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

Abstract

We analyze optimal monetary policy in a New Keynesian model with supply regimes—periods of persistently higher costs associated with commodity price supercycles, tariffs, or geopolitical fragmentation. Regime-dependent markup distortions generate bimodal output-gap distributions, while consumption expectations and saving incentives shift the natural rate across regimes. We solve globally for optimal policy using deep equilibrium nets. Under commitment, the central bank accommodates regime shifts through temporary inflation without subsequently reversing the price-level increase. Under discretion, greater distortions in bad times strengthen the incentive to stimulate the economy, raising inflation expectations and amplifying the inflationary bias.

JEL Classification: E32, E58, E63.

Keywords: Deep learning, Markov-switching Model, Cost-push shocks.

*This draft supersedes a previous version entitled "Monetary Policy with Persistent Supply Shocks". This paper was partially completed while Galo Nuño was visiting the BIS and Simon Scheidegger was a Research Fellow at the BIS. We thank Ryan Banerjee, Florin Bilbiie, Eduardo Davila, Davide Debortoli, Jose Elias Gallegos, Veronica Guerrieri, Boris Hofmann, Jesper Lindé, Joël Marbet, Leonardo Melosi, Dominik Thaler, Mathias Trabandt, and participants at various seminars and conferences for their valuable comments. All remaining errors are ours. We are grateful to Katti Irastorza for excellent research assistance. This work was generously supported by the Swiss National Science Foundation (SNSF), under the project "New methods for asset pricing with frictions".

1 Introduction

Economists have long recognized the importance of supply-side dynamics in shaping macroeconomic fluctuations. As early as the 19th century, Jevons' sunspot theory posited that 11-year agricultural cycles, driven by solar activity, were a major determinant of business cycles. Since then, the literature has broadened to encompass studies of commodity price super-cycles (e.g., Fernandez et al., 2020), characterized by extended phases of rising and declining prices driven by structural changes in global demand and supply.[1] More recently, attention has shifted to supply disruptions stemming from geopolitical conflicts (see, e.g., Federle et al., 2024; Fernandez-Villaverde et al., 2024a), trade tariffs (Werning et al., 2025), and other exogenous disturbances. These phenomena motivate modeling persistent shifts in supply conditions as distinct "supply regimes," alongside standard autoregressive shocks.

This paper examines the design of monetary policy in response to shifts in the supply regime by extending the standard New Keynesian framework (Clarida et al., 1999, Woodford, 2003) to allow for periods of sustained change in production costs. The central innovation of our model is the introduction of supply regimes that modify the mean of the cost-push shock according to a two-state Markov-switching process. The economy alternates between "normal times," when the mean of the cost-push shock is negative, and "bad times," when it is positive, reflecting a uniform increase in firms' production costs. This mechanism captures persistent, broad-based increases in marginal costs and serves as a reduced-form representation of structural cost shocks, such as rising energy prices, higher labor-market regulation, or trade fragmentation.[2]

Apart from this regime-switching feature, the model follows the canonical New Keynesian structure. A representative household consumes a continuum of differentiated goods and supplies labor in a centralized, frictionless market. Each good is produced by a monopolistically

[1] Even if global supply and demand affect commodity prices, changes in prices of key commodities such as oil or gas operate as supply shocks for non-producing economies.

[2] For example, Afrouzi et al. (2023) interpret a similar labor wedge as a proxy for shifts toward more regulated labor markets or a deceleration in globalization.

competitive firm using labor as the sole input. Nominal price rigidities are introduced *à la* Calvo (Calvo, 1983), such that firms can adjust prices only with a certain probability, while the remaining prices remain fixed. The economy is further subject to standard aggregate productivity and government spending shocks, both modeled as autoregressive processes in logs with zero-mean innovations. Monetary policy is conducted by a central bank that adjusts the short-term nominal interest rate. We study the optimal design of monetary policy under both discretion and commitment, and additionally assess the performance of Taylor-type policy rules in the Supplementary Material.

Deep equilibrium nets. The introduction of regime-switching dynamics substantially increases the model's computational complexity, requiring a global solution method that captures nonlinear interactions across regimes.[3] To address this challenge, we extend the framework of "deep equilibrium nets" of Azinovic et al. (2022) to incorporate Markov-switching dynamics. This methodological contribution enables accurate approximation of the model's nonlinear decision rules and efficient computation of equilibrium outcomes, thereby facilitating the analysis of optimal monetary policy across different supply regimes.

Preview of main mechanisms and policy results. Our contributions are threefold. First, we uncover a new precautionary-savings mechanism through which the supply regime affects the natural rate. In bad times, the higher labor wedge increases the distortion relative to the efficient allocation, reducing output and consumption and widening the negative output gap. Consequently, the economy exhibits bimodality, characterized by two distinct regime-dependent stochastic steady states (SSSs), each corresponding to one of the regimes.[4] The difference in consumption across the two regimes drives the dynamics of the natural interest rate.[5] In normal times, households anticipate a possible shift to a regime with lower average consumption and thus increase desired precautionary saving at a given interest rate.

[3]The presence of seven state variables in the optimal policy problem under commitment further exacerbates the computational burden due to the "curse of dimensionality" (Bellman, 1957).

[4]A regime-dependent stochastic steady state refers to the equilibrium state when continuous shock innovations are zero and the realized regime is fixed, and agents correctly anticipate the stochastic process.

[5]The natural interest rate is defined as the real interest rate in each SSS of a counterfactual economy without nominal rigidities.

With bonds in zero net supply, this lowers the natural rate. Conversely, when the economy enters the bad-times regime, consumption falls, the demand for savings declines as households expect higher consumption once the regime ends, and the natural rate rises.

Second, if the central bank can credibly commit to a state-contingent plan, the optimal policy "leans against the wind" (Galí, 2008). The central bank allows inflation to rise when the economy transitions to the bad-times regime, but subsequently tightens monetary policy gradually to bring inflation back to zero over the following quarters. The central bank thereby partially cushions the bad regime's adverse impact on the output gap, which converges gradually to the new, lower SSS. The optimal response to a supply-regime shift differs substantially from the response to a temporary cost-push shock. In the latter, the central bank commits to a deflationary episode following the inflation spike in order to restore the aggregate price level to its pre-shock value. Along the simulated transition, with subsequent regime realizations held fixed, no deflationary episode follows; the price level remains higher, consistent with a "bygones-are-bygones" policy. The comparison illustrates how the price-level response depends on the persistence of the change in the output gap.

Third, under discretion, the central bank cannot commit to future policy but internalizes how today's price dispersion constrains its future actions. We find that this gives rise to a state-contingent inflationary bias (Kydland and Prescott, 1977, Barro and Gordon, 1983). In normal times, the economy is less distorted and this inflationary bias is milder. When the economy enters the bad-times regime, however, the persistent supply shock increases the SSS distortion by raising the labor wedge. This creates an incentive for the central bank to loosen monetary policy to support output and consumption. Private agents anticipate this incentive and preemptively raise their prices, resulting in persistently positive inflation during bad times. Our model therefore generates an endogenous increase in inflation through a de-anchoring of inflation expectations, not through any change in the central bank's inflation target, but purely through the distortion introduced by the supply regime.

Model extension with endogenous pricing frequency. The fact that inflation

remains substantially above zero when the central bank's optimal policy is discretionary may raise concerns regarding the assumption in the New Keynesian model of a constant repricing frequency irrespective of the inflation level. To address these concerns, we replace Calvo time-dependent pricing by a mechanism *à la* Blanco et al. (2024). In that model, firms sell a continuum of products and choose how many, but not which, prices to adjust in every period, subject to an adjustment cost. Consistent with the empirical evidence, the fraction of price changes evolves endogenously over time and increases with inflation. We analyze optimal monetary policy under discretion in this model extension. Our results show that, for the baseline calibration, the outcomes remain broadly unchanged regardless of the pricing mechanism. Nonetheless, as we increase the severity of the negative supply shock during bad times, the impact on inflation is proportionally lower under endogenous pricing frequency. Higher shock magnitudes imply more negative output gaps and higher inflation, as private agents anticipate the central bank cushioning the shock. If firms can endogenously change the frequency of price changes, higher inflation also leads to higher frequency, reducing nominal rigidities and thus the ability of monetary policy to counteract the real effects of the negative supply shock. As a result, private agents understand that the central bank's incentives to stimulate the economy are proportionally less as the shock size increases, and incorporate this into their inflation expectations, thus moderating the increase in inflation relative to Calvo pricing.

