive expectations. For the model described in the main text, the ALM solves

$$A_t S_t = B_t E_t^* S_{t+1} + C_{at} S_{t-1} + D_t \varepsilon_t, \quad (34)$$

where E_t^* denotes expectations taken with respect to the PLM at date t . We replace this expectations operator with adaptive expectations,

$$E_t^a S_{t+1} = E_{t-1}^a S_t + k (S_t - E_{t-1}^a S_t), \quad (35)$$

where k is a smoothing factor that governs how fast expectations correct forecast errors. Because agents in this version of the model do not explicitly estimate the coefficients of the policy rule (which influence the steady state of the economy) or the steady state directly, we log-linearize the equilibrium conditions each period around the final steady state (instead of the perceived steady state each period). Otherwise the model is identical to the version presented in the main text.

After adding $E_t^a S_{t+1}$ as a state variable and dropping time t subscripts on the system matrices, the model can be expressed as

$$\begin{bmatrix} A & -B \\ -kI & I \end{bmatrix} \begin{bmatrix} S_t \\ E_t^a S_{t+1} \end{bmatrix} = \begin{bmatrix} C_a & 0 \\ 0 & (1-k)I \end{bmatrix} \begin{bmatrix} S_{t-1} \\ E_{t-1}^a S_t \end{bmatrix} + \begin{bmatrix} D \\ 0 \end{bmatrix} \varepsilon_t.$$

It follows that the equilibrium law of motion is

$$\begin{aligned} \begin{bmatrix} S_t \\ E_t^a S_{t+1} \end{bmatrix} &= \begin{bmatrix} A & -B \\ -kI & I \end{bmatrix}^{-1} \left\{ \begin{bmatrix} C_a & 0 \\ 0 & (1-k)I \end{bmatrix} \begin{bmatrix} S_{t-1} \\ E_{t-1}^a S_t \end{bmatrix} + \begin{bmatrix} D \\ 0 \end{bmatrix} \varepsilon_t \right\}, \\ &= \begin{bmatrix} (A - kB)^{-1} C_a & (1-k)(A - kB)^{-1} B \\ k(A - kB)^{-1} C_a & (1-k)(A - kB)^{-1} A \end{bmatrix} \begin{bmatrix} S_{t-1} \\ E_{t-1}^a S_t \end{bmatrix} + \begin{bmatrix} (A - kB)^{-1} D \\ k(A - kB)^{-1} D \end{bmatrix} \varepsilon_t. \end{aligned}$$

Expectations are initialized at the steady state from the old regime, $E_0^a S_1 = S_{old}^{ss}$. So that expectations are sluggish, we examine adjustment coefficients ranging from $k = 0.2$ to 0.5 . Because the model is backward looking, policies optimized for the learning environment considered in the main text work poorly in this case. We therefore seek optimal simple rules for each value of k .⁵ The results are shown in table C1.

⁵Losses under the optimal policy increase by several orders of magnitude if k is increased further.

Table C1: Optimal Simple Rules and Sacrifice Ratios under Adaptive Expectations

k	$\bar{\pi}$	ψ_{π}	ψ_y	ρ_i	Sacrifice Ratio
0.2	0	1.5	0	0.5	1.40
0.3	0	1.5	0	0.5	1.42
0.4	0	1.5	0	0.25	1.38
0.5	0.005	1.5	0	0.25	1.68

For instance, for $k = 0.2$, the optimum is $\bar{\pi} = 0$, $\psi_{\pi} = 1.5$, $\psi_y = 0$ and $\rho_i = 0.5$.⁶ As shown in figure C1, impulse response functions under this policy are smoother than under learning. The response of output is shallower but more prolonged, with the output gap remaining negative for 28 quarters. The cumulative output loss during this period is 6.3 percent. Since inflation declines permanently by 4.6 percent, the sacrifice ratio is 1.4. Results for other values of k are similar, with sacrifice ratios ranging from 1.38 to 1.68.

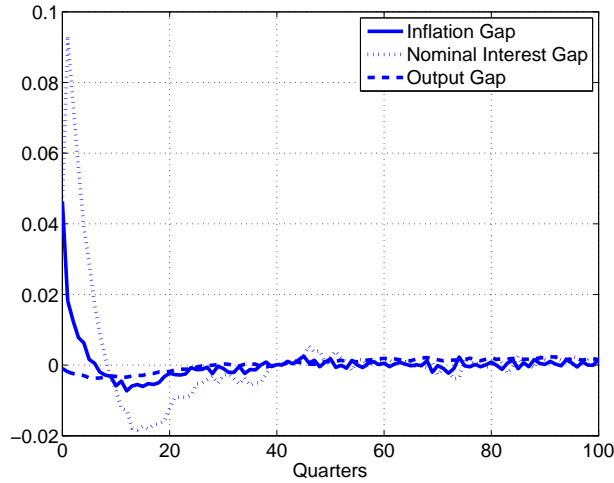


Figure C1: Impulse response functions under the policy optimized for adaptive expectations ($k = 0.2$)

⁶Restricting the search for the optimal policy rule to rules that feature $\rho_i = 1$ leads to substantially worse outcomes than in our benchmark. In contrast, when agents learn about the policy rule coefficients, the optimal ρ_i is close to 1.

