Supplementary Material for "Monetary Policy with Supply Regimes"

Galo Nuño[1] Philipp Renner[2] Simon Scheidegger[3]

[1]IE University; CEMFI; CEPR
[2]University of Lancaster
[3]University of Lausanne; Grantham Research Institute, LSE

October 17, 2026

This Supplementary Material is organized as follows. Appendix A presents the equilibrium conditions and derives optimal monetary policy under discretion and commitment. Appendix B describes the deep equilibrium nets algorithm, its adaptation to Markov switching, and numerical accuracy diagnostics. Appendix C examines standard and regime-contingent Taylor rules. Appendix D presents the model with endogenous pricing frequency and the associated optimal discretionary policy problem. Appendix E contains additional figures.

A Optimal policy derivations

This appendix first presents the private-sector equilibrium conditions and then derives the first-order conditions for the optimal monetary policy problems under discretion and commitment introduced in Section 2.7 of the main paper.

A.1 Equilibrium conditions

The equilibrium conditions are given by:

$$
\begin{aligned}
c_t^{-\gamma} &= \lambda_t, \\
h_t^{\omega} &= w_t \lambda_t, \\
\frac{\lambda_t}{1+i_t} &= \beta E_t \left[\left(\frac{\lambda_{t+1}}{1+\pi_{t+1}} \right) \right], \\
\Xi_t^N &= y_t w_t \left(1+\tau_t\right) \left(A_t\right)^{-1} + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} \left(1+\pi_{t+1}\right)^{\epsilon} \Xi_{t+1}^N \right], \\
\Xi_t^D &= y_t + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} \left(1+\pi_{t+1}\right)^{\epsilon-1} \Xi_{t+1}^D \right], \\
p_t^* &= \mathcal{M} \frac{\Xi_t^N}{\Xi_t^D}, \\
1 &= \theta \left(1+\pi_t\right)^{\epsilon-1} + \left(1-\theta\right) \left(p_t^*\right)^{1-\epsilon}, \\
\Delta_t &= \theta \left(1+\pi_t\right)^{\epsilon} \Delta_{t-1} + \left(1-\theta\right) \left(p_t^*\right)^{-\epsilon}, \\
y_t &= c_t + g_t, \\
y_t &= A_t h_t / \Delta_t.
\end{aligned}
$$

The stochastic processes for the exogenous shocks are specified as follows:

$$
\begin{aligned}
g_t &= \bar{g} \tilde{g}_t, \\
\log\left(A_t\right) &= \left(1-\rho^A\right) \left(-\frac{\left(\sigma^A\right)^2}{2\left(1-(\rho^A)^2\right)} \right) + \rho^A \log\left(A_{t-1}\right) + \varepsilon_t^A, \\
\log\left(\tilde{g}_t\right) &= \left(1-\rho^g\right) \left(-\frac{\left(\sigma^g\right)^2}{2\left(1-(\rho^g)^2\right)} \right) + \rho^g \log\left(\tilde{g}_{t-1}\right) + \varepsilon_t^g, \\
\xi_t &= \rho^{\tau} \xi_{t-1} + \varepsilon_t^{\tau}, \\
\tau_t &= \xi_t + \eta_t.
\end{aligned}
$$

Here, $\varepsilon_t^A \sim N\left(0, (\sigma^A)^2\right), \varepsilon_t^g \sim N\left(0, (\sigma^g)^2\right)$, $\varepsilon_t^{\tau} \sim N\left(0, (\sigma^{\tau})^2\right)$, and $\varepsilon_t \sim [[1-p_{12}, p_{12}], [p_{21}, 1-p_{21}]]$ define a Markov state $s_t \in \{1,2\}$ with $\Pr(s_{t+1} = j \mid s_t = i) = p_{ij}$, $p_{11} = 1 - p_{12}$, and $p_{22} = 1 - p_{21}$, and set $\eta_t = -\bar{\eta}_1$ in state 1 and $\eta_t = \bar{\eta}_2$ in state 2.

A.2 Optimal policy under discretion

We next derive the first-order conditions for the problem under discretion. When the central bank cannot commit to future policies, the optimization problem can be expressed as follows:

$$V\left(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t\right) = \max \frac{c_t^{1-\gamma}}{1-\gamma} - \frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{1+\omega} + \beta \mathbb{E}_t\left[V\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]$$

subject to

$$\begin{aligned} p_t^*\left[(c_t+g_t) + c_t^\gamma \mathbb{E}_t\left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right] = \mathcal{M}\Bigg[&(c_t+g_t)^{1+\omega}\left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^\gamma\left(1+\tau_t\right)\left(A_t\right)^{-1} \\ &+ c_t^\gamma \mathbb{E}_t\left[G\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\Bigg], \\ 1 = \theta\left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}&, \\ \Delta_t = \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}&. \end{aligned}$$

We define

$$\begin{aligned} F\left(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t\right) &= \theta\beta c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon-1}\Xi_t^D, \\ G\left(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t\right) &= \theta\beta c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon}\Xi_t^N. \end{aligned}$$

The Lagrangian is given by:

$$\begin{aligned} \mathcal{L} = &\frac{c_t^{1-\gamma}}{1-\gamma} - \frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{1+\omega} + \beta \mathbb{E}_t\left[V\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right] \\ &+ \mu_t\Bigg(p_t^*\left[(c_t+g_t) + c_t^\gamma \mathbb{E}_t\left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right] \\ &\quad - \mathcal{M}\left[(c_t+g_t)^{1+\omega}\left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^\gamma\left(1+\tau_t\right)\left(A_t\right)^{-1} + c_t^\gamma \mathbb{E}_t\left[G\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right]\Bigg) \\ &+ \rho_t\left[-1 + \theta\left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}\right] \\ &+ \zeta_t\left[-\Delta_t + \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}\right]. \end{aligned}$$

The complete system of equations solved by the deep equilibrium net algorithm, described in Appendix B, comprises the equilibrium conditions presented in Appendix A.1 and the first-order conditions derived below.

The first-order condition with respect to consumption reads:

$$\begin{aligned}0 = {} & c_t^{-\gamma} - (c_t + g_t)^{\omega} \left(\frac{\Delta_t}{A_t}\right)^{1+\omega} \\ & + \mu_t \left[p_t^* \left(1 + \gamma c_t^{\gamma-1} \mathbb{E}_t \left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right)\right. \\ & \quad - \mathcal{M} \left(\left((1+\omega)\, c_t + \gamma\left(c_t + g_t\right)\right)\left(c_t + g_t\right)^{\omega} c_t^{\gamma-1} \left(\frac{\Delta_t}{A_t}\right)^{\omega} \left(1+\tau_t\right)\left(A_t\right)^{-1} \right. \\ & \quad \left.\left. + \ \gamma c_t^{\gamma-1} \mathbb{E}_t \left[G\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right)\right].\end{aligned}$$

The first-order condition with respect to inflation is given by:

$$0 = \rho_t \theta \left(\epsilon - 1\right)\left(1 + \pi_t\right)^{\epsilon-2} + \zeta_t \theta \epsilon \left(1 + \pi_t\right)^{\epsilon-1} \Delta_{t-1}.$$

The first-order condition with respect to the optimal price is:

$$\begin{aligned}0 = {} & \mu_t \left[\left(c_t + g_t\right) + c_t^{\gamma} \mathbb{E}_t \left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]\right] \\ & + \rho_t (1-\theta)\left(1-\epsilon\right)\left(p_t^*\right)^{-\epsilon} - \zeta_t (1-\theta)\epsilon \left(p_t^*\right)^{-\epsilon-1}.\end{aligned}$$

We substitute $\Xi_t^D = \left(c_t + g_t\right) + c_t^{\gamma} \mathbb{E}_t \left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}\right)\right]$ and obtain:

$$0 = \mu_t \Xi_t^D + \rho_t (1-\theta)\left(1-\epsilon\right)\left(p_t^*\right)^{-\epsilon} - \zeta_t (1-\theta)\epsilon \left(p_t^*\right)^{-\epsilon-1}.$$

Finally, the first-order condition with respect to price dispersion is given by

$$\begin{aligned}0 = {} & -\frac{\left(\frac{\left(c_t+g_t\right)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} + \beta \mathbb{E}_t \left[\frac{\partial V}{\partial \Delta_t}\right] - \zeta_t \\ & + \mu_t \left[p_t^* c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial F}{\partial \Delta_t}\right]\right. \\ & \quad \left. - \mathcal{M} \left(\left(c_t + g_t\right)^{1+\omega} c_t^{\gamma} \frac{\omega}{\Delta_t} \left(\frac{\Delta_t}{A_t}\right)^{\omega} \left(1+\tau_t\right)\left(A_t\right)^{-1} + c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_t}\right]\right)\right].\end{aligned} \tag{1}$$

Note that the derivatives of V, F, and G are evaluated at the next-period state, which is left out for readability.

By the envelope theorem, the following condition holds at the optimum:

$$\frac{\partial V}{\partial \Delta_{t-1}} = \frac{\partial \mathcal{L}}{\partial \Delta_{t-1}} = \theta (1 + \pi_t)^{\varepsilon} \zeta_t.$$

Shifting the envelope condition forward by one period and taking conditional expectations

gives the substitution for equation (1):

$$\mathbb{E}_t\left[\frac{\partial V}{\partial \Delta_t}\right] = \mathbb{E}_t\left[\theta(1+\pi_{t+1})^{\varepsilon}\zeta_{t+1}\right].$$

The model has the states: $(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t)$ and the output policies are $(\Xi_t^N, \Xi_t^D, c_t, \rho_t)$. All other policies are derived by reducing the above equations.

A.3 Optimal policy under commitment

In this subsection, we derive the optimal policy under commitment. The corresponding Lagrangian is given by:

$$\begin{aligned}
\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t & \left[\frac{c_t^{1-\gamma}}{1-\gamma} - \frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{1+\omega}\right] \\
& + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \mu_t \left[-p_t^* \Xi_t^D + \mathcal{M}\Xi_t^N\right] \\
& + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \nu_t \left[-1 + \theta\left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}\right] \\
& + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \zeta_t \left[-\Delta_t + \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}\right] \\
& + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \vartheta_t \left[(c_t+g_t)^{1+\omega}\left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^{\gamma}\left(1+\tau_t\right)\left(A_t\right)^{-1}\right. \\
& \quad \left. + \mathbb{E}_t\left[\beta\theta c_t^{\gamma} c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right] - \Xi_t^N\right] \\
& + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \varrho_t \left[(c_t+g_t) + \mathbb{E}_t\left[\beta\theta c_t^{\gamma} c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right] - \Xi_t^D\right].
\end{aligned}$$

We only consider the continuation conditions, i.e., we assume that we inherit some commitment from the past.