Related literature. This paper relates to six strands of the literature. First, it contributes to the literature on optimal monetary policy design in the non-linear New Keynesian model, both under commitment (see, e.g., Benigno and Woodford, 2005; Yun, 2005; Benigno and Rossi, 2021) and discretion (see, e.g., Albanesi et al., 2003; King and Wolman, 2004; Zandweghe and Wolman, 2019; Arellano et al., 2020; Afrouzi et al., 2023; or Rios-Rull et al., 2026). We revisit the valuable insights of this literature and find that some key results, such as the price-level features of optimal policy under commitment, are modified in the presence of Markov switching regimes.

Second, the paper relates to the extensive literature analyzing monetary policy in regime-switching models (e.g., Schorfheide, 2005; Davig and Doh, 2014; Davig, 2016; Blake and Zampolli, 2011; Debortoli and Nunes, 2014; Bianchi and Melosi, 2017). These models typically analyze regime switches in linear models, implying that the dynamics are local around the deterministic steady state. Our focus instead is on how regime changes open the door to the multiplicity of regime-dependent stochastic steady states under optimal policy, with the economy transitioning among them in response to supply regimes.[6]

Third, while most of the literature analyzing the determinants of the natural rate has focused on how it depends on structural variables such as total factor productivity (TFP) growth or demographics (e.g., Cesa-Bianchi et al., 2022; Gagnon et al., 2021; Del Negro et al., 2017),[7] a number of new works highlight how policies, such as fiscal (Rachel and Summers, 2019; Bayer et al., 2023; Kaplan et al., 2023; Campos et al., 2024) and monetary policy (Fernández-Villaverde et al., 2024; Bianchi et al., 2021) may also affect the natural rate. Our results introduce a new dimension in the debate: supply regimes may strongly affect natural rates through precautionary motives.

Fourth, we analyze optimal discretionary monetary policy when the frequency of price changes is endogenous. This complements an emerging literature analyzing optimal monetary policy in state-dependent models, which includes Nakov and Thomas (2014), Adam and Weber (2019), Blanco (2021), Caratelli and Halperin (2024), or Karadi et al. (2024).

Fifth, our paper extends the emerging literature on the use of deep learning to solve high-dimensional general equilibrium models (e.g., Maliar et al., 2021; Han et al., 2021; Friedl et al., 2023; Gu et al., 2024; Kahou et al., 2024; Kase et al., 2025).[8] Specifically, we extend the "deep equilibrium nets" methodology by Azinovic et al. (2022) to the case of regime switching. Our deep learning algorithm allows us to analyze dynamics globally.

[6] Choi and Foerster (2021) use global methods to analyze optimal simple rules in a regime-switching model. Our focus is on fully state-contingent policy under commitment and discretion.

[7] Climate change, competition, and inequality have also been suggested as drivers of natural rates, see Sahuc et al. (2023), Romei et al. (2025) or Mian et al. (2021).

[8] See Scheidegger (2026) for a textbook treatment, and Fernandez-Villaverde et al. (2024b) for a recent review.

Sixth, our paper relates to the new literature on the monetary policy response to tariffs (e.g., Bergin and Corsetti, 2023; Bianchi and Coulibaly, 2025; Monacelli, 2025). In particular, Werning et al. (2025) highlights how tariff shocks can be seen as cost-push shocks.

The paper is organized as follows. Section 2 presents the model and optimal policy problems. Section 3 characterizes efficient and flexible-price allocations and natural rates. Section 4 presents the quantitative results and the extension with endogenous pricing frequency. Section 5 concludes. In addition, the Supplementary Material provides derivations of the optimal policy conditions (Appendix A), details of the deep equilibrium nets algorithm (Appendix B), an evaluation of monetary policy rules (Appendix C), the full analytical setup for the extension with endogenous pricing frequency (Appendix D), and additional figures (Appendix E).

2 Model

In this section, we present the formal structure of the model. It builds on the standard New Keynesian framework (Galí, 2008), augmented to incorporate regime-switching cost-push shocks. Time t is discrete, and the economy consists of three types of agents: households, firms, and a central bank.

2.1 Households

Households consume goods c_t, and supply labor h_t to maximize the expected discounted utility:

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\gamma}}{1-\gamma} - \frac{h_t^{1+\omega}}{1+\omega} \right) \right],$$

where β is the discount factor, γ denotes the coefficient of relative risk aversion, and ω is the inverse of the Frisch elasticity of labor supply. Households face the budget constraint:

$$p_t c_t + B_t \leq p_t w_t h_t + (1 + i_{t-1}) B_{t-1} + T_t,$$

where B_t are holdings of a nominal bond that pays a gross return of $1 + i_t$, w_t is the real wage, p_t is the price level, and T_t denotes nominal profits plus net lump-sum fiscal transfers. The labor wedge is treated as a tax or subsidy; its receipts finance government spending, with lump-sum taxes or transfers balancing the fiscal budget.

The optimal choices of consumption, labor, and bond holdings satisfy the following first-order conditions:

$$c_t^{-\gamma} = \lambda_t, \tag{1}$$

$$h_t^{\omega} = w_t \lambda_t, \tag{2}$$

$$\lambda_t = \beta E_t \left[\left(\frac{1 + i_t}{1 + \pi_{t+1}} \right) \lambda_{t+1} \right], \tag{3}$$

where λ_t is the Lagrange multiplier associated with the budget constraint in real terms, and inflation is $\pi_t \equiv \frac{p_t}{p_{t-1}} - 1$.

2.2 Final good producer

A competitive representative final-good producer aggregates a continuum of differentiated goods of measure one, indexed by j. Cost minimization implies that the demand for the j-th variety of the good satisfies

$$y_t(j) = \left(\frac{p_t(j)}{p_t} \right)^{-\epsilon} y_t, \tag{4}$$

where y_t denotes total output of the final good, ϵ is the elasticity of substitution across varieties, and p_t is the aggregate price index, defined as

$$p_t = \left[\int_0^1 \left(p_t(j) \right)^{1-\epsilon} dj \right]^{\frac{1}{1-\epsilon}}, \tag{5}$$

with $p_t(j)$ denoting the nominal price of variety j.

2.3 Retail good producers

There is a continuum of monopolistic firms, each producing a differentiated variety indexed by j. Production uses labor as the sole input according to

$$y_t(j) = A_t h_t(j),$$

where A_t denotes stochastic total factor productivity. Firms are subject to a temporary cost-push shock ξ_t, and a regime-switching component η_t that governs the prevailing supply regime, as described below.

We define the labor wedge as

$$(1 + \tau_t) \equiv (1 + \xi_t + \eta_t), \tag{6}$$

which lumps together the two cost-push shocks. The firm's total real cost is $(1 + \tau_t) w_t h_t(j)$.

Each firm has monopoly power over its variety and takes the demand schedule in (4) as given. We assume price stickiness *à la* Calvo: each retailer receives a random signal to adjust their prices with a probability $1 - \theta$, allowing them to choose a new price $p_t(j)$ to maximize the stream of expected profits, that is:

$$\max_{P_t^*(j)} \mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} \left[\frac{P_t^*(j)}{p_{t+k}} y_{t+k}(j) - \Psi\left(y_{t+k}(j)\right) \right],$$

subject to (4), where $\Psi\left(y_{t+k}(j)\right) \equiv (1 + \tau_{t+k}) w_{t+k} \left(\frac{y_{t+k}(j)}{A_{t+k}} \right)$ are total costs, and $\Lambda_{t,t+k} \equiv \beta^k \frac{\lambda_{t+k}}{\lambda_t}$ is the stochastic discount factor for payments between periods t and $t + k$. Using (4) the optimization problem can be written as:

$$\max_{P_t^*(j)} \mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} \left[\frac{P_t^*(j)}{p_{t+k}} \left(\frac{P_t^*(j)}{p_{t+k}} \right)^{-\epsilon} y_{t+k} - \Psi\left(\left(\frac{P_t^*(j)}{p_{t+k}} \right)^{-\epsilon} y_{t+k} \right) \right].$$

The corresponding first-order condition is

$$\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} y_{t+k} \left[\left(P_t^*(j)\right)^{-\epsilon} \left(p_{t+k}\right)^{\epsilon-1} - \mathcal{M} \Psi' \left(y_{t+k}(j)\right) \left(P_t^*(j)\right)^{-\epsilon-1} \left(p_{t+k}\right)^{\epsilon} \right] = 0,$$

where $\mathcal{M} \equiv \frac{\epsilon}{\epsilon-1}$, and where $\Psi' \left(y_{t+k}(j)\right) = w_{t+k} \left(1+\tau_{t+k}\right) \left(A_{t+k}\right)^{-1}$.