D A stripped-down example

Our minimalist model consists of three equations,

$$\pi_t = \beta E_t^* \pi_{t+1} + x_t + u_t, \quad (36)$$

$$x_t = -\psi \pi_{t-1} + \varepsilon_{xt}, \quad (37)$$

$$u_t = \rho_u u_{t-1} + \varepsilon_{ut}, \quad (38)$$

where π_t is inflation, u_t is an exogenous cost-push shock, and x_t is an abstract policy instrument. The expectations operator E_t^* signifies that forecasts are made with respect to a subjective probability model. Equation (36) is a stylized version of the NKPC, with x_t representing an abstract policy instrument, equation (37) is a simple policy rule, and equation (38) is the law of motion for the cost-push shock. The innovations $\varepsilon_t = [\varepsilon_{ut}, \varepsilon_{xt}]'$ are iid normal with mean zero and variance

$$\text{var}(\varepsilon_t) = \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_u^2 \end{bmatrix}. \quad (39)$$

The perceived policy is

$$x_t = -\psi_t \pi_{t-1} + \tilde{\varepsilon}_{xt}, \quad (40)$$

where $\tilde{\varepsilon}_{xt}$ is a perceived policy shock and ψ_t represents a beginning-of-period t estimate of ψ . By subtracting the actual policy from the perceived rule, the perceived shock can be expressed as

$$\tilde{\varepsilon}_{xt} = \varepsilon_{xt} + (\psi_t - \psi) \pi_{t-1}. \quad (41)$$

Agents believe that $\tilde{\varepsilon}_{xt}$ is iid normal with mean zero and variance σ_{xt}^2 .

To solve for the PLM, substitute the perceived policy into the NKPC and rearrange terms to obtain an expectational difference equation in π_t ,

$$E_t^*(1 - \beta^{-1}L - \beta^{-1}\psi_t L^2)\pi_{t+1} = -\beta^{-1}(u_t + \tilde{\varepsilon}_{xt}). \quad (42)$$

The lag polynomial can be factored as

$$1 - \beta^{-1}L - \beta^{-1}\psi_t L^2 = (1 - \lambda_{1t}L)(1 - \lambda_{2t}L), \quad (43)$$

where λ_{1t} and λ_{2t} are reciprocals of the roots of

$$1 - \beta^{-1}L - \beta^{-1}\psi_t L^2 = 0. \quad (44)$$

Because this difference equation involves a single forward-looking variable, a unique nonexplosive solution exists when one root lies outside the unit circle and the other lies inside. We restrict parameters so that $|\lambda_{1t}| < 1$ and $|\lambda_{2t}| > 1$. We then solve the stable factor λ_{1t} backward and the unstable factor λ_{2t} forward,

$$\begin{aligned} E_t^* [(1 - \lambda_{1t}L)(1 - \lambda_{2t}L)\pi_{t+1}] &= -\beta^{-1}(u_t + \tilde{\varepsilon}_{xt}), \\ E_t^* [(1 - \lambda_{1t}L)(1 - \lambda_{2t}^{-1}L^{-1})\pi_t] &= \frac{1}{\beta\lambda_{2t}}(u_t + \tilde{\varepsilon}_{xt}), \\ E_t^* [(1 - \lambda_{2t}^{-1}L^{-1})q_t] &= \frac{1}{\beta\lambda_{2t}}(u_t + \tilde{\varepsilon}_{xt}), \end{aligned} \quad (45)$$

where $q_t \equiv (1 - \lambda_{1t}L)\pi_t$. Solving forward for q_t , using the law of iterated expectations, and imposing a no-bubbles condition yields

$$q_t = \frac{1}{\beta\lambda_{2t}} \sum_{j=0}^{\infty} \lambda_{2t}^{-j} E_t^*(u_{t+j} + \tilde{\varepsilon}_{xt+j}). \quad (46)$$

Subjective forecasts satisfy

$$\begin{aligned} E_t^* u_{t+j} &= \rho_u^j u_t, \\ E_t^* \tilde{\varepsilon}_{xt+j} &= 0, \quad j > 0, \end{aligned} \quad (47)$$

implying

$$q_t = \frac{1}{\beta\lambda_{2t}} \tilde{\varepsilon}_{xt} + \frac{1}{\beta\lambda_{2t}} \sum_{j=0}^{\infty} \lambda_{2t}^{-j} \rho_u^j u_t = \frac{1}{\beta\lambda_{2t}} \tilde{\varepsilon}_{xt} + \frac{1}{\beta\lambda_{2t}} \frac{1}{1 - \rho_u/\lambda_{2t}} u_t. \quad (48)$$

The geometric sum converges because $|\rho_u| < 1$ and $|\lambda_{2t}| > 1$. Unpacking q_t and using the shock process delivers the first equation of a reduced-form VAR,

$$\pi_t = \lambda_{1t}\pi_{t-1} + \frac{\rho_u}{\beta\lambda_{2t}(1 - \rho_u/\lambda_{2t})} u_{t-1} + \frac{1}{\beta\lambda_{2t}} \tilde{\varepsilon}_{xt} + \frac{1}{\beta\lambda_{2t}(1 - \rho_u/\lambda_{2t})} \varepsilon_{ut}. \quad (49)$$