The first-order condition with respect to consumption reads:

$$\begin{aligned}0 = {} & c_t^{-\gamma} - (c_t + g_t)^{\omega} \left(\frac{\Delta_t}{A_t}\right)^{1+\omega} \\ & + \vartheta_t \left[\left((1+\omega)\, c_t + \gamma\,(c_t + g_t)\right)(c_t + g_t)^{\omega} c_t^{\gamma-1} \left(\frac{\Delta_t}{A_t}\right)^{\omega} (1+\tau_t)\,(A_t)^{-1} \right. \\ & \quad \left. + \mathbb{E}_t\left[\beta\theta\gamma c_t^{\gamma-1} c_{t+1}^{-\gamma} (1+\pi_{t+1})^{\epsilon} \Xi_{t+1}^N\right]\right] \\ & + \beta^{-1}\vartheta_{t-1}(-\gamma)\,\beta\theta c_{t-1}^{\gamma} c_t^{-\gamma-1} (1+\pi_t)^{\epsilon} \Xi_t^N \\ & + \varrho_t\left[1 + \mathbb{E}_t\left[\beta\theta\gamma c_t^{\gamma-1} c_{t+1}^{-\gamma} (1+\pi_{t+1})^{\epsilon-1} \Xi_{t+1}^D\right]\right] \\ & + \beta^{-1}\varrho_{t-1}(-\gamma)\,\beta\theta c_{t-1}^{\gamma} c_t^{-\gamma-1} (1+\pi_t)^{\epsilon-1} \Xi_t^D.\end{aligned}$$

The first-order condition with respect to price dispersion is given by:

$$\begin{aligned}0 = {} & -\frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} - \zeta_t + \beta\mathbb{E}_t\left[\zeta_{t+1}\theta\,(1+\pi_{t+1})^{\epsilon}\right] \\ & + \vartheta_t\omega\,(c_t+g_t)^{1+\omega} c_t^{\gamma} \left(\frac{\Delta_t}{A_t}\right)^{\omega} \frac{1}{\Delta_t}(1+\tau_t)\,(A_t)^{-1}.\end{aligned}$$

The first-order condition for inflation reads:

$$\begin{aligned}0 = {} & \nu_t\theta\,(\epsilon-1)\,(1+\pi_t)^{\epsilon-2} + \zeta_t\theta\epsilon\,(1+\pi_t)^{\epsilon-1}\Delta_{t-1} \\ & + \beta^{-1}\epsilon\vartheta_{t-1}\beta\theta c_{t-1}^{\gamma} c_t^{-\gamma}\,(1+\pi_t)^{\epsilon-1}\Xi_t^N \\ & + \beta^{-1}\varrho_{t-1}\beta\theta\,(\epsilon-1)\,c_{t-1}^{\gamma} c_t^{-\gamma}\,(1+\pi_t)^{\epsilon-2}\Xi_t^D.\end{aligned}$$

The first-order condition with respect to the optimal price is:

$$0 = -\mu_t\Xi_t^D + \nu_t(1-\theta)\,(1-\epsilon)\,(p_t^*)^{-\epsilon} + \zeta_t(1-\theta)\,(-\epsilon)\,(p_t^*)^{-\epsilon-1}.$$

The first-order condition with respect to Ξ_t^N is given by:

$$0 = \mu_t\mathcal{M} - \vartheta_t + \beta^{-1}\vartheta_{t-1}\beta\theta c_{t-1}^{\gamma} c_t^{-\gamma}\,(1+\pi_t)^{\epsilon}.$$

Finally, the first-order condition with respect to Ξ_t^D reads:

$$0 = -\mu_t p_t^* - \varrho_t + \beta^{-1}\varrho_{t-1}\beta\theta c_{t-1}^{\gamma} c_t^{-\gamma}\,(1+\pi_t)^{\epsilon-1}.$$

As in the discretion model, the system of equations solved by our algorithm (see Appendix B) comprises the equilibrium conditions presented in Appendix A.1, together with

the first-order conditions derived above.

The model has the states: $(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t, \vartheta_{t-1}, \varrho_{t-1}, c_{t-1})$ and the output policies are (Ξ_t^N, Ξ_t^D, c_t). All other policies are derived by reducing the above equations.

B Deep equilibrium nets algorithm

The deep equilibrium nets (DEQN) algorithm is a simulation-based solution method that uses deep neural networks to approximate the vector of policy and equilibrium functions $\mathbf{p} : X \to Y \subset \mathbb{R}^M$ of a dynamic model (Azinovic et al., 2022). The equilibrium conditions can be written as

$$\mathbf{G}(\mathbf{x}, \mathbf{p}) = \mathbf{0}, \qquad \mathbf{x} \in X \subset \mathbb{R}^d. \tag{2}$$

The unknown functions are approximated by a neural network, $\mathbf{p}(\mathbf{x}) \approx \mathcal{N}_{\boldsymbol{\nu}}(\mathbf{x})$, whose parameters $\boldsymbol{\nu}$ are trained to minimize violations of the equilibrium conditions. Evaluating these conditions generally requires the network's predictions at both the current state and possible successor states. Conditions imposed directly through substitution need not enter the training loss. Below, we describe the common algorithm, its application to each monetary policy problem, and the resulting residual distributions.

Neural network and training. We use densely connected feedforward neural networks. With $L-1$ hidden layers, the network takes the form

$$\mathcal{N}_{\boldsymbol{\nu}} = \mathcal{A}^L \circ \sigma_{L-1} \circ \mathcal{A}^{L-1} \circ \cdots \circ \sigma_1 \circ \mathcal{A}^1, \qquad \mathcal{A}^\ell(\mathbf{z}) = \mathbf{W}^\ell \mathbf{z} + \mathbf{b}^\ell, \tag{3}$$

where $\boldsymbol{\nu}$ collects the weights and biases and σ_ℓ denotes an elementwise activation function. All numerical experiments use two hidden layers with 512 neurons each and SELU activations.[1] Figure S.1 illustrates the architecture. Starting from randomly initialized parameters, the algorithm alternates between simulating the model under the current policy approximation and updating the network using the Adam optimizer (Kingma and Ba, 2014). The initial learning rate is 10^{-5}; training starts in single precision and subsequently continues in double precision.

Expectations under regime switching. Let $s_t \in \{1, 2\}$ denote the supply regime, with $\eta_t = -\bar{\eta}_1$ in normal times and $\eta_t = \bar{\eta}_2$ in bad times. The regime indicator is distinct from the value of the cost-push component η_t. Given the full current state $\mathbf{x}_t$, expectations

[1] The numbers of inputs and outputs vary across model specifications. For example, in the baseline New Keynesian model with Calvo pricing and monetary policy governed by the Taylor rule, the network can take the five state variables $(\Delta_{t-1}, A_t, g_t, \xi_t, \eta_t)$ as inputs and return three outputs: consumption c_t and the two auxiliary pricing functions Ξ_t^N and Ξ_t^D.

combine a sum over future regimes with quadrature over the continuous innovations:

$$\mathbb{E}_t[H(\mathbf{x}_{t+1})] \approx \sum_{k=1}^{2} p_{s_t k} \sum_{j=1}^{J} w_j H\left(\mathcal{T}_{\boldsymbol{\nu}}(\mathbf{x}_t, \boldsymbol{\varepsilon}^{(j)}, k)\right). \tag{4}$$

Here, $\mathcal{T}_{\boldsymbol{\nu}}$ is the model-implied transition under the current policy approximation, and $(\boldsymbol{\varepsilon}^{(j)}, w_j)$ are quadrature nodes and weights for the innovations. The innovation nodes are fixed, whereas the corresponding successor states depend on the current state and policy. The regime probabilities depend on s_t. This representation permits batched evaluation across states, quadrature nodes, and future regimes.

Training loss. For model q, let $\mathbf{R}_t^q$ collect the equations that remain after algebraic substitutions, and let $f^q_{m,\boldsymbol{\nu}}(\mathbf{x}_t)$ denote their signed residuals in the scaling used for training. Specifically, $f^q_{m,\boldsymbol{\nu}} = R^q_m / a^q_m$, where $a^q_m > 0$ is the equation-specific scale and $a^q_m = 1$ gives an unscaled residual. For a batch $\mathcal{B}$ of simulated states, the objective is

$$\ell^q_{\boldsymbol{\nu}}(\mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} \sum_{m=1}^{N_q} \left[f^q_{m,\boldsymbol{\nu}}(\mathbf{x})\right]^2. \tag{5}$$

Thus, training minimizes the average sum of squared equation residuals. The residual vector stacks the individual equations; it is not their scalar sum. This construction follows the equation-by-equation approach in Azinovic et al. (2022, Section 4) and Kübler et al. (2025, Appendix A.3).

The equations are scaled in such a manner that they have no units. That is the loss from all equations is comparable when training. More precisely we identify a part of the equation that is nonzero in equilibrium and divide the equation by that term. For the quadrature we use Gauss-Hermite with 3 quadrature nodes since we are dealing solely with the normal distribution.

Common states and pricing equations. Under Calvo pricing and without commitment, the state is

$$\mathbf{x}_t^0 = (\Delta_{t-1}, A_t, g_t, \xi_t, \eta_t).$$

It contains four continuous variables and one discrete regime variable. Keeping ξ_t and η_t separate preserves their distinct transition laws, while $\tau_t = \xi_t + \eta_t$ enters current production costs. For the linear production technology used in the quantitative analysis, the two pricing

recursions generate the residuals

$$\mathcal{C}_t^N = \frac{1}{\Xi_t^N}\left(\Xi_t^N - \frac{y_t w_t(1+\tau_t)}{A_t} - \mathbb{E}_t\left[q_{t+1}\Lambda_{t,t+1}(1+\pi_{t+1})^{\epsilon}\Xi_{t+1}^N\right]\right),$$
$$\mathcal{C}_t^D == \frac{1}{\Xi_t^D}\left(\Xi_t^D - y_t - \mathbb{E}_t\left[q_{t+1}\Lambda_{t,t+1}(1+\pi_{t+1})^{\epsilon-1}\Xi_{t+1}^D\right]\right), \tag{6}$$

where $q_{t+1} = \theta$ under Calvo pricing and $q_{t+1} = 1 - n_{t+1}$ under endogenous pricing frequency and Ξ_t^N, Ξ_t^D are the equation specific scaling factors. Expectations are evaluated using equation (4). The following reduced representations distinguish the functions to be approximated from variables recovered through the equilibrium conditions.