We assume a symmetric equilibrium in which all resetting firms choose the same price, and thus $P_t^*(j) = P_t^*$ holds. The optimal price is given by:

$$p_t^* \equiv \frac{P_t^*}{p_t} = \mathcal{M} \frac{\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} y_{t+k} \left(p_{t+k}/p_t\right)^{\epsilon} \Psi' \left(y_{t+k}\right)}{\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} y_{t+k} \left(p_{t+k}/p_t\right)^{\epsilon-1}}. \tag{7}$$

Finally, we can express the numerator and denominator in (7) as

$$\Xi_t^N = y_t w_t \left(1+\tau_t\right) \left(A_t\right)^{-1} + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} \left(1+\pi_{t+1}\right)^{\epsilon} \Xi_{t+1}^N \right], \tag{8}$$

$$\Xi_t^D = y_t + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} \left(1+\pi_{t+1}\right)^{\epsilon-1} \Xi_{t+1}^D \right]. \tag{9}$$

2.4 Central Bank

The central bank sets the nominal interest rate on bonds. Bonds are in zero net supply $B_t = 0$. In Section 2.7 below, we define the optimal monetary policy problem both under discretion and commitment.

2.5 Market clearing conditions

The market-clearing condition for goods is

$$y_t = c_t + g_t,$$

where g_t denotes government spending, which evolves according to a stochastic process.

Since all firms face the same probability θ of keeping prices fixed, the law of large numbers

ensures that a fraction θ of firms will keep their prices fixed, while the remaining fraction, $(1-\theta)$, will optimally reset their prices. As a result, equation (5) becomes:

$$p_t = \left\{\theta \left(p_{t-1}\right)^{1-\epsilon} + (1-\theta)\left(P_t^*\right)^{1-\epsilon}\right\}^{\frac{1}{1-\epsilon}},$$

which implies that

$$1 = \theta \left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}. \tag{10}$$

The market-clearing condition for labor is given by:

$$h_t = \int h_t\left(j\right) dj = \int \left(\frac{y_t\left(j\right)}{A_t}\right) dj = \left(\frac{y_t}{A_t}\right) \int \left(\frac{p_t\left(j\right)}{p_t}\right)^{-\epsilon} dj.$$

We define price dispersion as:

$$\Delta_t \equiv \int \left(\frac{p_t\left(j\right)}{p_t}\right)^{-\epsilon} dj = \theta\left(1+\pi_t\right)^{\epsilon} \Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon},$$

such that the aggregate production function becomes:

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

2.6 Shocks and regimes

In what follows, we define the stochastic processes governing the temporary shocks and the regime dynamics.

Temporary shocks. We consider TFP, government spending, and cost-push shocks, following AR(1) processes in logs for the first two and in levels for the third. First, we define:

$$g_t = \bar{g}\tilde{g}_t,$$

where $\bar{g}$ is a constant. Then, we have:

$$\log(A_t) = \left(1-\rho^A\right)\left(-\frac{\left(\sigma^A\right)^2}{2(1-(\rho^A)^2)}\right) + \rho^A \log(A_{t-1}) + \varepsilon_t^A,$$

$$\log(\tilde{g}_t) = (1-\rho^g)\left(-\frac{(\sigma^g)^2}{2(1-(\rho^g)^2)}\right) + \rho^g \log(\tilde{g}_{t-1}) + \varepsilon_t^g,$$

and:[9]

$$\xi_t = \rho^\tau \xi_{t-1} + \varepsilon_t^\tau,$$

where $\varepsilon_t^A \sim N\left(0, (\sigma^A)^2\right)$, $\varepsilon_t^g \sim N\left(0, (\sigma^g)^2\right)$, and $\varepsilon_t^\tau \sim N\left(0, (\sigma^\tau)^2\right)$.

Regimes. We assume that the shock η_t follows a two-state Markov chain. The regime component shifts the labor wedge $\tau_t = \xi_t + \eta_t$: in regime 1 ("normal times"), $\eta_t = -\bar{\eta}_1 < 0$; in regime 2 ("bad times"), $\eta_t = \bar{\eta}_2 > 0$. The regime component therefore affects costs in both regimes. The transition probabilities between the two regimes are defined as

$$p_{12} = \mathbb{P}\left(\eta_t = \bar{\eta}_2 \mid \eta_{t-1} = -\bar{\eta}_1\right), \qquad p_{21} = \mathbb{P}\left(\eta_t = -\bar{\eta}_1 \mid \eta_{t-1} = \bar{\eta}_2\right). \tag{11}$$

All expectations are taken with respect to both the AR(1) shocks and the Markov regimes.

2.7 Optimal monetary policy: problem statement

We now turn to the definition of the optimal monetary policy problem under both discretion and commitment.

2.7.1 Discretion

First, consider the case in which the central bank maximizes household welfare under discretion. The central bank value function is $V(\Delta, A, \tau, g, \eta)$, where price dispersion Δ_{t-1} is the endogenous state variable, $\{A_t, \tau_t, g_t\}$ is a vector of TFP, cost-push, and government

[9] The constant terms in the log processes ensure that the ergodic means of A_t and $\tilde{g}_t$ equal one.

spending shocks.[10] This problem, therefore, includes one endogenous state variable, three continuous exogenous variables, and the discrete regime. The value function satisfies the Bellman equation, that is,

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

subject to the equilibrium conditions,

$$c_t^{-\gamma} = h_t^{\omega} / w_t, \tag{12}$$

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

$$\Delta_t = \theta \left(1+\pi_t\right)^{\epsilon} \Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}, \tag{14}$$

$$y_t = A_t h_t \left(\Delta_t\right)^{-1}, \tag{15}$$

$$y_t = c_t + g_t, \tag{16}$$

$$p_t^* = \mathcal{M} \frac{y_t w_t \left(1+\tau_t\right)\left(A_t\right)^{-1} + \mathbb{E}_t\left[\theta \Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon} \Xi_{t+1}^N\right]}{y_t + \mathbb{E}_t\left[\theta \Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon-1} \Xi_{t+1}^D\right]}. \tag{17}$$

Under discretion, the central bank cannot commit to future policy paths. Future policy functions are taken as given. It, however, understands how its policies may affect price dispersion, which in turn affects expectations by constraining the actions of the central bank itself in the future. The complete set of first-order conditions can be found in Appendix A.2 of the Supplementary Material.

[10] We use $\tau_t = \xi_t + \eta_t$ and retain η_t separately because it affects the transition law.

2.7.2 Commitment

We also consider the case in which the central bank maximizes household welfare under commitment. In this scenario, the central bank solves the Ramsey problem,

$$\max_{\{c_t,h_t,w_t,\pi_t,p_t^*,\Delta_t\}_{t\geq 0}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\gamma}}{1-\gamma} - \frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{1+\omega} \right],$$

subject to the equilibrium conditions (12)-(16) and the following constraints:

$$\begin{aligned} p_t^* \Xi_t^D &= \mathcal{M} \Xi_t^N, \\ \Xi_t^N &= (c_t+g_t)^{1+\omega} \left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^{\gamma} (1+\tau_t) (A_t)^{-1} + \mathbb{E}_t \left[\beta\theta c_t^{\gamma} c_{t+1}^{-\gamma} (1+\pi_{t+1})^{\epsilon} \Xi_{t+1}^N\right], \quad (18) \\ \Xi_t^D &= (c_t+g_t) + \mathbb{E}_t \left[\beta\theta c_t^{\gamma} c_{t+1}^{-\gamma} (1+\pi_{t+1})^{\epsilon-1} \Xi_{t+1}^D\right]. \quad (19) \end{aligned}$$

The first-order conditions can be found in Appendix A.3 of the Supplementary Material. There, it can be seen that the recursive solution to the Ramsey problem depends on seven state variables: the five already present in the discretionary case, plus the two backward-looking Lagrange multipliers associated with equations (18)–(19).

We focus on the optimal "timeless policy" (Woodford, 2003), which assumes that the economy has been operating under the optimal policy for a long period of time. Consequently, the central bank is bound by a past history of commitments. The backward-looking Lagrange multipliers, associated with equations (18) and (19), evolve around their steady-state values.[11]

3 Analytical Characterization

In this section, we characterize the efficient allocation, the flexible-price equilibrium, and the regime-dependent natural rates. We begin by deriving the efficient allocation chosen by

[11] The timeless-policy case contrasts with the case of the optimal "time-0" policy, in which the backward-looking Lagrange multipliers start at zero, reflecting the absence of pre-commitments.

a social planner, which serves as the welfare benchmark of the economy. We then analyze the flexible-price equilibrium, where cost-push shocks distort the marginal conditions, and conclude by characterizing the natural rates that emerge under the Markov-switching supply regimes.