The second equation of the VAR is just the shock process. It follows that the PLM is

$$\begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = \begin{bmatrix} \lambda_{1t} & \frac{\rho_u}{\beta\lambda_{2t}(1 - \rho_u/\lambda_{2t})} \\ 0 & \rho_u \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} \frac{1}{\beta\lambda_{2t}} & \frac{1}{\beta\lambda_{2t}(1 - \rho_u/\lambda_{2t})} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{\varepsilon}_{xt} \\ \varepsilon_{ut} \end{bmatrix}. \quad (50)$$

To find the ALM, we exploit the relation between perceived and actual shocks,

$$\begin{bmatrix} \tilde{\varepsilon}_{xt} \\ \varepsilon_{ut} \end{bmatrix} = \begin{bmatrix} \varepsilon_{xt} \\ \varepsilon_{ut} \end{bmatrix} + \begin{bmatrix} \psi_t - \psi & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ u_{t-1} \end{bmatrix}. \quad (51)$$

Substituting this into the PLM yields,

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

Thus the ALM is

$$\begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = \begin{bmatrix} \lambda_{1t} + \frac{1}{\beta\lambda_{2t}}(\psi_t - \psi) & \frac{\rho_u}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & \rho_u \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} \frac{1}{\beta\lambda_{2t}} & \frac{1}{\beta\lambda_{2t}(1-\rho_u/\lambda_{2t})} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_{xt} \\ \varepsilon_{ut} \end{bmatrix}. \quad (52)$$

E McCallum's information constraint

McCallum (1999) contends that monetary policy rules should be specified in terms of lagged variables because the Fed lacks good current-quarter information about inflation, output, and other arguments of policy reaction functions. For instance, the Bureau of Economic Analysis released the advance estimate of 2013.Q4 GDP on January 30, 2014, one month after the end of the quarter. This constraint plays a critical role in our analysis. To highlight its importance, we contrast the backward-looking Taylor rule in equation (1) in the text with one involving contemporaneous feedback to inflation and output growth,

$$i_t - i_{t-1} = \psi_\pi(\pi_t - \bar{\pi}) + \psi_y(y_t - y_{t-1}) + \varepsilon_{it}. \quad (53)$$

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

There is also a slight change in the timing protocol. Private agents still enter period t with beliefs about policy coefficients inherited from $t - 1$, and they treat estimated parameters as if they were known with certainty when updating decision rules. But now the central bank and private sector simultaneously execute their contingency plans when period t shocks are realized. After observing current-quarter outcomes, private agents update estimates and carry them forward to $t + 1$.

E.1 The perceived law of motion

Because of the change in timing protocol, the perceived and actual laws of motion differ from those in the text. As before, we assume that private agents know the form of the policy rule but not the policy coefficients, so that at any given date their perceived policy rule is

$$i_t - i_{t-1} = \psi_{\pi t}(\pi_t - \bar{\pi}_t) + \psi_{y t}(y_t - y_{t-1}) + \tilde{\varepsilon}_{it}, \quad (54)$$

where

$$\tilde{\varepsilon}_{it} = \varepsilon_{it} + (\psi_\pi - \psi_{\pi t})\pi_t + (\psi_y - \psi_{y t})\Delta y_t + \psi_{\pi t}\bar{\pi}_t - \psi_\pi\bar{\pi}. \quad (55)$$

is a perceived policy shock that depends on the actual policy shock ε_{it} and the estimated policy coefficients.

The private sector's model of the economy can be represented as a system of linear expectational difference equations,

$$A_t S_t = B_t E_t^* S_{t+1} + C_t S_{t-1} + D_t \tilde{\varepsilon}_t, \quad (56)$$

where S_t is the model's state vector, defined as before, and $\tilde{\varepsilon}_t$ is a vector of perceived innovations. The matrices A_t, B_t, C_t , and D_t entering the PLM are now defined as:

$$A_t = \begin{bmatrix} 1 & 0 & -\varsigma_t & -1 & -\kappa_t & 0 & 0 & 0 & 0 & \tilde{\kappa}_t \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda_{1t} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ -\psi_{\pi t} & 0 & 0 & 0 & -\psi_{yt} & 0 & \psi_{yt} & 1 & \psi_{\pi t} \bar{\pi}_t & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -\xi_1 & \xi_2 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (57)$$

$$B_t = \begin{bmatrix} \beta_t + \gamma_{1t}(\theta - 1) & \gamma_{1t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \gamma_{2t}(\theta - 1) & \gamma_{2t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \beta \xi_2 & \beta \xi_2 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (58)$$

$$C_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} & 0 \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} & 0 \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_\gamma & 0 & 0 & \mu_\gamma & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \xi_2 & 0 & 0 & 0 & \mu_\xi & 0 \end{bmatrix}, \quad (59)$$

$$D_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (60)$$

Note that the parameters of the policy rule enter the matrix A_t , rather than C_t . The expressions for the intercepts in C_t are

$$\begin{aligned} \mu_{\pi t} &= [1 - \beta_t - \gamma_{1t}(\theta - 1)]\bar{\pi}_t - \kappa_t \bar{y}_t - \varsigma_t \bar{\delta}_t + \tilde{\kappa} \xi, \\ \mu_{\phi_t} &= -\gamma_{2t}(\theta - 1)\bar{\pi}_t, \\ \mu_{\delta t} &= (1 - \lambda_{2t})\bar{\delta}_{t-1} - \lambda_{1t}\bar{\pi}_t, \\ \mu_y &= r - \gamma, \\ \mu_\gamma &= (1 - \rho_\gamma)\gamma, \\ \mu_\xi &= \xi - (\xi_1 + (1 + \beta)\xi_2)\bar{y}_t + \gamma\xi_2(1 - \beta), \end{aligned} \quad (61)$$