Taylor rules. For both the standard and regime-contingent Taylor rules in Appendix C, a reduced policy vector is

$$\mathbf{p}_t^{\mathrm{TR}} = (\Xi_t^D, \Xi_t^N, c_t).$$

The reset price follows from $p_t^* = \mathcal{M}\Xi_t^N/\Xi_t^D$, inflation from the price-index equation, and current dispersion from its law of motion. Output, labor, and wages then follow from market clearing and household labor supply. The policy rule determines i_t. The three remaining residuals are

$$\mathbf{R}_t^{\mathrm{TR}} = (\mathcal{E}_t, \mathcal{C}_t^N, \mathcal{C}_t^D), \qquad \mathcal{E}_t = c_t^{-\gamma} - \beta\left(\frac{(1+\bar{\pi})}{\beta} + \psi\left(\pi_t - \bar{\pi}\right)\right)\mathbb{E}_t\left[\frac{c_{t+1}^{-\gamma}}{1+\pi_{t+1}}\right].$$

Discretion with Calvo pricing. The state remains $\mathbf{x}_t^0$. A reduced policy vector augments the Taylor-rule representation with the multiplier on the price-index constraint:

$$\mathbf{p}_t^{\mathrm{D}} = (\Xi_t^D, \Xi_t^N, c_t, \rho_t).$$

The inflation and reset-price optimality conditions recover the remaining current multipliers. The residual system consists of the two pricing recursions and the consumption and price-dispersion optimality conditions from Appendix A.2:

$$\mathbf{R}_t^{\mathrm{D}} = \left(\mathcal{L}_c^{\mathrm{D}}, \mathcal{L}_\Delta^{\mathrm{D}}, \mathcal{C}_t^N, \mathcal{C}_t^D\right).$$

The dispersion condition includes the effect of current policy on continuation policies. The envelope condition eliminates the derivative of the value function, and the household Euler equation determines the implementing nominal interest rate.

Commitment with Calvo pricing. Commitment adds three inherited states associated with the forward-looking pricing equations. Defining ϑ_t, c_t and ϱ_t gives the eight-component

state

$$\mathbf{x}_t^{\mathrm{C}} = (\mathbf{x}_t^0, \vartheta_t, c_t, \varrho_t).$$

A reduced policy vector is $\mathbf{p}_t^{\mathrm{C}} = (\Xi_t^D, \Xi_t^N, c_t, \nu_t)$, where ν_t is the price-index multiplier in Appendix A.3. The other multiplier conditions can be substituted out, leaving

$$\mathbf{R}_t^{\mathrm{C}} = \left(\mathcal{L}_c^{\mathrm{C}}, \mathcal{L}_\Delta^{\mathrm{C}}, \mathcal{C}_t^N, \mathcal{C}_t^D\right).$$

The current multipliers become next period's inherited states, as does current consumption. They are initialized randomly at first and then via simulation results.

Discretion with endogenous pricing frequency. With the within-period timing in the main paper, n_t is a current choice. The physical state remains $\mathbf{x}_t^0$. The reduced policy vector is

$$\mathbf{p}_t^{\mathrm{EF}} = (\Xi_t^D, \Xi_t^N, c_t, n_t, \rho_t, \chi_t),$$

where χ_t denotes the multiplier on the firm's pricing-frequency condition. Relative to discretionary Calvo pricing, the residual system adds the firm's frequency condition, denoted $\mathcal{P}_t$, and the corresponding monetary-policy optimality condition:

$$\mathbf{R}_t^{\mathrm{EF}} = \left(\mathcal{L}_\Delta^{\mathrm{EF}}, \mathcal{L}_c^{\mathrm{EF}}, \mathcal{L}_n^{\mathrm{EF}}, \mathcal{C}_t^N, \mathcal{C}_t^D, \mathcal{P}_t\right).$$

The model and private pricing conditions are given in Appendix D. When solving jointly over different shock magnitudes, we append $\bar{\eta}$ as a fixed parameter input, with $\bar{\eta}_{t+1} = \bar{\eta}_t$.

Numerical accuracy. Table S.1 reports equation-level residual distributions from a simulation of 10,000 periods with 100 period burn in starting from a random state in the training set. Following the reporting approach of Azinovic et al. (2022) and Kübler et al. (2025), we distinguish average accuracy from the upper tail. For each equation, the table reports the mean, 95th percentile, and maximum of the squared residual f_m^2. Equations satisfied through algebraic substitution are omitted. The mean squared residuals are the equation-level contributions to the objective. The equations were normalized and so all residuals are unit free and therefore comparable.

Across the reported equations, mean residuals range from 6.4×10^{-5} to 5.2×10^{-8}. The largest reported 95th percentile is 1.9×10^{-4}, and the largest maximum is 4.0×10^{-3}, both in the endogenous-frequency model. Note that the endogenous frequency model has non-differentiable residuals due to the max operator and so higher residuals are to be expected. These statistics describe accuracy on the simulated states.

The residual magnitudes shown here are broadly comparable to those reported in recent deep-learning solutions of nonlinear stochastic equilibrium models (see, e.g., Azinovic et al.,

Equation	Mean	95th percentile	Maximum
A. Taylor rule, Calvo pricing			
Euler equation [%]	5.2×10^{-8}	2.4×10^{-7}	5.7×10^{-7}
Pricing numerator [%]	1.0×10^{-6}	2.5×10^{-6}	4.1×10^{-6}
Pricing denominator [%]	2.0×10^{-7}	4.5×10^{-7}	1.9×10^{-6}
B. Discretion, Calvo pricing			
$\mathcal{L}_c$ [%]	3.5×10^{-6}	1.3×10^{-5}	8.4×10^{-5}
$\mathcal{L}_\Delta$ [%]	4.8×10^{-6}	2.0×10^{-5}	1.1×10^{-4}
Pricing numerator [%]	5.5×10^{-6}	3.0×10^{-5}	1.5×10^{-4}
Pricing denominator [%]	5.6×10^{-6}	3.0×10^{-5}	1.3×10^{-4}
C. Commitment, Calvo pricing			
$\mathcal{L}_c$ [%]	4.2×10^{-6}	7.6×10^{-6}	3.0×10^{-5}
$\mathcal{L}_\Delta$ [%]	1.8×10^{-6}	4.5×10^{-6}	1.2×10^{-5}
Pricing numerator [%]	2.0×10^{-7}	6.1×10^{-7}	4.4×10^{-6}
Pricing denominator [%]	6.7×10^{-8}	2.2×10^{-7}	2.7×10^{-6}
D. Discretion, endogenous pricing frequency			
$\mathcal{L}_\Delta$ [%]	3.4×10^{-5}	3.7×10^{-5}	3.5×10^{-3}
$\mathcal{L}_c$ [%]	6.5×10^{-6}	1.5×10^{-5}	6.6×10^{-4}
$\mathcal{L}_n$ [%]	1.3×10^{-5}	3.7×10^{-5}	7.7×10^{-4}
Pricing numerator [%]	3.0×10^{-6}	1.2×10^{-5}	1.6×10^{-4}
Pricing denominator [%]	7.5×10^{-6}	2.2×10^{-5}	4.6×10^{-4}
Pricing frequency [%]	6.4×10^{-5}	1.9×10^{-4}	4.0×10^{-3}

Table S.1: Squared equation-level residual distributions across monetary policy models.

Note: Statistics are based on a simulation of 10,000 periods. All numerical entries are reported as percentages of the corresponding implemented squared residual measure.[1]

[1]The displayed "Maximum" denotes the 99.9th percentile of the corresponding squared residual distribution, following the reporting convention in Brumm and Scheidegger (2017) and references therein. This measure reduces sensitivity to isolated extreme observations and to the random seed used to generate the evaluation paths.

2022, Kübler et al., 2025, Kase et al., 2025), allowing for differences in residual definitions and normalization across applications. The small average and upper-tail residuals indicate that the equilibrium and optimality conditions are closely satisfied over the simulated region of the state space, supporting the numerical reliability of our policy comparisons. Furthermore, our numerical experiments show that reducing numerical errors, for example by using double precision, leaves the computed policy functions and the resulting economic conclusions essentially unchanged.

C Monetary policy rules

Taylor rule. We consider a Taylor rule of the form

$$i_t = \frac{(1+\bar{\pi})}{\beta} - 1 + \psi\left(\pi_t - \bar{\pi}\right), \tag{7}$$

where $\bar{\pi} = 0$ is the inflation target and $\psi = 2$ is the policy response coefficient to deviations of inflation from its target.

Figure S.5 presents the corresponding ergodic distributions. While the results preserve the bimodality of the output gap and the real interest rate observed under optimal policy with commitment, inflation is no longer centered around zero in each regime (panel a). In the normal-times regime (blue bars), inflation is negative (centered around -0.3 percent), whereas in bad times (orange bars) it is positive (centered around 0.3 percent). These deviations from the central bank's target $\bar{\pi} = 0$ arise because the Taylor rule does not track the appropriate natural rate of interest. Evaluating the rule in the stochastic steady state of regime $n \in \{1, 2\}$ yields

$$i_{n,ss} \simeq \left(\frac{1}{\beta} - 1\right) + \bar{\pi} + \psi\left(\pi_{n,ss} - \bar{\pi}\right) = \bar{r} + (1-\psi)\,\bar{\pi} + \psi\pi_{n,ss},$$

where $\bar{r} = \frac{1}{\beta} - 1$, corresponding to approximately 1 percent on an annualized basis, is the central bank's implicit real-rate target, which under the Taylor rule (7) coincides with the real rate in the non-stochastic steady state of the efficient allocation, $\bar{r} = \hat{r}$. Substituting the approximate Fisher equation, $i_{n,ss} \simeq r_{n,ss} + \pi_{n,ss}$, gives

$$\pi_{n,ss} \simeq \bar{\pi} + \frac{r_{n,ss} - \bar{r}}{\psi - 1}. \tag{8}$$

Equation (8) shows how long-run inflation deviates from the central bank's target when the policy rule anchors the nominal rate to an incorrect long-term real rate. This relationship was also obtained by Campos et al. (2024) in the context of HANK models, where the natural

rate depends endogenously on fiscal policy. In contrast, in our model the regime-switching cost-push shock drives shifts in the natural rate. As a result, a sizable gap emerges between the natural rate in each regime and the central bank's implicit real-rate target. In normal times, the bank uses a real-rate intercept that is too high, leading to excessively tight policy and persistently low inflation ($\pi_{1,ss} \simeq 0.7\% - 1\% = -0.3\%$, where 0.7% is the mean natural rate in normal times).[2] In bad times, by contrast, the intercept is too low, generating inflation above target ($\pi_{2,ss} \simeq 1.3\% - 1\% = 0.3\%$, where 1.3% is the mean natural rate in bad times).