3.1 Efficient allocation and flexible-price equilibrium

Efficient allocation. We begin by analyzing the efficient allocation in the model, defined as the allocation chosen by a social planner who maximizes household welfare subject to the economy's technological constraints. The efficient allocation equates the marginal rate of substitution between consumption and labor, $\hat{h}_t^{\omega}\hat{c}_t^{\gamma}$, to the corresponding marginal rate of transformation, which corresponds to the marginal product of labor, A_t:[12]

$$\hat{h}_t^{\omega}\hat{c}_t^{\gamma} = A_t.$$

Combining this result with the aggregate budget constraint,

$$\hat{c}_t + g_t = \hat{y}_t = A_t\hat{h}_t = A_t^{\frac{1}{\omega}+1}\hat{c}_t^{-\frac{\gamma}{\omega}},$$

implicitly defines efficient consumption $\hat{c}_t$ through the equation

$$\left(\frac{\hat{c}_t + g_t}{A_t}\right)^{\omega} - A_t\hat{c}_t^{-\gamma} = 0.$$

Notice that the efficient allocation is affected by TFP and government expenditure shocks, but not by cost-push shocks. The real interest rate in the non-stochastic steady state is given by $\hat{r} = \frac{1}{\beta} - 1$. We define the efficient output gap as the difference between (log-)consumption

[12] Variables in the efficient allocation are denoted with a hat.

in the baseline model and in the efficient allocation:

$$x_t \equiv \log(c_t) - \log(\hat{c}_t).$$

Flexible-price equilibrium. We next consider the counterfactual equilibrium with flexible prices, that is, when $\theta = 0$. In this case, the optimal relative reset price and price dispersion both equal one, $p_t^* = \Delta_t = 1$, indicating that individual prices are always set optimally. This setup introduces a potential wedge between the marginal rate of substitution between consumption and labor and the marginal rate of transformation, given by:[13]

$$h_t^{*\omega} c_t^{*\gamma} = \frac{A_t}{(1+\tau_t)\,\mathcal{M}}.$$

With the zero-mean temporary shock independent of the regime, $(1+\tau_t)\mathcal{M}$ has a mean of $(1-\bar{\eta}_1)\mathcal{M}$ in the normal-times regime, but a mean of $(1+\bar{\eta}_2)\mathcal{M}$ in the bad-times regime. Consumption in this case satisfies

$$c_t^* + g_t = y_t^* = A_t h_t^* = A_t^{\frac{1}{\omega}+1} \left[(1+\tau_t)\,\mathcal{M}\right]^{-\frac{1}{\omega}} c_t^{*-\frac{\gamma}{\omega}}. \tag{20}$$

This expression can be rewritten as

$$\left(\frac{c_t^* + g_t}{A_t}\right)^{\omega} - \frac{A_t}{(1+\tau_t)\mathcal{M}(c_t^*)^{\gamma}} = 0. \tag{21}$$

Notice that consumption now depends on the cost-push shock. We define a *regime-dependent stochastic steady state* (SSS) in this economy as the fixed point obtained when continuous shock innovations are set to zero and the realized regime is held fixed, while agents continue to internalize the shock dynamics, including the Markov transition law. This is an adaptation of the standard concept of the SSS to the case of a Markov-switching model, which is instrumental in understanding model dynamics. Given that the regime-switching shock has

[13] We denote variables in the flexible-price equilibrium with a star.

two different values, the economy exhibits two SSSs, one in each regime.

3.2 Natural rates

The real interest rate in the flexible-price economy with zero inflation satisfies the Euler equation

$$1 = \beta E_t \left[\frac{{c_t^*}^{\gamma}}{{c_{t+1}^*}^{\gamma}} \right] (1 + r_t^*) .$$

If the economy is in regime 1 (normal times), this equation implies that

$$\frac{1}{\beta (1 + r_t^*)} = {c_{1,t}^*}^{\gamma} \left(p_{12} E_t \left[\frac{1}{{c_{2,t+1}^*}^{\gamma}} \right] + (1 - p_{12}) E_t \left[\frac{1}{{c_{1,t+1}^*}^{\gamma}} \right] \right),$$

where the notation $z_{n,t}$ denotes variable z at time t and regime $n \in \{1, 2\}$. Approximating next-period consumption by its regime-specific SSS value gives

$$1 + r_{1,ss}^* \approx \frac{1}{\beta} \frac{{c_{2,ss}^*}^{\gamma}}{\left(p_{12} {c_{1,ss}^*}^{\gamma} + (1 - p_{12}) {c_{2,ss}^*}^{\gamma} \right)}, \tag{22}$$

so that, in this approximation, if the economy remains in regime 1, consumption remains at its SSS value $c_{1,ss}^*$, whereas if a regime change occurs in the next period, consumption jumps to the SSS in regime 2, $c_{2,ss}^*$. We denote the SSS value of the real interest rate in the flexible-price economy as the *natural rate*, following Obstfeld (2023), who defines it as "real rate prevailing over a long-run equilibrium in the absence of nominal rigidities."

Each of these consumption levels is a solution to the SSS case of equation (20):

$$\left(\frac{c_{n,ss}^* + g_{ss}}{A_{ss}} \right)^{\omega} - \frac{A_{ss}}{\mathcal{M}(1 + \tau_n) {c_{n,ss}^*}^{\gamma}} = 0,$$

where $\mathcal{M}(1 + \tau_n)$ equals $(1 - \bar{\eta}_1)\mathcal{M}$ in regime 1, and $(1 + \bar{\eta}_2)\mathcal{M}$ in regime 2. Due to $(1 - \bar{\eta}_1) < (1 + \bar{\eta}_2)$ and ${c_{n,ss}^*}^{\gamma} \left(c_{n,ss}^* + g_{ss} \right)^{\omega}$ being an increasing function in $c_{n,ss}^*$, in the bad state (state 2), we have $c_{2,ss}^* < c_{1,ss}^*$. Therefore, the denominator on the right-hand side of

equation (22) exceeds the numerator $(c^*_{2,ss})^\gamma$. Within this approximation, this implies that the natural rate in regime 1, $r^*_{1,ss}$, is *lower* than the non-stochastic benchmark, $1/\beta - 1$. Conversely, the natural rate in regime 2, $r^*_{2,ss}$, is *higher* than this benchmark. In our calibration (cf. Section 4.1.1), these values are 0.7 percent and 1.3 percent, respectively, compared to the non-stochastic benchmark of 1 percent, with all three rates expressed on an annualized basis, as discussed below.

The differences in the natural rate as a result of the cost-push shock regime reflect households' consumption-smoothing and precautionary-saving incentives. In normal times, households anticipate that, with a certain probability, the economy may shift to the other regime, where consumption will be lower. In anticipation of this event, they attempt to save more, but with bonds in zero net supply, their increased demand for savings merely leads to a fall in the bond return, that is, in the natural rate. During bad times, expected consumption recovery reduces desired saving at a given interest rate, raising the natural rate.

4 Quantitative analysis

In this section, we quantify the model. We first present the calibration and the global solution method based on deep equilibrium nets adapted to Markov switching. We then study optimal policy under commitment and under discretion. Next, we evaluate simple interest-rate rules, contrasting a standard Taylor rule with a regime-contingent variant. Finally, we extend the framework to endogenous pricing frequency.

4.1 Calibration and solution method

4.1.1 Calibration

The model presented in Section 2 is calibrated at a quarterly frequency, with parameter values reported in Table 1. The calibration follows standard values from the literature whenever possible. Regarding preferences, the quarterly discount factor is set to 0.9975, implying an

annualized real interest rate of approximately 1 percent in the deterministic steady state. The elasticity of substitution across products is $\epsilon = 7$, yielding a frictionless net markup of 1/6. The inverse of the intertemporal elasticity of substitution, γ, is set to 2, and the inverse of the Frisch elasticity of labor supply, ω, is set to 1. The unconditional mean of productivity, A_t, is normalized to one, and the government spending constant, $\bar{g}$, is set to 0.2.

Parameter		Value
Inverse Frisch elasticity	ω	1
Inverse of intertemporal elasticity of substitution	γ	2
Discount factor	β	0.9975
Elasticity of substitution among varieties	ϵ	7
Government spending constant	$\bar{g}$	0.2
Calvo constant	θ	0.75
Mean of cost-push shock in good times	$-\bar{\eta}_1$	-0.025
Mean of cost-push shock in bad times	$\bar{\eta}_2$	0.025
Transition probability from normal to bad times	p_{12}	1/52
Transition probability from bad to normal times	p_{21}	1/52
Persistence of TFP shock	ρ^A	0.99
Persistence of cost-push shock	ρ^τ	0.90
Persistence of government spending shock	ρ^g	0.97
Standard deviation of TFP shock innovations	σ^A	0.009
Standard deviation of cost-push shock innovations	σ^τ	0.0014
Standard deviation of government spending shock innovations	σ^g	0.0052

Table 1: Key parameters of the model.

The numerical values of the innovation standard deviations and the persistence parameters for government spending and cost-push shocks follow Coibion et al. (2012). We set $\rho^A = 0.99$ and use the shock specifications in Section 2.