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

The PLM is the reduced-form VAR associated with (56),

$$S_t = F_t S_{t-1} + G_t \tilde{\varepsilon}_t, \quad (62)$$

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

E.2 The actual law of motion

The actual law of motion, obtained again by stacking the actual policy rule along with the *IS* curve, the aggregate supply block, and the shock processes, is another system of expectational difference equations,

$$A_{at} S_t = B_t E_t^* S_{t+1} + C_t S_{t-1} + D_t \varepsilon_t. \quad (63)$$

The matrices B_t , C_t , and D_t are the same as in equations (58), (59), and (60). All rows of the matrix A_{at} agree with those of A_t except for the one corresponding to the

monetary-policy rule (row 8). In that row, the true policy coefficients ψ replace the estimated coefficients ψ_t ,

$$A_{at}(8, \cdot) = \begin{bmatrix} -\psi_\pi & 0 & 0 & 0 & -\psi_y & 0 & \psi_y & 1 & \psi_\pi \bar{\pi} & 0 \end{bmatrix}. \quad (64)$$

To find the ALM, substitute $E_t^* S_{t+1} = F_t S_t$ from the PLM into (63) and re-arrange terms,⁷

$$S_t = H_t S_{t-1} + J_t \varepsilon_t, \quad (66)$$

where

$$\begin{aligned} H_t &= (A_{at} - B_t F_t)^{-1} C_t, \\ J_t &= (A_{at} - B_t F_t)^{-1} D_t. \end{aligned} \quad (67)$$

Comparing this expression to the one for the backward-looking rule, we observe that the matrix H_t does not depend directly upon the perceived coefficients of the policy rule, only indirectly through the PLM matrix F_t .

E.3 Quantitative analysis

All other aspects of the model remain the same, including the prior on $(\bar{\pi}, \psi_\pi, \psi_y, \sigma_i)$.⁸ When private agents know the new policy, the optimal simple rule sets $\bar{\pi} = 0$, $\psi_\pi = 2.4$, and $\psi_y = 0.1$. Like the model with a backward policy rule, the economy is

⁷The 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 (63) from (56),

$$D_t \tilde{\varepsilon}_t = D_t \varepsilon_t + (A_t - A_{at}) S_t.$$

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

$$\begin{aligned} S_t &= (A_t - B_t F_t)^{-1} C_t S_{t-1} + (A_t - B_t F_t)^{-1} D_t \tilde{\varepsilon}_t, \\ (A_t - B_t F_t) S_t &= C_t S_{t-1} + D_t \tilde{\varepsilon}_t, \\ &= C_t S_{t-1} + D_t \varepsilon_t + (A_t - A_{at}) S_t, \\ (A_{at} - B_t F_t) S_t &= C_t S_{t-1} + D_t \varepsilon_t, \\ S_t &= (A_{at} - B_t F_t)^{-1} C_t S_{t-1} + (A_{at} - B_t F_t)^{-1} D_t \varepsilon_t. \end{aligned} \quad (65)$$

⁸Results differ only slightly if the prior is based on estimates of the contemporaneous rule (equation 53) for the period 1966-1981.

highly fault tolerant under full information (see the left column of figure E1). There is less overshooting, however, and the cumulative output gap and sacrifice ratio are both smaller (see figure E2). Inflation again falls by 4.6 percentage points, but the cumulative output loss is just 1.3 percent, implying a sacrifice ratio of 0.28 percent.

When agents learn about the new policy, the optimized Taylor rule sets $\bar{\pi} = 0$, $\psi_{\pi} = 1.4$, and $\psi_y = 0.1$. Although the central bank reacts less aggressively to inflation than under full information, the response is quite a bit stronger than for the backward-looking rule. The bank can afford to react more aggressively because explosive dynamics vanish and the learning economy becomes highly fault tolerant (see the right-hand column of figure E1). The model therefore behaves more like its full-information counterpart than did the economy with a backward-looking rule. The learning transition is also less volatile than for the backward-looking rule, and the sacrifice ratio is smaller (compare figure E3 with figure 6 in the text). Learning is slower than for the backward-looking rule (compare figure E4 with figure 7 in the text), but that is because there is less transitional volatility in inflation and output.

From these calculations we draw two conclusions. First, many of the difficulties reported in the text follow from the fact that the central bank cannot observe current quarter output and inflation. If a rule with contemporaneous feedback to inflation and output were feasible, it would be superior. Secondly, the difference between contemporaneous and backward-looking Taylor rules is more pronounced under learning than under full information. It is also worth emphasizing that, even when there are no locally-explosive dynamics, as in the case of contemporaneous rule, the optimal simple rule under learning turns out to be less aggressive on inflation than under full information.