Although the central bank fails to stabilize inflation across regimes, it achieves its price stability objective approximately on average. The ergodic mean of inflation is approximately 0 percent, and the average real interest rate is approximately 1 percent, consistent with equation (8) on average ($1\% - 1\% = 0\%$).

These results offer a possible rationale for the joint behavior of inflation and long-term interest rates before and after the COVID-19 pandemic (Benigno et al., 2024). Before the pandemic, long-term inflation expectations in advanced economies, such as the euro area, were persistently below target, while estimates of the natural rate were near zero or negative. Since the pandemic, both inflation expectations and natural-rate estimates have risen sharply, in line with adverse supply shocks such as the wars in Ukraine and Iran and the ongoing process of de-globalization.

Modified Taylor rule. We now consider an alternative, regime-contingent monetary policy rule. The modified Taylor rule is given by

$$i_t = \bar{r}_t + \psi \left(\pi_t - \bar{\pi} \right), \tag{9}$$

where the intercept of the rule equals the natural rate in each regime,

$$\bar{r}_t = r_n^*, \quad \text{if the regime at time } t \text{ is } n.$$

Here, r_n^* denotes the natural rate in regime n, defined as the SSS real interest rate in the counterfactual flexible-price allocation. The modified rule requires the central bank to infer the prevailing supply regime and its associated natural rate. In the empirical interpretation used for our calibration, this information can be tied to observable, persistent external cost indicators such as the real oil price.

In this case, it is straightforward to verify that equation (8) is consistent with a zero inflation target, $\bar{\pi} = 0$. The red and green bars in Figure S.5 show the ergodic distributions under this modified Taylor rule, with the inflation distribution centered around zero. The

[2] Notice that we approximating the equilibrium real rate under the Taylor rule in Equation (8), $r_{n,ss}$ with the flexible-price natural rate, r_n^*.

within-regime variances of the main macroeconomic variables are similar under the original and modified Taylor rules.

The policy implication is clear: in the presence of supply regimes, the central bank should adjust the intercept of its interest-rate rule to track the natural rate, which itself becomes a regime-contingent variable.

Sensitivity to regime length. Finally, Figure S.6 shows the sensitivity of SSS inflation and real interest rates to alternative expected durations of the bad-times regime under the Taylor rule (7). The solid blue line corresponds to the normal-times SSS in the baseline model. As discussed earlier, the parameter p_{21} governs the probability of transitioning from regime 2 (bad times) to regime 1 (normal times), and thus the expected duration of bad times is $1/p_{21}$. In the baseline calibration, this duration is 52 quarters.

The SSS values of inflation and the real interest rate exhibit strong nonlinearity with respect to changes in the expected duration of the bad-times regime. When the duration is reduced, making the regime less persistent, the SSS values of inflation and real rates in normal times converge to those of the efficient allocation (zero inflation and a 1 percent real rate). This occurs because, as $p_{21} \to 1$, agents perceive regime changes as equivalent to one-period i.i.d. shocks, and the SSS therefore approximately converges to that of a model without regimes. Conversely, when the duration increases, the values of inflation and real rates gradually approach those in the baseline calibration.

The dashed red line in Figure S.6 depicts the bad-times SSS in a counterfactual scenario where $p_{12} \to 1$, that is, where the expected duration of the normal-times regime is one period. In this case, convergence to the SSS of the efficient allocation occurs when the duration of the bad-times regime *increases*. The logic is thus reversed relative to the previous case: the economy spends most of its time in the bad-times regime but experiences only brief transitions to the normal-times regime, corresponding to positive i.i.d. cost-push shocks with temporarily reduced markups.

D The economy with endogenous pricing frequency

This appendix presents the full analytical derivation of the economy with endogenous pricing frequency, following Blanco et al. (2024). It details the model's structure, equilibrium conditions, and the characterization of optimal monetary policy under discretion.

D.1 Model

We now describe the structure of the economy under endogenous pricing frequency.

Technology. There is a continuum of monopolistically competitive firms, each producing a differentiated variety indexed by j. Each firm produces a continuum of products k according to the production technology

$$y_{jkt} = A_t \left(l_{jkt} \right),$$

where A_t denotes stochastic total factor productivity. Firms hire labor to produce output. A perfectly competitive sector aggregates the intermediate goods y_{jkt} into a composite final good:

$$y_t = \left[\int_0^1 \int_0^1 \left(y_{jkt} \right)^{\frac{\epsilon-1}{\epsilon}} dk dj \right]^{\frac{\epsilon}{\epsilon-1}},$$

where ϵ is the elasticity of substitution across varieties. The corresponding demand for each individual product is given by

$$y_{jkt} = \left(\frac{P_{jkt}}{P_t} \right)^{-\epsilon} y_t, \tag{10}$$

where P_{jkt} is the price of product k in variety j, and the aggregate price index is defined as

$$P_t = \left[\int_0^1 \int_0^1 \left(P_{jkt} \right)^{1-\epsilon} dk dj \right]^{\frac{1}{1-\epsilon}}. \tag{11}$$

Profits. Nominal operating profits are given by

$$P_{jkt} y_{jkt} - \left(1 + \tau_t \right) P_t w_t l_{jkt},$$

where τ_t denotes the cost-push shock (including the temporary and the Markov-switching elements, as in the main model). Real operating profits are therefore

$$\left(\frac{P_{jkt}}{P_t} \right)^{1-\epsilon} y_t - \left(1 + \tau_t \right) w_t \left(\frac{P_{jkt}}{P_t} \right)^{-\epsilon} \left(\frac{y_t}{A_t} \right).$$

We define firm-level composite output as

$$y_{jt} = \left[\int_0^1 \left(y_{jkt} \right)^{\frac{\epsilon-1}{\epsilon}} dk \right]^{\frac{\epsilon}{\epsilon-1}},$$

and total production labor as

$$l_{jt} = \int_0^1 l_{jkt} \, dk.$$

The firm-level production function can then be expressed as

$$y_{jt} = \left(\frac{X_{jt}}{P_{jt}}\right)^{\epsilon} A_t\left(l_{jt}\right),$$

where

$$P_{jt} = \left[\int_0^1 \left(P_{jkt}\right)^{1-\epsilon} dk\right]^{\frac{1}{1-\epsilon}},$$

and

$$X_{jt} = \left[\int_0^1 \left(P_{jkt}\right)^{-\epsilon} dk\right]^{-1/\epsilon}.$$

Price adjustment cost. The key assumption is that, although the firm can choose how many prices to adjust in a given period, it cannot choose which specific prices to change. We assume that the firm faces a (virtual) convex cost of price adjustment, denominated in units of labor:[3]

$$\frac{\xi}{2}\left(\max\{0, n_{jt} - \bar{n}\}\right)^2 w_t.$$

Price setting.

The value of the firm equals the present discounted sum of its flow profits:

$$\mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s} \left\{ \int_0^1 \left[\left(\frac{P_{jkt+s}}{P_{t+s}}\right)^{1-\epsilon} y_{t+s} - \left(1+\tau_{t+s}\right) w_{t+s} \left(\frac{P_{jkt+s}}{P_{t+s}}\right)^{-\epsilon} \left(\frac{y_{t+s}}{A_{t+s}}\right) \right] dk \right. \\ \left. - \frac{\xi}{2}\left(\max\{0, n_{jt+s} - \bar{n}\}\right)^2 w_{t+s} \right\}, \tag{12}$$

where $\Lambda_{t,t+s} \equiv \beta^s \frac{\lambda_{t+s}}{\lambda_t}$ is the stochastic discount factor between periods t and $t+s$. The firm chooses the fraction of prices n_{jt} to reset in each period and the corresponding reset price P_{jt}^*. Because all products share the same demand and cost conditions, $P_{jkt} = P_{jt}^*$ for all products whose prices are reset, while the remaining prices stay constant. Using the relationships

$$\int_0^1 \left(P_{jkt}\right)^{-\epsilon} dk = \left(X_{jt}\right)^{-\epsilon}, \quad \text{and} \quad \int_0^1 \left(P_{jkt}\right)^{1-\epsilon} dk = \left(P_{jt}\right)^{1-\epsilon},$$

we can express the objective, including adjustment costs, as

$$\mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s} \left\{ \left[\left(\frac{P_{jt+s}}{P_{t+s}}\right)^{1-\epsilon} y_{t+s} - \left(1+\tau_{t+s}\right) w_{t+s} \left(\frac{X_{jt+s}}{P_{t+s}}\right)^{-\epsilon} \left(\frac{y_{t+s}}{A_{t+s}}\right) \right] \right. \\ \left. - \frac{\xi}{2}\left(\max\{0, n_{jt+s} - \bar{n}\}\right)^2 w_{t+s} \right\}.$$

[3]The virtual cost is rebated back to households.