To calibrate the Markov chain, we consider the oil price super-cycle. Figure S.2 in the Supplementary Material displays the real price of oil. Shaded areas indicate periods in which prices are above 60 USD per barrel for more than a year. We identify three distinct periods: (i) from January 1974 to January 1988, oil prices were high. This period follows the Yom Kippur War and includes the Iranian Revolution and the Iran–Iraq War. The average real oil price was 90 USD per barrel. (ii) From February 1988 to December 2003, real oil prices were

low. The average price was 47 USD. (iii) From January 2004 to November 2015, oil prices were again high, averaging 116 USD.[14]

The two high-price periods average approximately 13 years (52 quarters). We impose this expected duration in both regimes, setting $p_{12} = p_{21} = 1/52$. The average real price in the low-price period is 47 USD, while the average of the two high-price period means is 103 USD, approximately 119 percent higher. To calibrate the labor wedge, we use a heuristic scaling factor of 0.04, motivated by the ratio of the oil-driven standard deviations of wage inflation and real oil prices in the post-1984 sample reported by Blanchard and Galí (2007) (Table 4). Applying this factor to the approximately 119 percent price difference motivates a gap of about five percentage points in the labor wedge between regimes. With symmetric transition probabilities, imposing a zero unconditional mean then gives $-\bar{\eta}_1 = -2.5$ percent and $\bar{\eta}_2 = 2.5$ percent.

4.1.2 Global solution method: deep equilibrium nets

Motivation. Standard perturbation methods provide local approximations (see, e.g., Adjemian et al., 2011, and references therein). Because the model features two distinct regime-dependent stochastic steady states, an approximation around a single steady state may lose accuracy when the economy moves between regimes. This motivates our global approach, which complements earlier analyses based on local approximations (e.g., Schorfheide, 2005; Davig and Doh, 2014; Blake and Zampolli, 2011; Bianchi and Melosi, 2017).

Deep equilibrium nets. The model has three exogenous autoregressive shocks, one endogenous state variable, and two distinct regimes. Optimal monetary policy under commitment introduces two additional endogenous state variables corresponding to the Lagrange multipliers associated with forward-looking equations. This state space makes tensor-product grid methods costly because their computational requirements grow exponentially with the number of continuous states (Bellman, 1957). To address this challenge, we extend the

[14]We discard the period after 2015, which includes events such as the COVID-19 pandemic and the Strait of Hormuz blockade, as it is too early to conclude whether the next phase of the supercycle is over.

deep equilibrium nets (DEQN) approach of Azinovic et al. (2022), adapting it to efficiently handle Markov-switching models. Appendix B of the Supplementary Material outlines the basic DEQN algorithm, following the exposition in Azinovic et al. (2022) and Friedl et al. (2023), and describes our extension.[15] It also provides a thorough assessment of numerical convergence and solution accuracy using standard diagnostics from the literature on deep learning for equilibrium models (see, e.g., Azinovic et al., 2022, Kübler et al., 2025, Kase et al., 2025).

4.2 Optimal policy analysis

4.2.1 Commitment

Ergodic distribution and dynamics. We begin by analyzing the ergodic distribution of the main macroeconomic variables under commitment. The distribution shown in Figure 1 is obtained by simulating the economy over a large number of periods. The blue bars represent the share of observations corresponding to the normal-times regime, while the orange bars indicate those associated with bad times.

Several results emerge. First, the output gap (panel b) and the real interest rate (panel c) exhibit *bimodal* distributions: the realizations cluster around two distinct points. These points approximately correspond to the regime-dependent stochastic steady states of each variable, as reported in Table 2.[16] As discussed in Section 3.1, the bimodality in the output gap and real interest rates arises from the distortion generated by the regime-switching supply shock. Given the lack of an optimal labor subsidy, the steady-state markup is distorted. This distortion widens in bad times and shrinks in normal times. As a result, consumption falls relative to the efficient allocation, but relatively more in bad times, producing a more negative output gap. Anticipation of such shifts in consumption also induces a precautionary

[15] See also Azinovic and Žemlička (2024), who build on DEQN to solve a rare-disaster model with overlapping generations and multiple assets.

[16] Because the model is nonlinear, the distributions are not symmetric around the mean, and the mean does not necessarily coincide with the SSS.

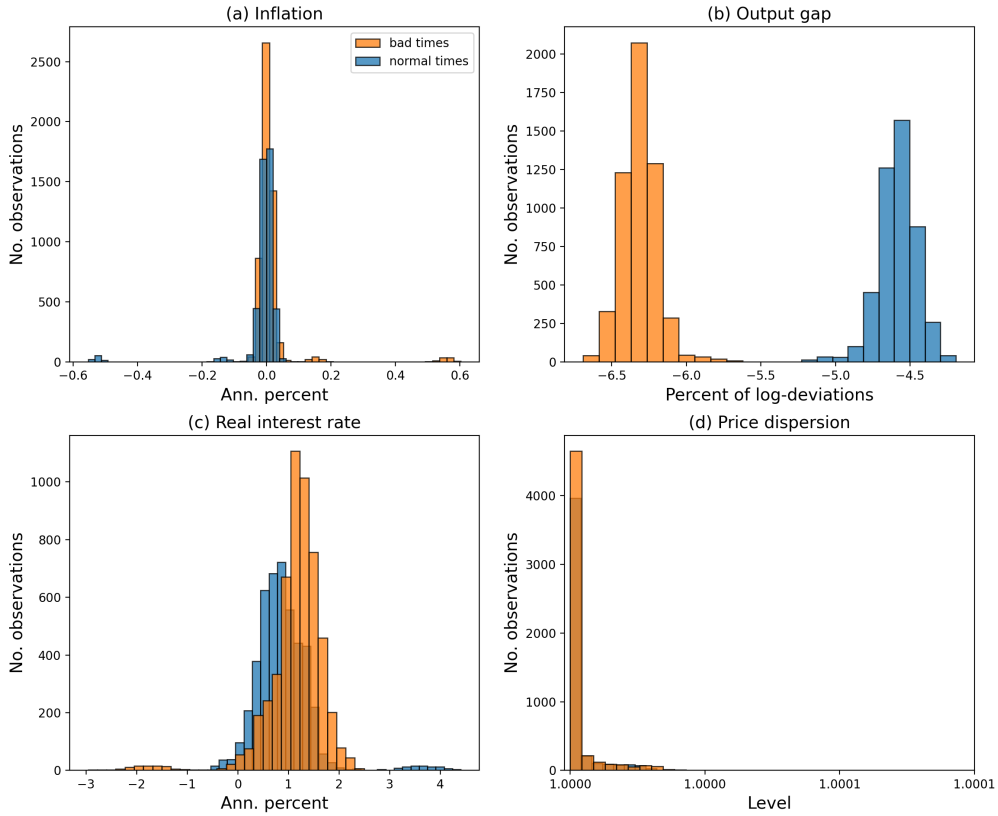

Figure 1: Ergodic distribution: commitment.

Note: The figure shows the ergodic distribution of the model under commitment. Colors indicate the two regimes: blue corresponds to normal times and orange to bad times. The distribution is obtained by simulating the model over a large number of periods.

savings motive, which accounts for the corresponding jumps in the real interest rate.

Second, inflation clusters around zero in both regimes (panel a). This is a well-established result in monetary economics: the steady-state inflation rate under the optimal Ramsey policy with commitment is zero, regardless of whether the steady state is distorted or efficient (Benigno and Woodford, 2005). In our case, the fact that the economy transitions between a more distorted and a less distorted SSS does not alter this result. In response to government spending and TFP shocks, the optimal policy keeps inflation at zero at all times. By contrast, in response to cost-push shocks, the optimal policy "leans against the wind" (Galí, 2008), allowing inflation to move in the opposite direction to the output gap. These policy prescriptions hold in both regimes.

	Flex. prices				Commitment				Discretion			
	SSS	Mean	Std	Skew	SSS	Mean	Std	Skew	SSS	Mean	Std	Skew
Inflation												
normal times	0.00%	0.00%	0.00%	-	-0.00%	-0.01%	0.08%	-5.82	2.59%	2.57%	0.08%	-0.72
bad times	0.00%	0.00%	0.00%	-	-0.00%	0.01%	0.08%	6.33	3.79%	3.80%	0.08%	0.71
Output gap												
normal times	-4.57%	-4.56%	0.12%	0.12	-4.72%	-4.58%	0.13%	-0.64	-4.94%	-4.79%	0.12%	-0.04
bad times	-6.34%	-6.34%	0.11%	-0.15	-6.48%	-6.31%	0.12%	0.71	-6.81%	-6.63%	0.10%	-0.07
Real rates												
normal times	0.68%	0.72%	0.40%	-0.17	0.77%	0.89%	0.57%	2.14	0.77%	0.80%	0.41%	-0.17
bad times	1.23%	1.28%	0.40%	-0.41	1.16%	1.14%	0.57%	-2.55	1.22%	1.25%	0.39%	-0.31
Nominal rates												
normal times	0.68%	0.72%	0.40%	-0.17	0.77%	0.87%	0.52%	1.50	3.36%	3.38%	0.40%	-0.14
bad times	1.23%	1.28%	0.40%	-0.41	1.16%	1.15%	0.52%	-1.92	5.01%	5.05%	0.38%	-0.25

Table 2: Key model moments.