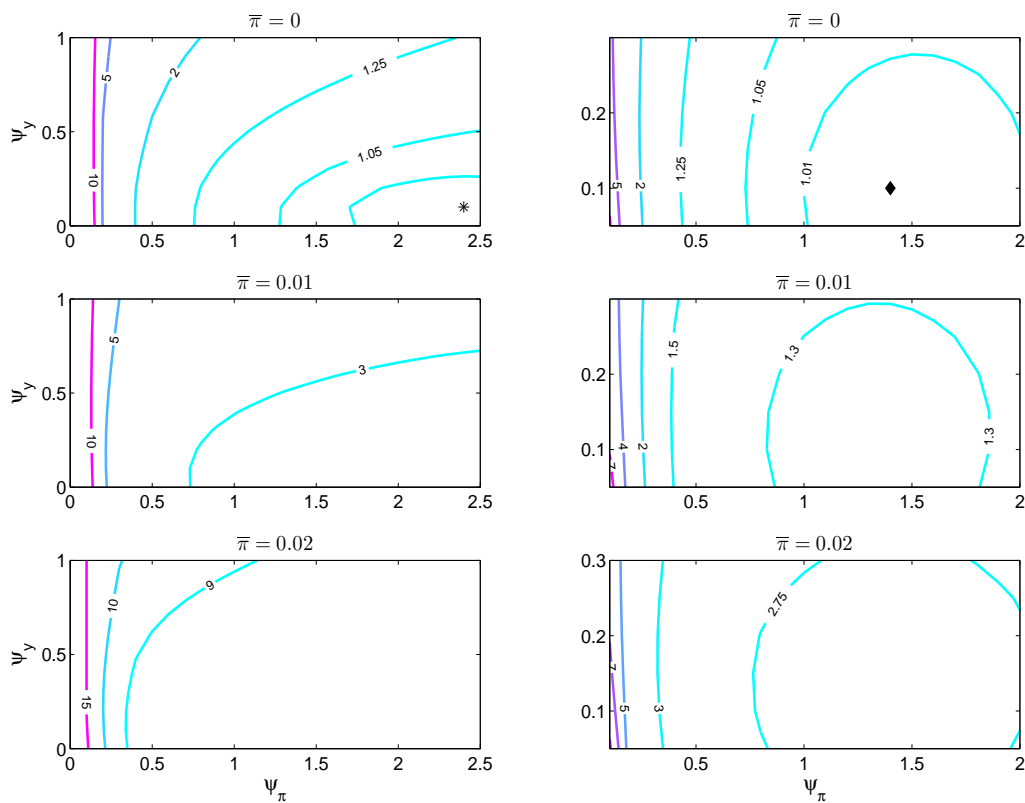


Figure E1: Isoloss contours for a contemporaneous Taylor rule. The left and right columns depict outcomes under full information and learning, respectively. The optimal rules are marked by an asterisk (full information) and diamond (learning).




Figure E2: Average responses under a contemporaneous Taylor rule optimized for full information.




Figure E3: Average responses under the contemporaneous Taylor rule optimized for learning.




Figure E4: Average estimates of policy coefficients under the contemporaneous Taylor rule optimized for learning.

F Policy shocks

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

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

Furthermore, since the initial nonexplosive region depends neither on σ_i nor on prior beliefs about σ_i , the central bank's main challenge in a $\sigma_i = 0$ economy would be the same. Hence the optimized rules are similar. When $\sigma_i = 0$ with all other aspects of the baseline economy held constant, the optimized rule sets $\bar{\pi} = 0.01$, $\psi_\pi = 0.15$, and $\psi_y = 0.1$. Thus target inflation and the response to output growth are about the same, and the response to inflation is a bit weaker. The same is true when $\sigma_i = 0$ and the prior standard deviation for σ_i is 50 percent lower than the baseline value. In both cases, isoloss contours, impulse response functions, and average estimates of policy coefficients under the optimized rule resemble those in the text (see figures F1-F5).

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




Figure F1: Average responses under the Taylor rule optimized for learning: $\sigma_i = 0$, baseline prior.




Figure F2: Average estimates of policy coefficients under the Taylor rule optimized for learning: $\sigma_i = 0$, baseline prior.




Figure F3: Average responses under the Taylor rule optimized for learning: $\sigma_i = 0$, prior standard deviation equal to half its baseline value.




Figure F4: Average estimates of policy coefficients under the Taylor rule optimized for learning: $\sigma_i = 0$, prior standard deviation equal to half its baseline value.