Notice that

$$\frac{\partial \left(P_{jt+s}\right)^{1-\epsilon}}{\partial P_{jt}^{*}} = (1-\epsilon)\left(P_{jt}^{*}\right)^{-\epsilon}\prod_{i=1}^{s}\left(1-n_{jt+i}\right)n_{jt},$$

and

$$\frac{\partial \left(X_{jt+s}\right)^{-\epsilon}}{\partial P_{jt}^{*}} = -\epsilon\left(P_{jt}^{*}\right)^{-\epsilon-1}\prod_{i=1}^{s}\left(1-n_{jt+i}\right)n_{jt}.$$

The optimality condition is

$$\begin{aligned}0 = \mathbb{E}_t\sum_{s=0}^{\infty}\prod_{i=1}^{s}\left(1-n_{jt+i}\right)n_{jt}\Lambda_{t,t+s}y_{t+s}\Big[&\left(P_{jt}^{*}\right)^{-\epsilon}\left(P_{t+s}\right)^{\epsilon-1}\\ &-\mathcal{M}\left(1+\tau_{t+s}\right)\frac{w_{t+s}}{y_{t+s}}\left(\frac{y_{t+s}}{A_{t+s}}\right)\left(P_{jt}^{*}\right)^{-\epsilon-1}\left(P_{t+s}\right)^{\epsilon}\Big],\end{aligned}$$

where $\mathcal{M} \equiv \frac{\epsilon}{\epsilon-1}$. The optimal reset price is therefore given by

$$\left(P_{jt}^{*}\right) = \mathcal{M}\frac{\mathbb{E}_t\sum_{s=0}^{\infty}\prod_{i=1}^{s}\left(1-n_{jt+i}\right)\Lambda_{t,t+s}\left(1+\tau_{t+s}\right)w_{t+s}\left(\frac{y_{t+s}}{A_{t+s}}\right)\left(P_{t+s}\right)^{\epsilon}}{\mathbb{E}_t\sum_{s=0}^{\infty}\prod_{i=1}^{s}\left(1-n_{jt+i}\right)\Lambda_{t,t+s}y_{t+s}\left[\left(P_{t+s}\right)^{\epsilon-1}\right]},$$

and dividing both sides by P_t, we obtain

$$\left(\frac{P_{jt}^{*}}{P_t}\right) = \mathcal{M}\frac{\Xi_{jt}^{N}}{\Xi_{jt}^{D}},$$

where

$$\begin{aligned}\Xi_{jt}^{N} &= w_t\left(1+\tau_t\right)\left(\frac{y_t}{A_t}\right) + \mathbb{E}_t\left[\left(1-n_{jt+1}\right)\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{jt+1}^{N}\right],\\ \Xi_{jt}^{D} &= y_t + \mathbb{E}_t\left[\left(1-n_{jt+1}\right)\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{jt+1}^{D}\right].\end{aligned}$$

Similarly, using

$$\frac{\partial \left(P_{jt+s}\right)^{1-\epsilon}}{\partial n_{jt}} = \left[\left(P_{jt}^{*}\right)^{1-\epsilon} - \left(P_{jt-1}\right)^{1-\epsilon}\right]\prod_{i=1}^{s}\left(1-n_{jt+i}\right),$$

and

$$\frac{\partial \left(X_{jt+s}\right)^{-\epsilon}}{\partial n_{jt}} = \left[\left(P_{jt}^{*}\right)^{-\epsilon} - \left(X_{jt-1}\right)^{-\epsilon}\right]\prod_{i=1}^{s}\left(1-n_{jt+i}\right),$$

the optimal fraction of prices to reset satisfies

$$\left[\left(\frac{P_{jt}^*}{P_t}\right)^{1-\epsilon} - \left(\frac{P_{jt-1}}{P_t}\right)^{1-\epsilon}\right] \mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s} \prod_{i=1}^{s} (1 - n_{jt+i}) \left(\frac{P_{t+s}}{P_t}\right)^{\epsilon-1} y_{t+s}$$
$$- \left[\left(\frac{P_{jt}^*}{P_t}\right)^{-\epsilon} - \left(\frac{X_{jt-1}}{P_t}\right)^{-\epsilon}\right] \mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s} \prod_{i=1}^{s} (1 - n_{jt+i}) (1 + \tau_{t+s}) w_{t+s} \left(\frac{P_{t+s}}{P_t}\right)^{\epsilon} \left(\frac{y_{t+s}}{A_{t+s}}\right)$$
$$= \xi w_t \left(\max\{0, n_{jt} - \bar{n}\}\right),$$

or equivalently,

$$\left[\left(\frac{P_{jt}^*}{P_t}\right)^{1-\epsilon} - \left(\frac{P_{jt-1}}{P_t}\right)^{1-\epsilon}\right] \Xi_{jt}^D - \left[\left(\frac{P_{jt}^*}{P_t}\right)^{-\epsilon} - \left(\frac{X_{jt-1}}{P_t}\right)^{-\epsilon}\right] \Xi_{jt}^N = \xi w_t \left(\max\{0, n_{jt} - \bar{n}\}\right).$$

Market clearing conditions.

The market-clearing condition for goods is

$$y_t = c_t + g_t,$$

where g_t denotes government spending. We assume a symmetric equilibrium in which all firms are identical, such that $n_{jt} = n_t$ and $P_{jt}^* = P_t^*$. Within each firm, product prices are selected randomly for adjustment. Thus, a fraction $1 - n_t$ of product prices remains unchanged, while the remaining fraction n_t is reset to P_t^*:

$$P_t = \left\{(1 - n_t) (P_{t-1})^{1-\epsilon} + n_t (P_t^*)^{1-\epsilon}\right\}^{\frac{1}{1-\epsilon}},$$

which implies

$$1 = (1 - n_t) (1 + \pi_t)^{\epsilon-1} + n_t (p_t^*)^{1-\epsilon}, \tag{13}$$

where $p_t^* \equiv \frac{P_t^*}{P_t}$.

The market-clearing condition for labor is

$$\begin{aligned} h_t &= \int_0^1 l_{jt} dj = \int_0^1 \int_0^1 l_{jkt} dk dj = \int_0^1 \int_0^1 \left(\frac{y_{jkt}}{A_t}\right) dk dj \\ &= (P_t)^{\epsilon} \left(\frac{y_t}{A_t}\right) \int_0^1 \int_0^1 (P_{jkt})^{-\epsilon} dk dj \\ &= (P_t)^{\epsilon} \left(\frac{y_t}{A_t}\right) \int_0^1 (X_{jt})^{-\epsilon} dj = \left(\frac{y_t}{A_t}\right) \int_0^1 \left(\frac{X_{jt}}{P_t}\right)^{-\epsilon} dj, \end{aligned}$$

so that the aggregate production function is

$$y_t = A_t h_t \left(\int_0^1 \left(\frac{X_{jt}}{P_t} \right)^{-\epsilon} dj \right)^{-1}.$$

Since firms are identical in equilibrium, we define price dispersion as

$$\Delta_t \equiv \left(\int_0^1 \left(\frac{X_{jt}}{P_t} \right)^{-\epsilon} dj \right) = \left(\frac{X_{jt}}{P_t} \right)^{-\epsilon}.$$

It follows that

$$\Delta_t = \int_0^1 \int_0^1 \left(\frac{P_{jkt}}{P_t} \right)^{-\epsilon} dk dj = (1 - n_t) (1 + \pi_t)^{\epsilon} \Delta_{t-1} + n_t (p_t^*)^{-\epsilon},$$

and therefore,

$$y_t = A_t h_t (\Delta_t)^{-1}.$$

D.2 Complete set of equations

The model's equilibrium conditions are collected below. The exogenous shock processes are those specified in Section 2 of the main paper.

$$c_t^{-\gamma} = h_t^{\omega}/w_t, \tag{14}$$

$$c_t^{-\gamma} = \beta E_t \left[\left(\frac{1+i_t}{1+\pi_{t+1}} \right) c_{t+1}^{-\gamma} \right], \tag{15}$$

$$(p_t^*) = \mathcal{M} \frac{\Xi_t^N}{\Xi_t^D}, \tag{16}$$

$$\Xi_t^N = w_t (1+\tau_t) \left(\frac{y_t}{A_t} \right) + \mathbb{E}_t \left[(1-n_{t+1}) \Lambda_{t,t+1} (1+\pi_{t+1})^{\epsilon} \Xi_{t+1}^N \right], \tag{17}$$

$$\Xi_t^D = y_t + \mathbb{E}_t \left[(1-n_{t+1}) \Lambda_{t,t+1} (1+\pi_{t+1})^{\epsilon-1} \Xi_{t+1}^D \right], \tag{18}$$

$$y_t = c_t + g_t, \tag{19}$$

$$1 = (1-n_t)(1+\pi_t)^{\epsilon-1} + n_t (p_t^*)^{1-\epsilon}, \tag{20}$$

$$\Delta_t = (1-n_t)(1+\pi_t)^{\epsilon} \Delta_{t-1} + n_t (p_t^*)^{-\epsilon}, \tag{21}$$

$$y_t = A_t h_t (\Delta_t)^{-1}, \tag{22}$$

$$w_t \max\{0, n_t - \bar{n}\} = \frac{1}{\xi} \left\{ \left[(p_t^*)^{1-\epsilon} - (1+\pi_t)^{\epsilon-1} \right] \Xi_t^D \right. \tag{23}$$

$$\left. - \left[(p_t^*)^{-\epsilon} - (\Delta_{t-1})(1+\pi_t)^{\epsilon} \right] \Xi_t^N \right\}. \tag{24}$$

Notice that the baseline New Keynesian model (Section 2 of the main paper) is recovered when $\xi \to \infty$ because the last equation (equation (24)) disappears.

D.3 Optimal discretionary policy

Central bank problem.

The central bank maximizes household welfare under discretion. The central bank's value function $V(\Delta, A, \tau, g, \eta)$ satisfies the Bellman equation

$$V(\Delta_{t-1}, A_t, \tau_t, g_t, \eta_t) = \max_{c_t, h_t, w_t, \pi_t, p_t^*, \Delta_t} \frac{c_t^{1-\gamma}}{1-\gamma} - \frac{h_t^{1+\omega}}{1+\omega} + \beta \mathbb{E}_t \left[V(\Delta_t, A_{t+1}, \tau_{t+1}, g_{t+1}, \eta_{t+1}) \right]$$

subject to

$$
\begin{aligned}
c_t^{-\gamma} &= h_t^{\omega}/w_t, \\
p_t^* &= \frac{\mathcal{M}\left[w_t\left(1+\tau_t\right)\left(\frac{y_t}{A_t}\right)+\mathbb{E}_t\left[\left(1-n_{t+1}\right)\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right]\right]}{\left(y_t+\mathbb{E}_t\left[\left(1-n_{t+1}\right)\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right]\right)}, \\
y_t &= c_t+g_t, \\
1 &= \left(1-n_t\right)\left(1+\pi_t\right)^{\epsilon-1}+n_t\left(p_t^*\right)^{1-\epsilon}, \\
\Delta_t &= \left(1-n_t\right)\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1}+n_t\left(p_t^*\right)^{-\epsilon}, \\
y_t &= A_t h_t\left(\Delta_t\right)^{-1}, \\
w_t\max\{0,n_t-\bar{n}\} &= \frac{1}{\xi}\left\{\left[\left(p_t^*\right)^{1-\epsilon}-\left(1+\pi_t\right)^{\epsilon-1}\right]\Xi_t^D-\left[\left(p_t^*\right)^{-\epsilon}-\left(\Delta_{t-1}\right)\left(1+\pi_t\right)^{\epsilon}\right]\Xi_t^N\right\}.
\end{aligned}
$$

For notational convenience, define

$$
\begin{aligned}
F\left(\Delta_{t-1},A_t,\tau_t,g_t,\eta_t\right) &\equiv \left(1-n_t\right)\beta c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon-1}\Xi_t^D, \\
G\left(\Delta_{t-1},A_t,\tau_t,g_t,\eta_t\right) &\equiv \left(1-n_t\right)\beta c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon}\Xi_t^N.
\end{aligned}
$$