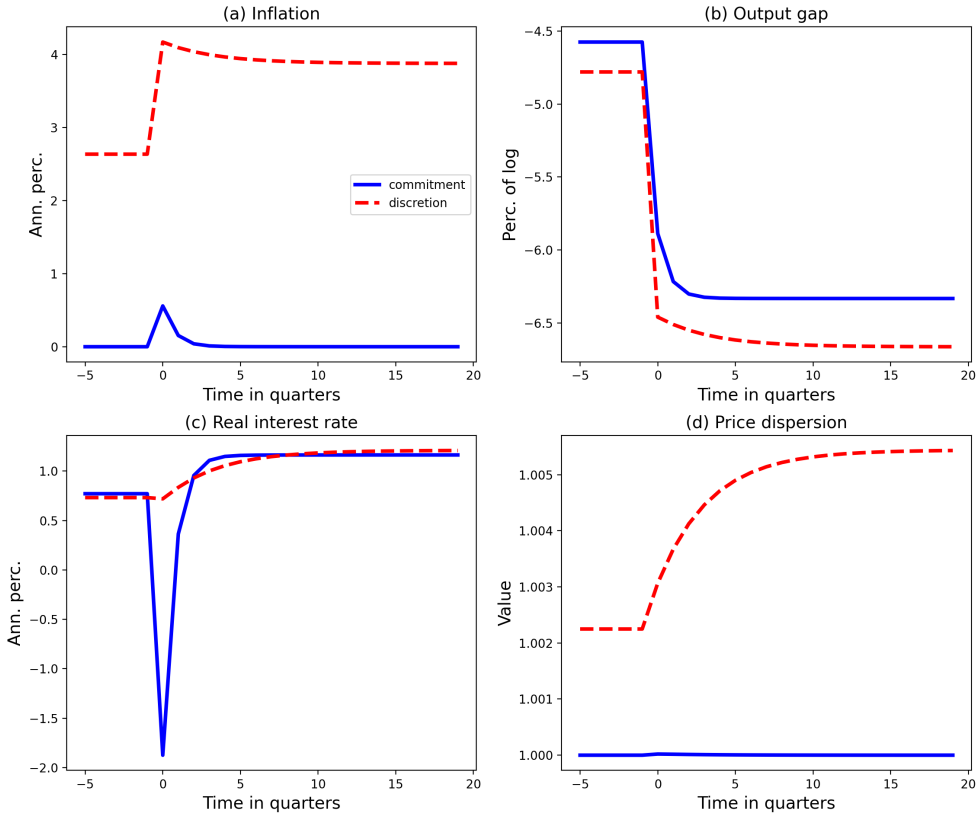

Figure 2: Response to a regime change: commitment versus discretion.

Note: The figure shows the transition from the normal-times regime to the bad-times regime under the optimal policy with commitment (solid blue line) and discretion (dashed red line). Innovations to the temporary shocks are set to zero. The economy starts from the SSS of the normal-times regime.

The comparison with the counterfactual flexible-price economy is informative (see Table 2). In both cases, steady-state inflation is zero; however, in the presence of nominal rigidities,

inflation deviates from zero in response to cost-push shocks. This accounts for the positive conditional standard deviation of inflation under commitment in both normal and bad times, whereas it is zero under flexible prices.

Third, the distribution of price dispersion (panel d) is similar in both regimes, and it is strongly asymmetric. It concentrates around one and has a right tail.

Notice that while the SSS of the output gap in bad times is lower than under flexible prices (-6.48 percent versus -6.34 percent), its conditional mean is higher (-6.31 percent versus -6.34 percent). This implies that, on average, the consumption loss in the economy with nominal rigidities is smaller than under flexible prices during bad times. To understand this result, Figure 2 (solid blue line) shows the transitional dynamics following a regime change. The simulation assumes that the economy starts from the SSS of the normal-times regime and transitions to the bad-times regime at time zero and remains there, with all other shocks set to zero. In response to this transition, the central bank tolerates a temporary increase in inflation (panel a) by initially lowering the real interest rate (panel c). This policy results in a smoother adjustment of the output gap towards its new SSS (panel b). The smoother transition of the output gap under nominal rigidities contrasts with the case of flexible prices, where the output gap immediately "jumps" to the new SSS (not shown), as discussed in equation (21).

Another important difference between the baseline model under commitment and the flexible-price economy is that the gap between the two real-rate regimes is narrower under commitment. While under flexible prices the two real-rate means are 0.72 percent and 1.28 percent, under commitment they are 0.89 percent and 1.14 percent. This occurs because precautionary motives are weaker under commitment: by conducting state-contingent monetary policy, the central bank provides partial consumption insurance to households, as discussed above.

"Bygones are bygones". As established above, the optimal response to autoregressive shocks to TFP and government spending implies strict price stability. In the case of

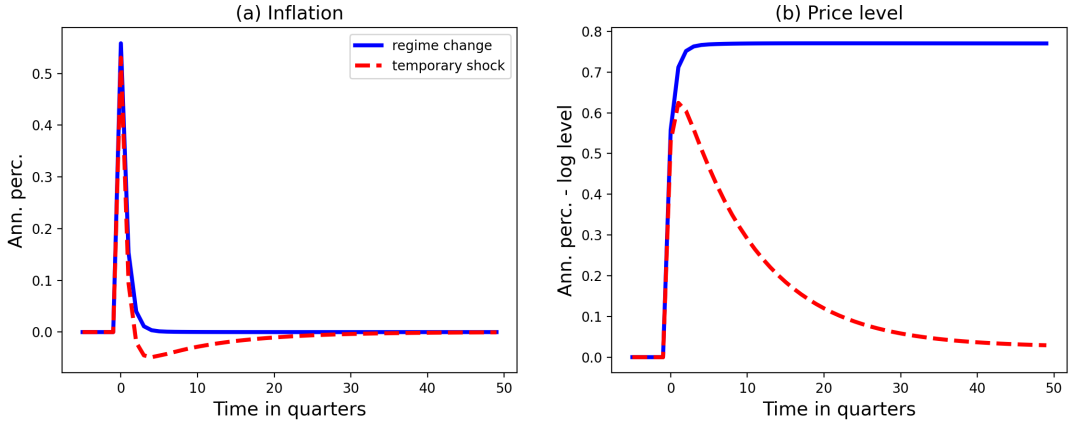

Figure 3: Comparison between a persistent (Markov-switching) shock and a temporary shock. *Note*: The figure shows (i) the transition from the normal-times regime to the bad-times regime under the optimal policy with commitment in the baseline case (solid blue line); (ii) and the dynamics following a temporary cost-push shock calibrated to generate the same initial inflation as the persistent shock (dashed red line). In both cases, the economy starts from the SSS of the normal-times regime. Innovations to the other temporary shocks are set to zero.

autoregressive cost-push shocks, the central bank leans against the wind. This differs from the dynamics in the flexible-price equilibrium, where inflation is always zero. By tolerating the rise in inflation, the central bank aims to minimize the distortions associated with the temporary increase in markups, albeit at the cost of higher price dispersion. When the central bank can commit, it ensures that the initial rise in inflation is later offset by a deflationary period, bringing the price level back to its pre-shock position. The dashed red line in Figure 3 shows the optimal response to a temporary AR(1) shock calibrated to produce the same inflation on impact as the switching (persistent) shock (solid blue line). The optimal policy exhibits this price-level targeting feature, as the price level returns to its initial value.

The response to an autoregressive cost-push shock exhibits several similarities to the response to a regime change in the mean of the same shock, illustrated by the blue lines in Figure 2. In both cases, the central bank cushions the impact on the output gap by allowing inflation to deviate temporarily from zero. The key difference between the two scenarios is that, under a regime change, no deflationary phase follows, and the price level therefore rises permanently after the shock (absent any new transition). In this case, "bygones are bygones,"

and the central bank refrains from any price-level targeting, as shown by the comparison of the two price levels in panel (b) of Figure 3.