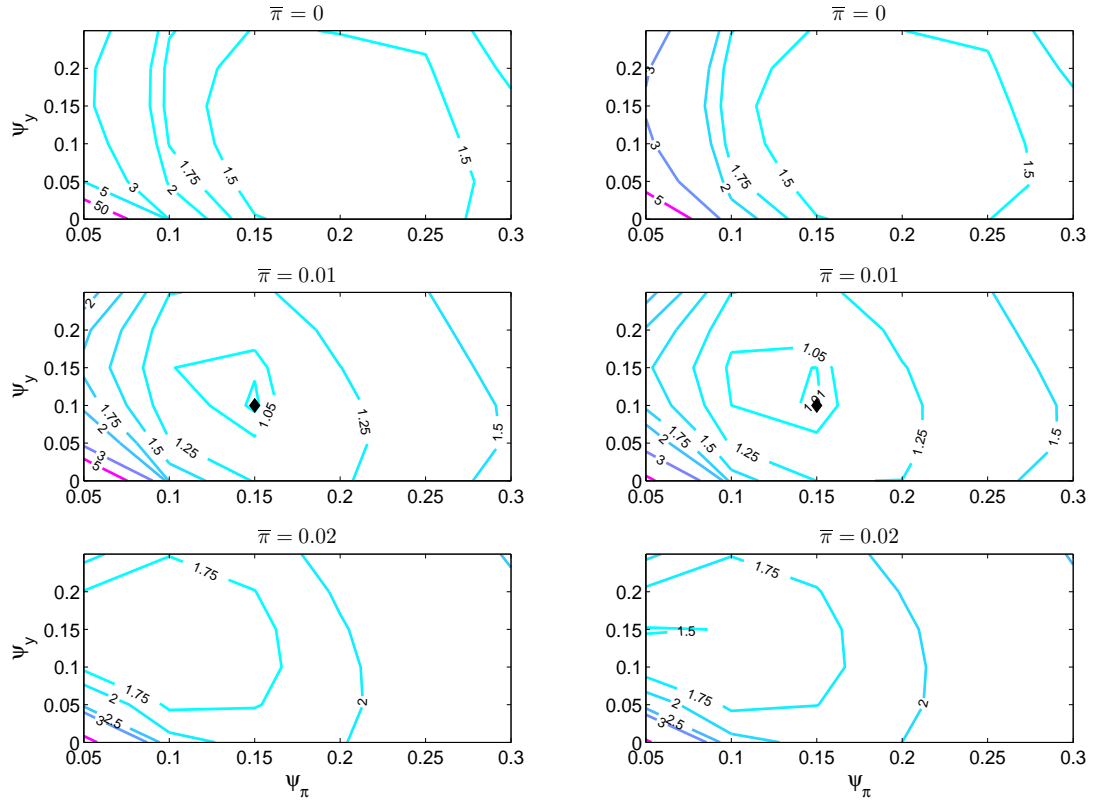


Figure F5: Iso-loss contours when $\sigma_i = 0$. The left column portrays results for the baseline prior, and the right column reduces the prior standard deviation for σ_i by half. Diamonds mark the optimized policy in each case.

G 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. Here we analyze 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.

In particular, we assume that for a period whose length is exogenous the policymaker reduces $\bar{\pi}$ but continues with response coefficients inherited from the old regime. After this initial period, when beliefs about $\bar{\pi}$ have had a chance to adjust, the policymaker adjusts the reaction coefficients. Once again, all other aspects of the baseline specification remain the same. Here we examine models in which the first stage lasts 10 and 20 quarters, respectively. Figures G1-G5 portray the results.

The two-tier approach prolongs the transition and raises expected loss. Delaying an adjustment of response coefficients postpones but does not circumvent the problem of coping with locally-explosive dynamics. The potential for explosive dynamics 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. In stage 1, beliefs about ψ_π and ψ_y harden around old-regime values because agents observe more weak responses to inflation and output growth, and this hinders learning about ψ_π and ψ_y in stage 2. Less obviously, the separation of reforms also retards learning about $\bar{\pi}$ 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 . The two-tier approach also postpones this until stage 2.

Optimized Taylor rules set $\bar{\pi} = 2$ percent per annum, $\psi_\pi = 0.15$, and $\psi_y = 0.15$ or 0.2 (see the diamonds in figure G5). Target inflation is therefore slightly higher than for simultaneous reforms, the inflation response is a bit weaker, and reaction to output growth is about the same. The transition is longer and more volatile (compares figures G1 and G3 with figure 6 in the text), learning is slower (compare figures G2 and G4 with figure 7 in the text), and expected loss is 5 times greater. Furthermore, expected loss is higher the longer is the first stage.




Figure G1: Average responses under the Taylor rule optimized for first stage lasting 10 quarters.




Figure G2: Average estimates of policy coefficients under a Taylor rule optimized for a first stage lasting 10 quarters.




Figure G3: Average responses under the Taylor rule optimized for first stage lasting 20 quarters.




Figure G4: Average estimates of policy coefficients under a Taylor rule optimized for a first stage lasting 20 quarters.