The first-order conditions with respect to inflation is

$$
\begin{aligned}
(\pi_t+1)^{\epsilon-2}\Bigg\{\frac{\alpha_t}{\xi}\Bigg[&\epsilon(\pi_t+1)\Delta_{t-1}c_t^{\gamma}\left(\frac{(\tau_t+1)(c_t+G_t)}{A_t}\left(\frac{\Delta_t y_t}{A_t}\right)^{\omega}+\mathbb{E}_t\big(G(\Omega_{t+1})\big)\right) \\
&-(\epsilon-1)\left(c_t^{\gamma}\mathbb{E}_t\big(F(\Omega_{t+1})\big)+c_t+G_t\right)\Bigg] \\
&-(n_t-1)\left[\epsilon(\pi_t+1)\Delta_{t-1}\zeta_t+(\epsilon-1)\rho_t\right]\Bigg\}=0
\end{aligned}
$$

The conditions with respect to p_t^* is

$$
\begin{aligned}
&\left[\mu_t-\frac{(\epsilon-1)\alpha_t}{\xi}(p_t^*)^{-\epsilon}\right]\left[c_t^{\gamma}\mathbb{E}_t\big(F(\Omega_{t+1})\big)+c_t+G_t\right] \\
&+\frac{\epsilon\alpha_t c_t^{\gamma}}{\xi}(p_t^*)^{-\epsilon-1}\left[\frac{(\tau_t+1)(c_t+G_t)}{A_t}\left(\frac{\Delta_t y_t}{A_t}\right)^{\omega}+\mathbb{E}_t\big(G(\Omega_{t+1})\big)\right] \\
&-n_t(p_t^*)^{-\epsilon-1}\left[\epsilon\zeta_t+(\epsilon-1)\rho_t p_t^*\right]=0
\end{aligned}
$$

The conditions with respect to Δ_t is

$$
\begin{aligned}
& c_t^{\gamma}\left\{\frac{\alpha_t}{\xi}\left[(p_t^*)^{1-\epsilon}-(\pi_t+1)^{\epsilon-1}\right]+\mu_t p_t^*\right\}\frac{\partial \mathbb{E}_t(F)}{\partial \Delta_t} \\
& -c_t^{\gamma}\left\{\frac{\alpha_t}{\xi}\left[(p_t^*)^{-\epsilon}-(\pi_t+1)^{\epsilon}\Delta_{t-1}\right]+\mu_t\mathcal{M}\right\} \\
& \quad\times\left[\frac{\omega(\tau_t+1)(c_t+G_t)}{A_t\Delta_t}\left(\frac{\Delta_t y_t}{A_t}\right)^{\omega}+\frac{\partial \mathbb{E}_t(G)}{\partial \Delta_t}\right] \\
& -\frac{y_t}{A_t}\left(\frac{\Delta_t y_t}{A_t}\right)^{\omega}-\zeta_t+\beta\frac{\partial \mathbb{E}_t(V)}{\partial \Delta_t}=0
\end{aligned}
$$

The conditions with respect to c_t is

$$
\begin{aligned}
& \frac{\alpha_t}{\xi}\Bigg[A_t^{-\omega-1}c_t^{\gamma-1}(p^*)_t^{-\epsilon}\Big((\pi_t+1)^{\epsilon}\Delta_{t-1}(p^*)_t^{\epsilon}-1\Big) \\
& \qquad\times\Big(\gamma A_t^{\omega+1}\mathbb{E}_t\big[G(\Omega_{t+1})\big]+(\tau_t+1)\Delta_t^{\omega}y_t^{\omega}(\gamma c_t+c_t+\gamma G_t)\Big) \\
& \qquad+\Big((p^*)_t^{1-\epsilon}-(\pi_t+1)^{\epsilon-1}\Big)\Big(\gamma c_t^{\gamma-1}\mathbb{E}_t\big[F(\Omega_{t+1})\big]+1\Big)\Bigg] \\
& +\mu_t p_t^*\Big(\gamma c_t^{\gamma-1}\mathbb{E}_t\big[F(\Omega_{t+1})\big]+1\Big) \\
& -\mathcal{M}\mu_t A_t^{-\omega-1}c_t^{\gamma-1}\Big(\gamma A_t^{\omega+1}\mathbb{E}_t\big[G(\Omega_{t+1})\big]+(\tau_t+1)\Delta_t^{\omega}y_t^{\omega}(\gamma c_t+c_t+\gamma G_t)\Big) \\
& -\gamma\alpha_t A_t^{-\omega}c_t^{\gamma-1}\Delta_t^{\omega}y_t^{\omega}\mathrm{ReLU}(n_t-\bar{n}) \\
& +\lambda_t=0
\end{aligned}
\tag{25}
$$

where $\mathrm{ReLU}(x)=\max\{0,x\}$. The conditions with respect to n_t is

$$
\zeta_t\Big((p_t^*)^{-\epsilon}-(1+\pi_t)^{\epsilon}\Delta_{t-1}\Big)+\rho_t\Big((p_t^*)^{1-\epsilon}-(1+\pi_t)^{\epsilon-1}\Big)-\alpha_t A_t^{-\omega}c_t^{\gamma}\Delta_t^{\omega}y_t^{\omega}\,\mathrm{ReLU}'(n_t-\bar{n})=0
$$

The envelope conditions with respect to Δ_{t-1} is

$$
\frac{\partial V}{\partial \Delta_{t-1}}=(\pi_t+1)^{\epsilon}\left(\frac{\alpha_t c_t^{\gamma}\left((\tau_t+1)\,A_t^{-\omega-1}\,(c_t+G_t)\,\Delta_t^{\omega}y_t^{\omega}+\mathbb{E}_t\left(G\left(\Omega_{t+1}\right)\right)\right)}{\xi}-n_t\zeta_t+\zeta_t\right)
$$

E Additional figures

In this appendix, we present additional figures that illustrate the main results of the paper. Figure S.1 depicts a stylized representation of the neural network architecture, while Figure S.2 plots the real oil price used to motivate the regime calibration. Figures S.3–S.8 report

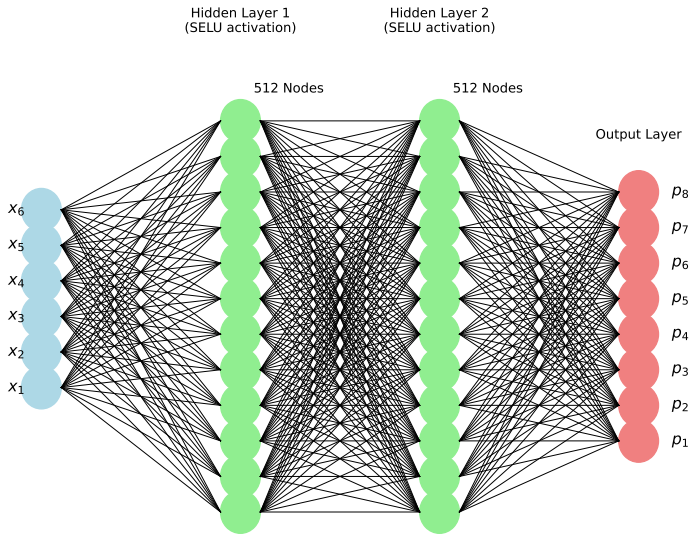

Figure S.1: The figure depicts a stylized neural network with an input $\mathbf{x} \in \mathbb{R}^6$. It consists of two hidden layers, each containing 512 neurons, using SELU activation functions, and produces an output $\mathbf{p}(\mathbf{x}) \in \mathbb{R}^8$.

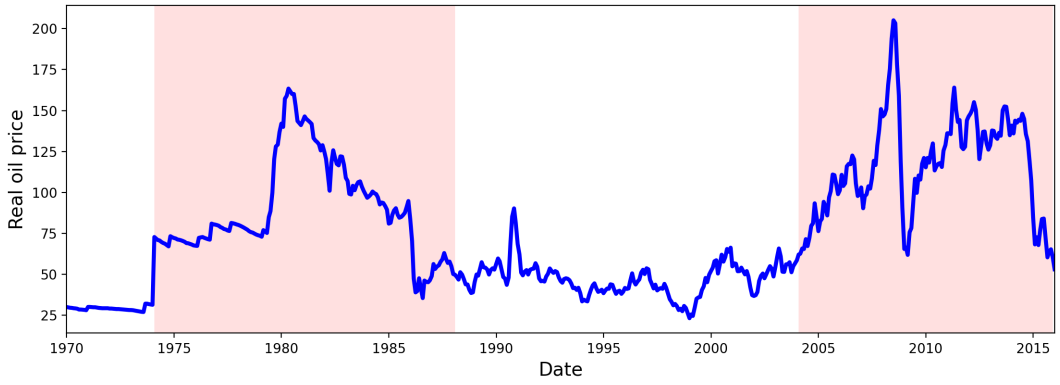

Figure S.2: Real oil prices.

Note: West Texas Intermediate (WTI) prices deflated by the consumer price index (CPI). Shaded areas correspond to periods in which the real oil price exceeds 60 USD per barrel.

additional results referenced in Section 4 of the main paper.