To provide intuition for this result, Figure S.3 in the Supplementary Material also reports the optimal response under commitment to an autoregressive cost-push shock with an autoregressive coefficient of 0.99, that is, much more persistent than in the baseline case. The figure shows that as temporary cost-push shocks become more persistent, the optimal policy produces a smaller but more prolonged deflationary phase, leading to persistent deviations of the price level from its pre-shock value. This pattern is consistent with the theoretical results derived from the log-linear approximation around the non-stochastic steady state commonly found in textbooks (e.g., Galí, 2008). In that framework, the optimal policy under commitment satisfies

$$\pi_t = -\frac{1}{\epsilon}\left(x_t - x_{t-1}\right),$$

implying that inflation moves in the opposite direction of changes in the output gap. As the shock becomes more persistent, the change in the output gap becomes smaller. In the limit, x_t drops on impact by $-\Delta$ and then remains constant, implying that inflation rises on impact by Δ/ϵ and subsequently stays at zero. Although this reasoning is local in nature, it provides a useful approximation to the global dynamics observed in response to a regime change, as shown in Figure 2.

Persistent shocks or regimes? Given the previous result, one may wonder whether a continuous state variable, for instance a combination of two autoregressive shocks with one of them extremely persistent, can replicate the dynamics of a discrete Markov-switching model. The short answer is no. The model without regimes fails to reproduce key features such as the bimodality of the output gap and real interest rates, as well as the pronounced skewness in price dispersion associated with shocks during bad times. This is illustrated in Figure S.4 in the Supplementary Material, which compares the ergodic distribution in the baseline model (blue bars) with that of a counterfactual specification that replaces the Markov transition

with a highly persistent AR(1) process calibrated to the same mean and variance.[17] The differences are evident at first glance. Inflation (panel a), for example, becomes substantially more volatile around zero. More importantly, the bimodality of output gap (panel b) and the real interest rate (panel c) disappears almost entirely under very persistent supply shocks. We conclude that, although highly persistent autoregressive shocks can generate prolonged price-level deviations, they cannot capture the richer nonlinear dynamics induced by regime switching.

4.2.2 Discretion

Ergodic distribution and dynamics. We now turn to the case with discretion. Figure 4 shows the corresponding ergodic distributions. Inflation now clusters around two distinct points: 2.6 percent inflation in normal times and 3.8 percent inflation in bad times (panel a and Table 2). This difference arises from the "inflationary bias" originally identified by Kydland and Prescott (1977) and Barro and Gordon (1983). The bias results from the regime-dependent markup distortion from monopolistic competition, which widens or shrinks in each of the regimes. In bad times, the central bank has a larger incentive to surprise private agents by loosening monetary policy and generating inflation in an attempt to reduce the markup. However, the private sector anticipates this incentive and incorporates it into expectations, resulting in persistently positive inflation. Consequently, the economy exhibits a systematically higher level of inflation in the absence of commitment in both regimes, but especially in bad times.

Compared with the case under commitment, a discretionary central bank provides less macroeconomic insurance to households. This is illustrated in Figure 2, where the dashed red line shows the transition from normal to bad times under discretion. When the shock occurs, inflation jumps, overshooting 3.9 percent before gradually converging to that level. In contrast to the commitment case, the real interest rate does not decline significantly, reflecting

[17] Figure S.8 provides the equivalent figure for discretion (discussed next).

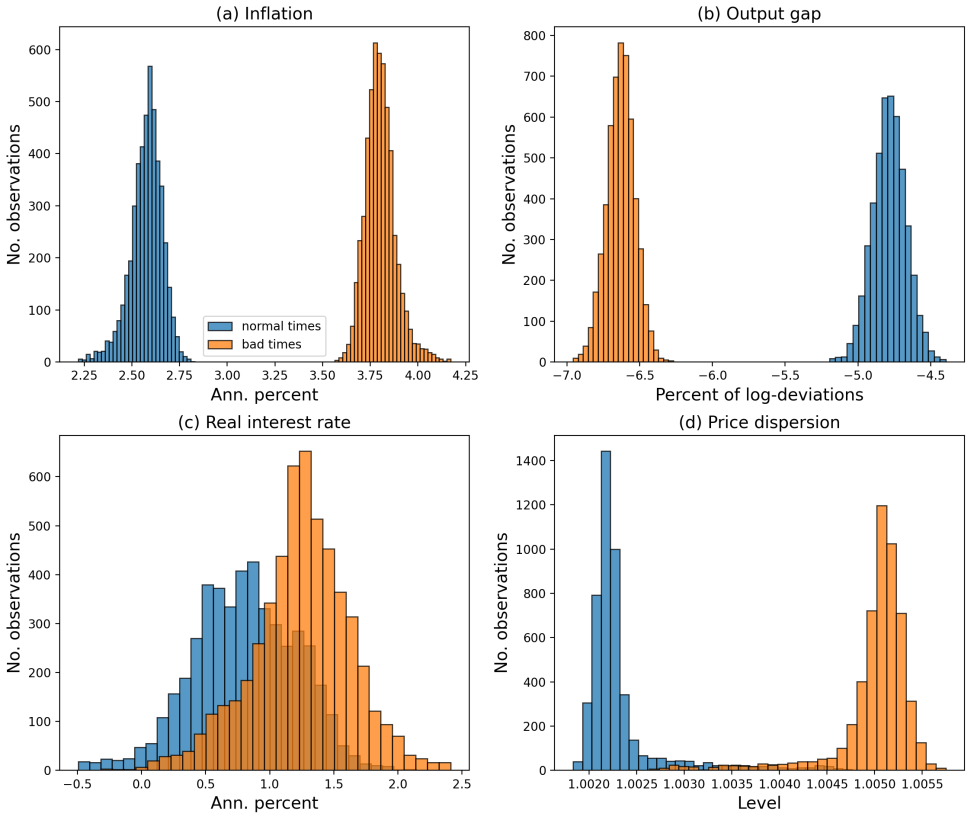

Figure 4: Ergodic distribution: discretion.

Note: The figure shows the ergodic distribution of the model under discretion. Colors indicate the two regimes: blue corresponds to normal times and orange to bad times. The distribution is obtained by simulating the model over a large number of periods.

the sharper drop in consumption and the output gap. As shown in Table 2, the mean output gap in bad times is more negative than that in the flexible-price economy. Households are therefore worse off: they experience greater markup (and labor) misallocation due to higher price dispersion, while receiving little additional macroeconomic insurance.

Afrouzi et al. (2023) analyze the transitional dynamics under discretion in a similar framework. They consider a shock that moves the economy from a deterministic steady state with low markups to another with high markups—an exercise analogous to a transition from normal to bad times. Their analysis is based on a one-off, unanticipated ("MIT") shock under perfect foresight, rather than the fully stochastic environment considered here. They similarly find that inflation overshoots before converging to the high-inflation steady state.

The key difference between the two analyses concerns the behavior of the real interest rate, which reflects different assumptions about the nature of risk. In Afrouzi et al. (2023), the real rate falls on impact, as the central bank initially stimulates the economy to cushion the shock, before gradually returning to its initial value, corresponding to the deterministic steady state. In our model, by contrast, the final real rate is above its initial level, as precautionary motives cause the real rate to rise on impact (see Section 3.1). Thus, in our model, the real interest rate instead converges to a higher, regime-dependent level, reflecting households' expected consumption dynamics and exposure to regime-switching risk.

Monetary policy rules. Appendix C in the Supplementary Material compares a standard and a modified Taylor rule in an economy with regime-switching supply shocks. Under the standard Taylor rule, the central bank sets interest rates in response to deviations of inflation from target but keeps a fixed long-run real-rate target. Because the natural rate of interest differs across regimes, policy becomes systematically too tight in normal times and too loose in bad times, generating persistent inflation below target in the former and above target in the latter, even though average inflation remains close to zero over the long run. The model shows that long-run inflation deviations arise when the policy rule fails to track changes in the natural rate. Motivated by this result, we consider a modified Taylor rule whose intercept adjusts to the regime-specific natural rate. This regime-contingent rule successfully centers inflation around its target while leaving within-regime fluctuations broadly similar. The key policy implication is that, in economies subject to changing supply conditions, central banks should adapt the intercept of their policy rule to the evolving natural rate rather than keeping it fixed.

Welfare. We assess the welfare cost of supply regimes by comparing household welfare in the baseline economy with welfare in a counterfactual economy without regimes, in which the regime-switching shock is set to zero. We consider monetary policy under discretion, holding all other parameters constant.[18] The welfare loss of supply regimes is 0.38 percent

[18]Welfare is computed over the respective ergodic distributions.

in consumption-equivalent terms; that is, households in an economy with regimes would be willing to forego 0.38 percent of their consumption to live in an economy without regimes.

We also analyze the gains from monetary policy commitment. They are larger in the economy with supply regimes. In the counterfactual economy without regimes the welfare gain from a policy under commitment is 0.17 percent (in consumption-equivalent terms of the baseline economy under discretion). This gain increases to 0.42 percent in the economy with regimes. This is because commitment improves the ability to cope both with AR(1) shocks and with regime transitions.