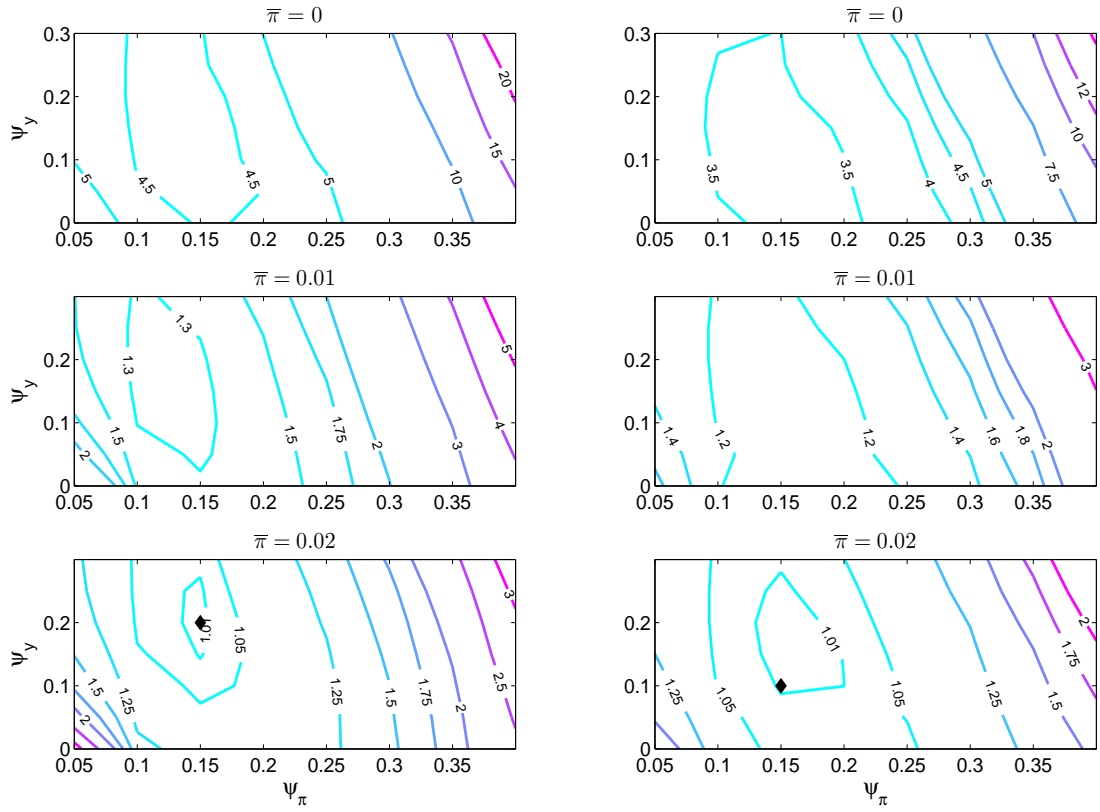


Figure G5: Iso-loss contours for two-stage disinflations. The left and right columns portray results for simulations in which the first stage lasts 10 and 20 quarters, respectively. Diamonds mark the optimal simple rule in each case.

H 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. Here we lighten their burden by assuming that agents estimate the policy rule by recursive least squares with either constant or decreasing gain. All other aspects of the baseline specification remain the same, including the prior. But we now assume that agents neglect cross-equation restrictions and work with the conditional likelihood function for the policy equation,

$$\ln p(\Delta i^t | \psi, \pi^t, \Delta y^t) = -\frac{1}{2} \sum_{j=1}^t \varphi^j \left\{ \ln \sigma_i^2 + \frac{(\Delta i_j - \psi_\pi(\pi_{j-1} - \bar{\pi}) - \psi_y \Delta y_{j-1})^2}{\sigma_i^2} \right\}. \quad (68)$$

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

Although estimates of policy coefficients sometimes differ from those in the baseline learning model, the optimal policies are essentially the same (see figure H1). Hence the choice of policy does not depend on whether private agents use single-equation or full-system estimators, nor on whether past observations are discounted.

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. Somewhat to our surprise, little is to be gained by exploiting cross-equation restrictions in the learning economy. Single-equation learning is almost as good.