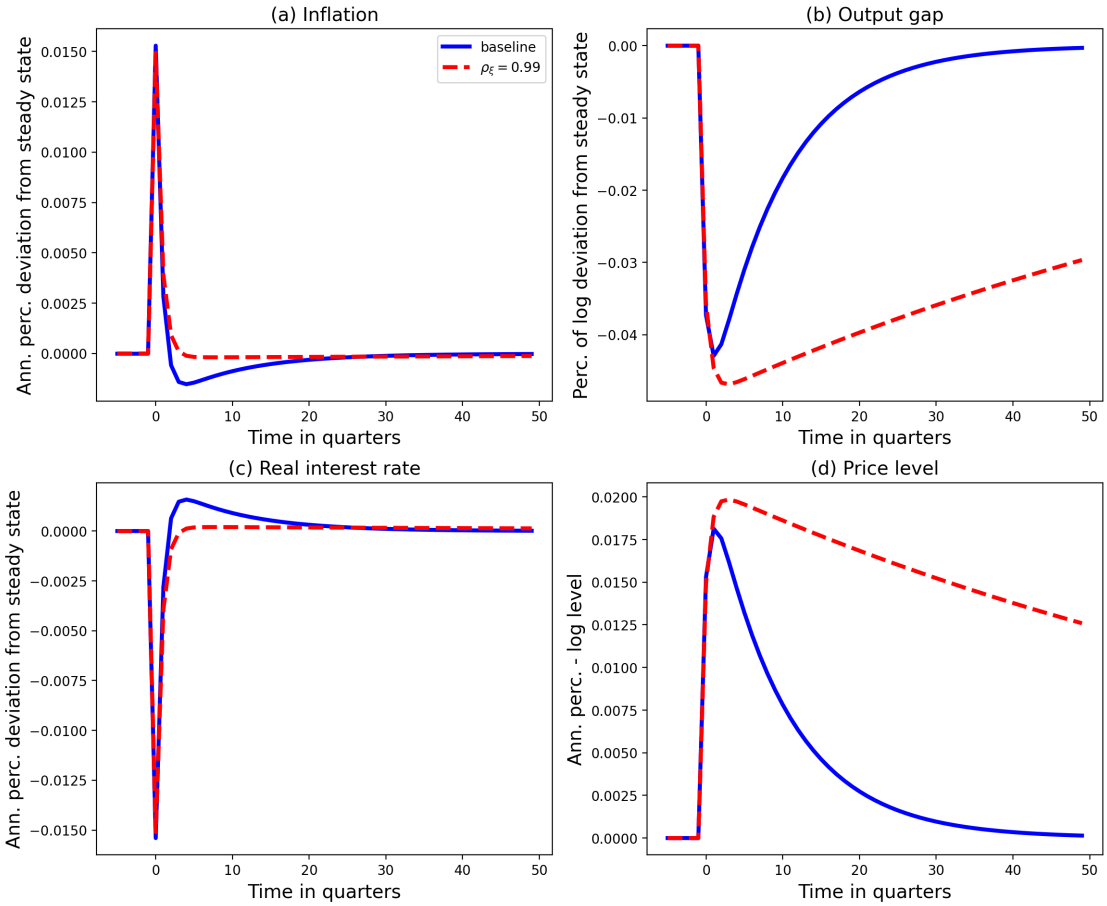

Figure S.3: Impulse response to a transitory cost-push shock with different levels of persistence.

Note: The figure displays the impulse responses to a temporary cost-push shock under the optimal policy with commitment. We consider two shocks with different autoregressive coefficients in a model without regime-switching cost-push shocks. The solid blue line corresponds to the baseline with $\rho^\tau = 0.90$ and the dashed red line to the more persistent case with $\rho^\tau = 0.99$.

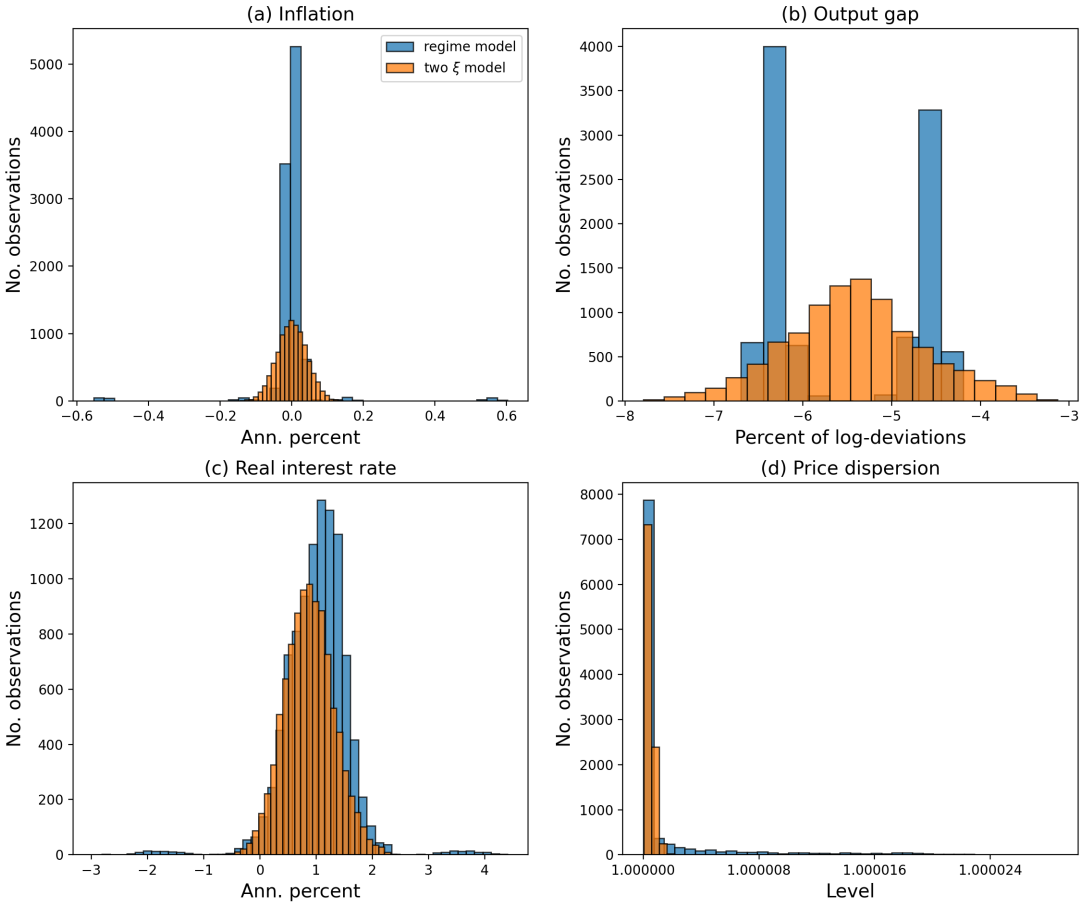

Figure S.4: Ergodic distributions under commitment: two AR(1) cost-push shocks (orange) versus supply regimes (blue).

Note: The figure shows the ergodic distribution in the model with two AR(1) cost-push shocks under commitment. Colors distinguish the two models: blue corresponds to the baseline model with regime switching, and orange to the standard New Keynesian model calibrated with two cost-push shocks with different autoregressive coefficients, $\rho^{\tau} = 0.90$ and $\rho^{\tau} = 0.99$.

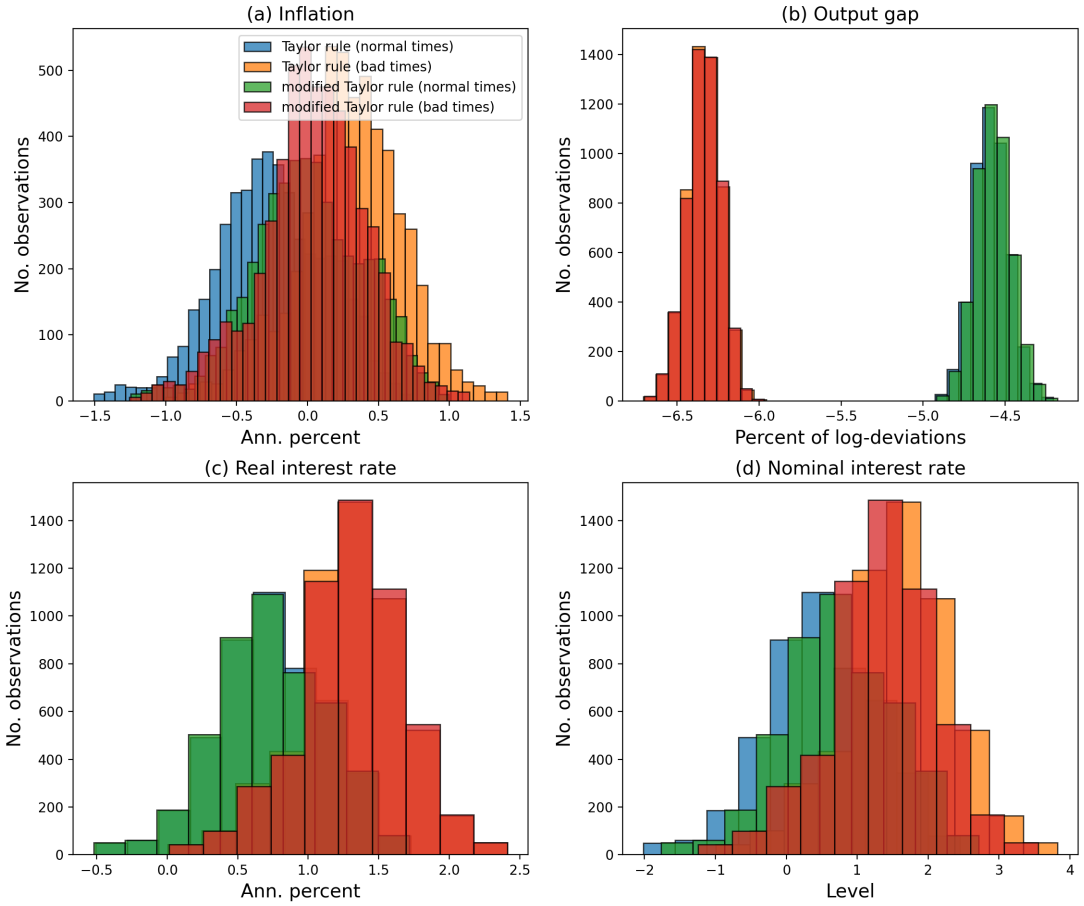

Figure S.5: Ergodic distribution (Taylor rule).

Note: The figure shows the ergodic distribution of the model under a standard and a modified Taylor rule. Colors indicate the two regimes: blue and orange correspond to the standard rule in normal and bad times, respectively, while green and red represent the modified rule in normal and bad times. The distribution is obtained by simulating the model over a large number of periods.

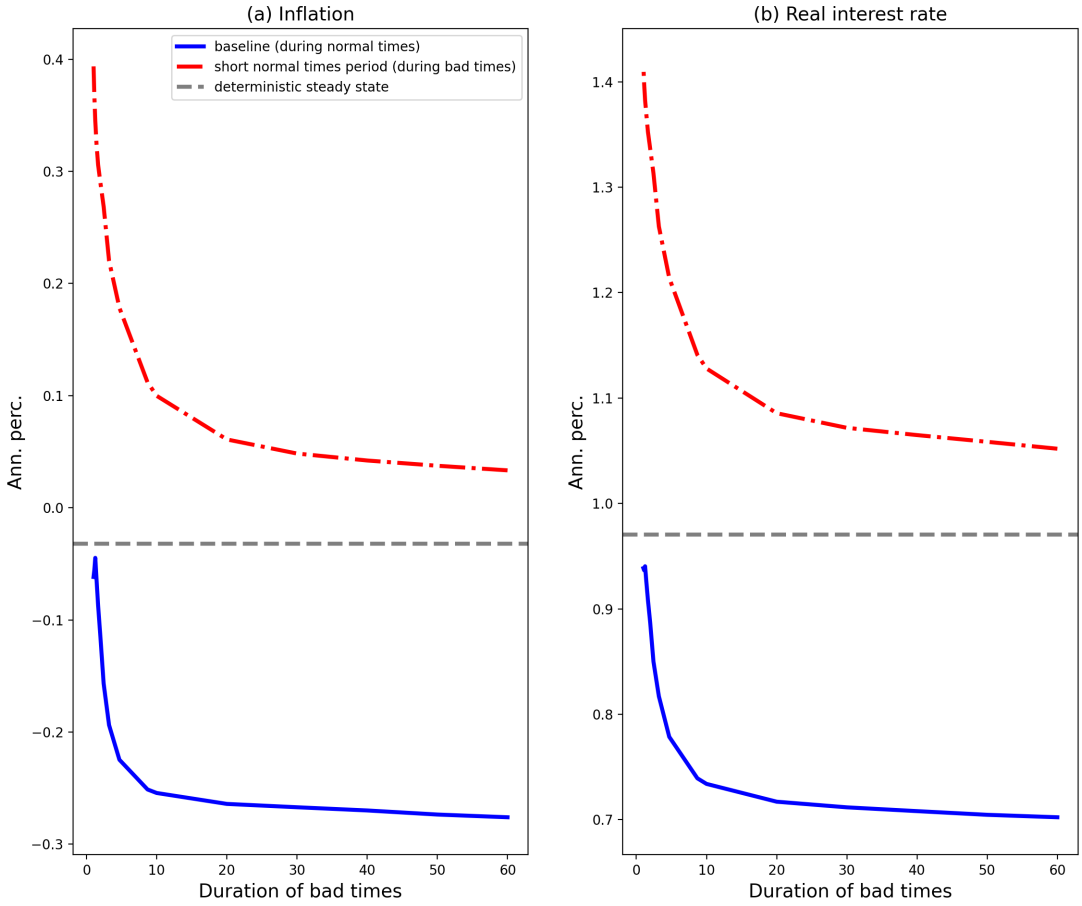

Figure S.6: Sensitivity to regime length (Taylor rule).

Note: The figure shows the SSS values of inflation and the real interest rate for alternative average durations of the bad-times regime. The solid blue line represents the normal-times SSS in the baseline model, the dash-dotted red line the bad-times SSS in a counterfactual case where $p_{12} \to 1$, and the dashed gray line the values in the efficient allocation.

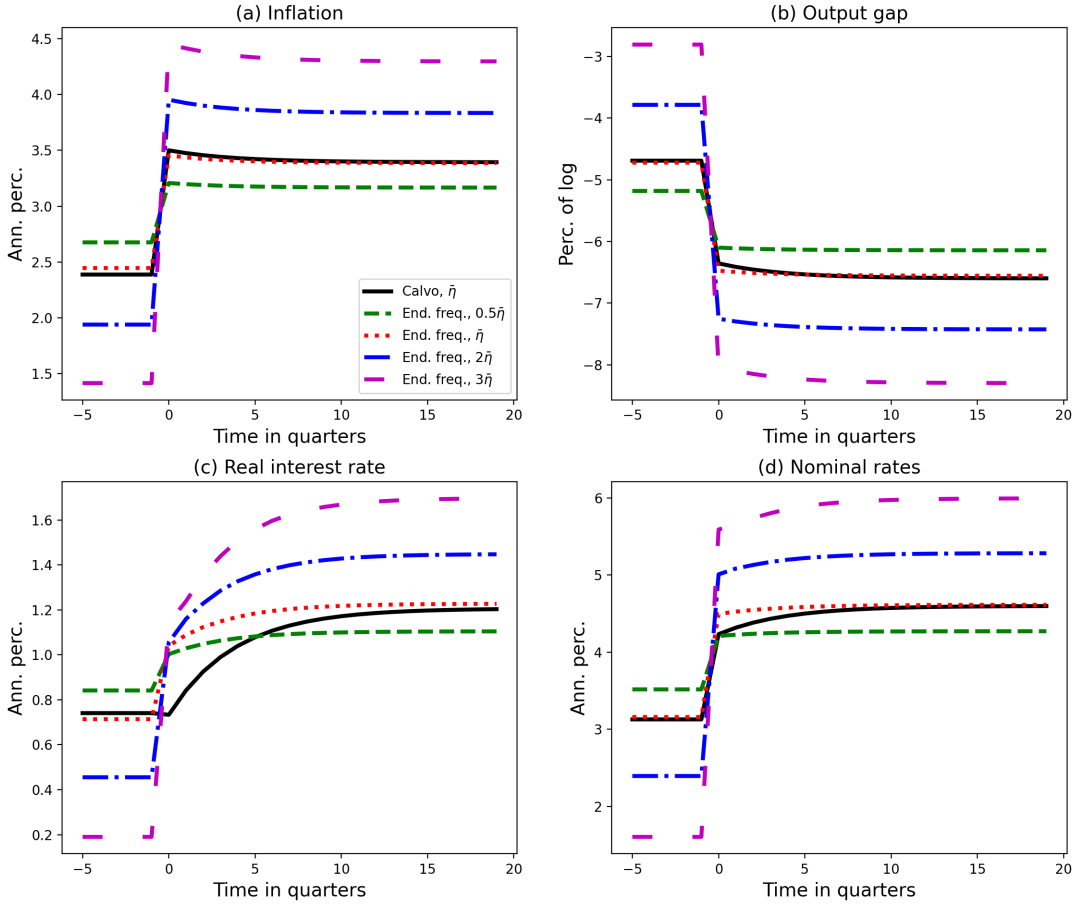

Figure S.7: Response to a regime change with endogenous pricing frequency for different Markovian shocks compared to Calvo pricing.

Note: The figure shows the transition from the normal-times regime to the bad-times regime under the optimal policy with discretion for different magnitudes of the Markovian cost-push shock: $\bar{\eta} = 0.025$ (dotted red line for endogenous frequency and solid black for Calvo), $0.5\bar{\eta}$ (dashed green line, endogenous frequency), $2\bar{\eta}$ (dash-dotted blue line, endogenous frequency) and $3\bar{\eta}$ (wide-dashed magenta line, endogenous frequency). Innovations to the temporary shocks are set to zero, and the economy starts from the SSS of the normal-times regime.

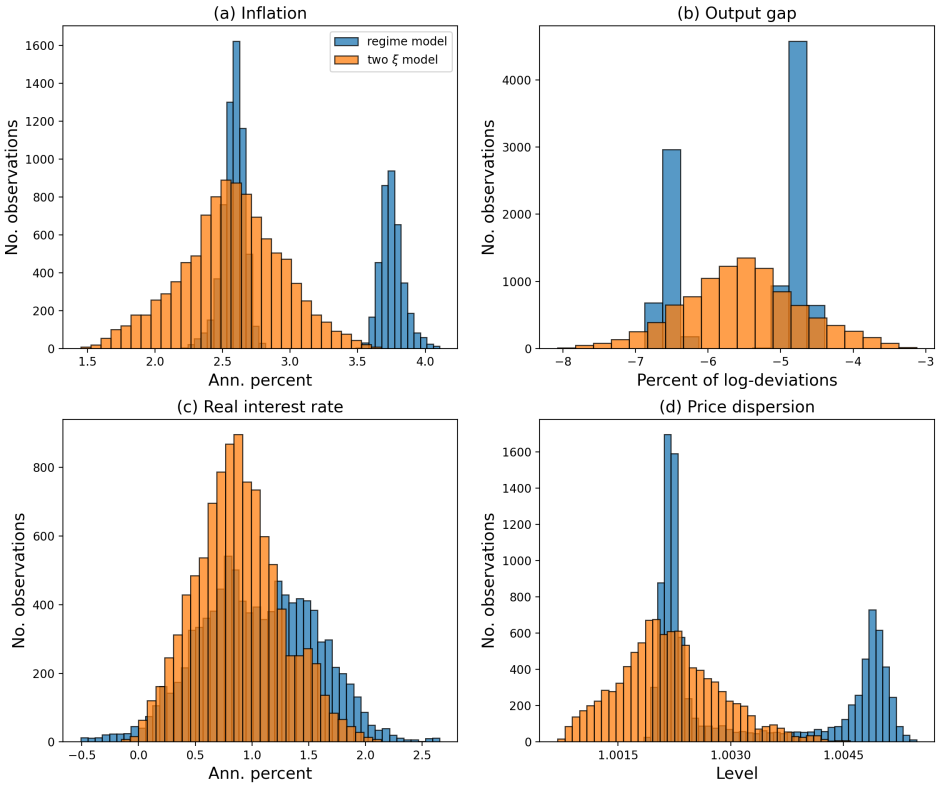

Figure S.8: Ergodic distributions under discretion: two AR(1) cost-push shocks (orange) versus supply regimes (blue).

Note: The figure shows the ergodic distribution in the model with two AR(1) cost-push shocks under discretion. Colors distinguish the two models: blue corresponds to the baseline model with regime switching, and orange to the standard New Keynesian model calibrated with two cost-push shocks with different autoregressive coefficients, $\rho^{\tau} = 0.90$ and $\rho^{\tau} = 0.99$.

References

Azinovic, M., L. Gaegauf, and S. Scheidegger (2022): "Deep Equilibrium Nets," *International Economic Review*, 63, 1471–1525.

Benigno, G., B. Hofmann, G. N. Barrau, and D. Sandri (2024): "Quo vadis, r*? The natural rate of interest after the pandemic," *BIS Quarterly Review*.

Blanco, A., C. Boar, C. J. Jones, and V. Midrigan (2024): "The Inflation Accelerator," NBER Working Papers 32531, National Bureau of Economic Research, Inc.

Brumm, J. and S. Scheidegger (2017): "Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models," *Econometrica*, 85, 1575–1612.

Campos, R. G., J. Fernández-Villaverde, G. Nuño, and P. Paz (2024): "Navigating by Falling Stars: Monetary Policy with Fiscally Driven Natural Rates," NBER Working Papers 32219, National Bureau of Economic Research, Inc.

Kase, H., L. Melosi, and M. Rottner (2025): "Estimating nonlinear heterogeneous agent models with neural networks," BIS Working Papers 1241, Bank for International Settlements.

Kingma, D. P. and J. Ba (2014): "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*.

Kübler, F., S. Scheidegger, and O. Surbek (2025): "Using machine learning to compute constrained optimal carbon tax rules," *arXiv preprint arXiv:2507.01704*.