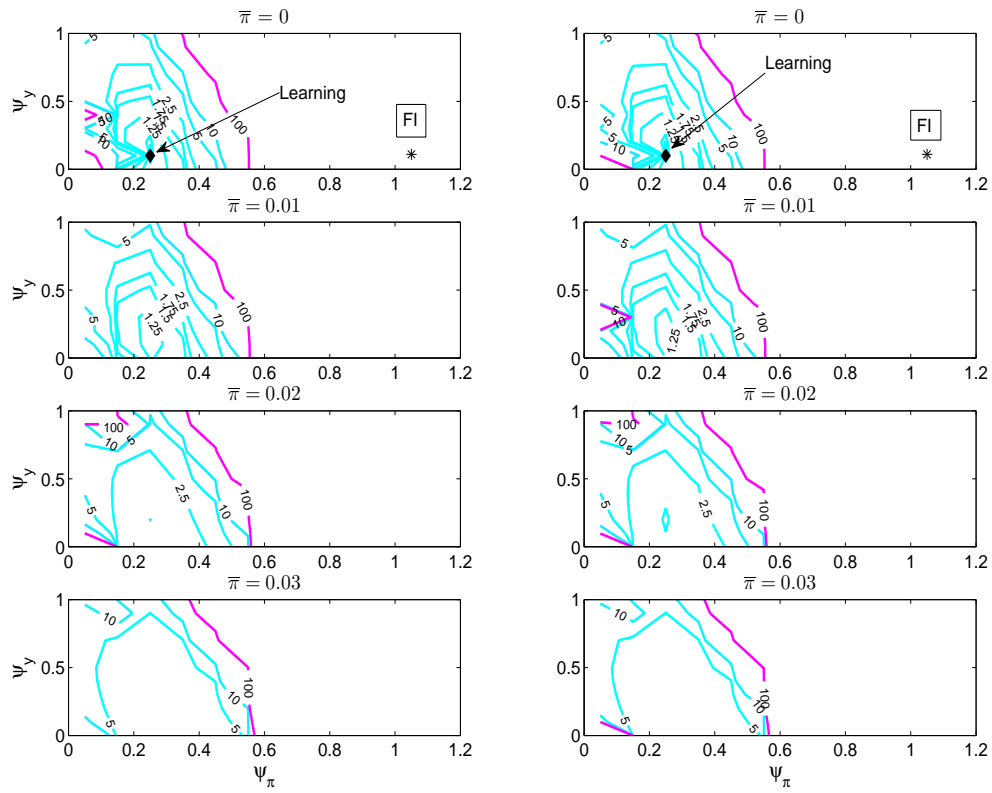


Figure H1: Iso-expected loss contours. The left and right columns refer to undiscounted and discounted least squares, respectively.

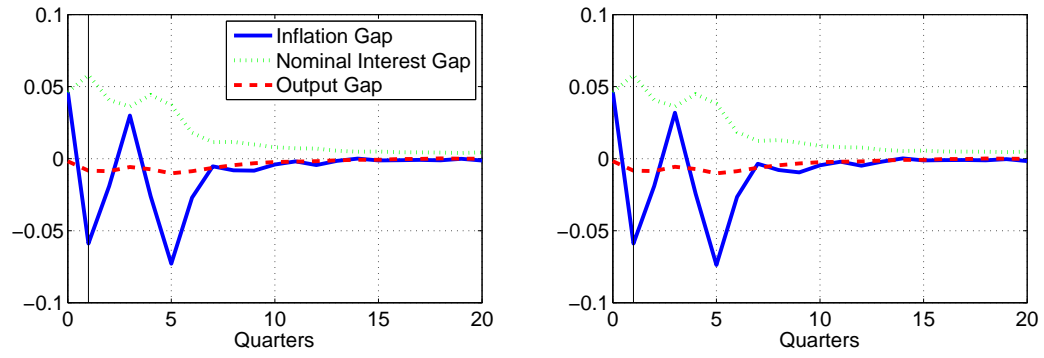


Figure H2: Average responses under the optimized rule. The left and right columns refer to undiscarded and discounted least squares, respectively.

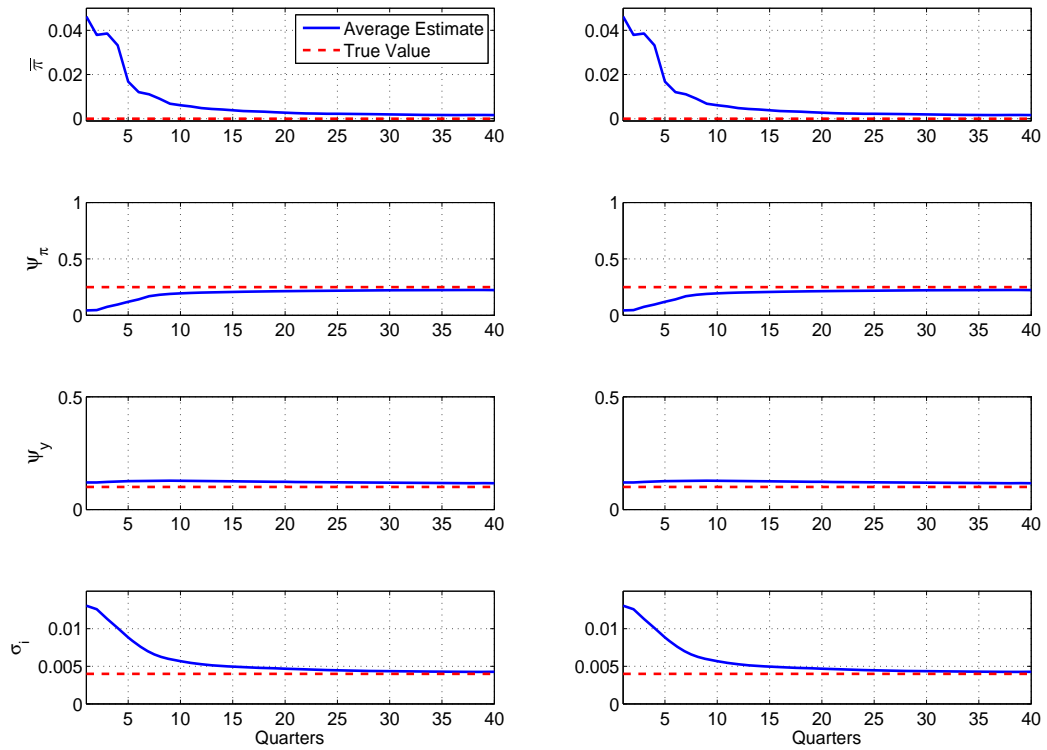


Figure H3: Average parameter estimates under the optimized rule. The left and right columns refer to undiscarded and discounted least squares, respectively.