TFP supply regimes. So far we have concentrated in cost-push shocks. Nonetheless, our computational framework allows us to analyze any Markov chain. We also analyze the case in which the regimes are linked to TFP dynamics, a different supply shock. In this case, inflation remains constant across regimes, both under commitment and discretion. This result extends what happens in response to AR(1) shocks to Markov chain dynamics.

4.3 Extension: the case with endogenous pricing frequency

In Section 4.2, we showed that inflation remains persistently positive in both regimes when the central bank cannot commit. A potential concern is whether the assumption of a constant frequency of price changes, regardless of the inflation level, remains valid in an environment of persistently positive inflation. In such a setting, it would be reasonable to expect firms to adjust prices more frequently.

To address this concern, we replace the Calvo time-dependent pricing assumption with a mechanism *à la* Blanco et al. (2024). In that framework, firms sell a continuum of products and decide how many, but not which, prices to adjust each period, subject to an adjustment cost. Consistent with empirical evidence, the fraction of price changes evolves endogenously over time and increases with the level of inflation.

4.3.1 Model

A detailed description of the model is presented in Appendix D of the Supplementary Material. Here, we briefly highlight the two main differences relative to the baseline economy.

First, as in the standard New Keynesian model, there is a continuum of monopolistically competitive firms producing differentiated varieties indexed by j. In this setting, each firm produces a continuum of products k according to the technology

$$y_{jkt} = A_t l_{jkt}.$$

Each firm chooses the fraction of prices, n_{jt}, across its products to reset each period, as well as the new reset price P^*_{jt}. Because all products are symmetric, $P_{jkt} = P^*_{jt}$ for all products whose prices are reset.

Second, although each firm can choose *how many* prices to change in a given period, it cannot choose *which* prices to change. We assume that adjusting prices entails a convex cost, denominated in units of labor, given by

$$\frac{\xi}{2} \left(\max\{0, n_{jt} - \bar{n}\}\right)^2 w_t. \tag{23}$$

Here, $\bar{n}$ denotes the fraction of prices that a firm can adjust at no cost.[19]

We assume a symmetric equilibrium in which all firms are identical, and thus $n_{jt} = n_t$ and $P^*_{jt} = P^*_t$. Product prices are selected randomly for adjustment, so a fraction $1 - n_t$ remains unchanged, and a fraction n_t is reset to P^*_t:

$$P_t = \left[(1 - n_t) \left(P_{t-1}\right)^{1-\epsilon} + n_t \left(P^*_t\right)^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}}.$$

The resulting model can be viewed as an extension of the baseline economy in which the repricing frequency is n_t instead of $1 - \theta$. The frequency evolves endogenously with

[19] The cost is transferred back to households.

macroeconomic conditions, as firms trade off the benefits of updating a larger share of prices each period against the convex menu costs (23). We calibrate $\bar{n} = 1 - \theta = 0.25$ and $\xi = 1.767$, following Blanco et al. (2024). The baseline economy, that is, the model with Calvo pricing, thus emerges as a limiting case of this model when $\xi \to \infty$.

4.3.2 Results

To assess the effects of endogenous pricing in response to regime changes, we extend the experiment shown in Figure 2. That figure illustrates the transitional dynamics following a regime change, assuming that the economy starts from the SSS of the normal-times regime and transitions to the bad-times regime at time zero, with all other shocks set to zero. In Figure S.7 (Supplementary Material), we repeat this exercise under endogenous-frequency pricing, considering different magnitudes of the cost-push shock in bad times. The optimal policy is computed under discretion.

Two results emerge. First, for the baseline mean of the cost-push shock ($\bar{\eta} = \bar{\eta}_1 = \bar{\eta}_2$), the responses under time-dependent (Calvo, solid black line) and endogenous pricing frequency (dotted red line) are nearly identical. The magnitudes of inflation, and the output gap, are similar in both cases, confirming that our previous results are robust to the choice of pricing mechanism.

Second, the effects of a regime change do not scale linearly with the magnitude of the mean cost-push shock. To illustrate this, we recompute the transition dynamics for several alternative values of $\bar{\eta}$: $2\bar{\eta}$ (dash-dotted blue line), $3\bar{\eta}$ (wide-dashed magenta line), and $0.5\bar{\eta}$ (dashed green line). Figure S.7 shows that as the shock magnitude increases, the responses of the different variables scale less than linearly: doubling the shock size does not imply that the response of macro variables also doubles. This pattern reflects the nonlinearity of the Phillips curve under endogenous pricing frequency.

To better understand this result, Figure 5 compares the SSS in bad times for different values of the mean cost-push shock ($\bar{\eta}$). To this end, we solve the model over a grid of $\bar{\eta}$

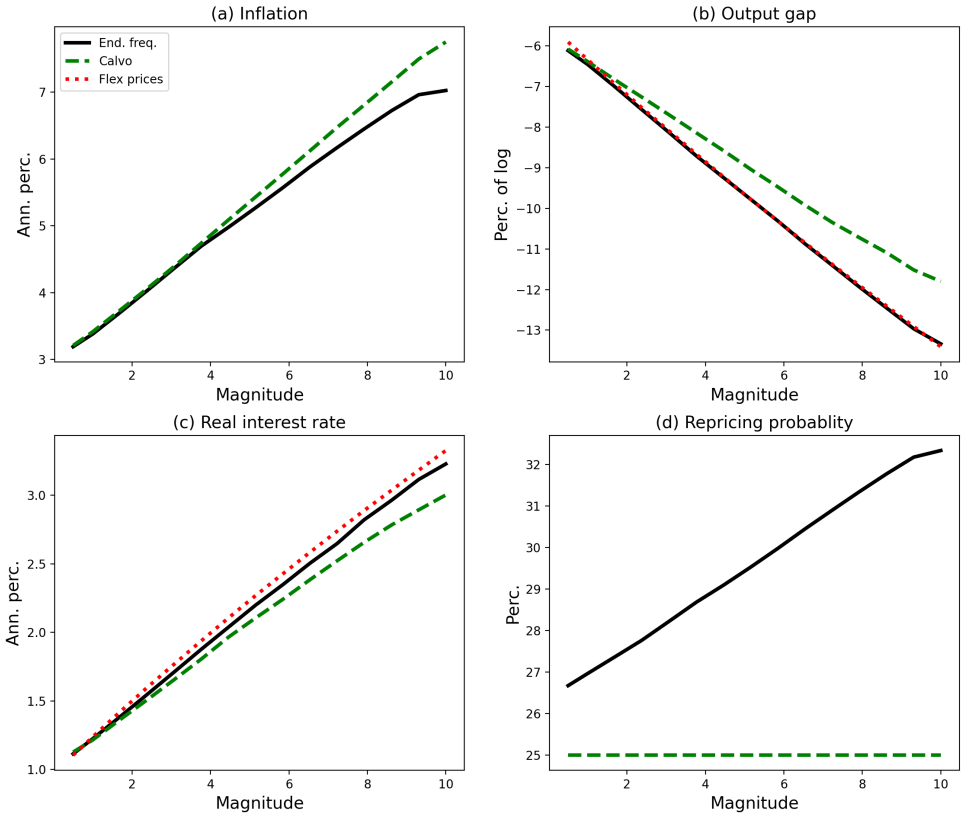

Figure 5: SSS in bad times for different means of the cost-push shock ($\bar{\eta}$).
Note: The figure shows the SSS in the bad-times regime under three pricing assumptions: time-dependent (Calvo) pricing (dashed green line), endogenous pricing frequency (solid black line), and flexible pricing (dotted red line).

values under three pricing assumptions: time-dependent (Calvo) pricing, endogenous pricing frequency, and flexible prices.[20] Under flexible pricing (dotted red line), the output gap declines approximately (log-)linearly with the magnitude of the shock, following the SSS relationship in equation (21) (panel b):

$$\omega \log\left(c^*_{ss} + g_{ss}\right) + \gamma \log\left(c^*_{ss}\right) = (\omega + 1) \log A_{ss} - \log \mathcal{M} - \log(1 + \tau_{ss}), \tag{24}$$

and the real interest rate rises (panel c).

In the baseline case with Calvo pricing (dashed green line), the decline in the output gap

[20]This can be done efficiently using deep learning techniques, treating $\bar{\eta}$ as an additional variable in the problem, as discussed by Chen et al. (2026).

as the shock increases is less pronounced than under flexible prices. This reflects the fact that the nonlinear Phillips curve is not completely vertical: higher steady-state inflation (panel a) affects average markups, partly cushioning the negative impact of the cost-push shock.

Under endogenous pricing (solid black line), the frequency of price changes increases with the shock magnitude (panel d), making prices more flexible as firms adjust them more frequently. This greater flexibility weakens the transmission of monetary policy to real variables—that is, the Phillips curve becomes steeper than under Calvo pricing. As a result, the central bank has less scope to cushion the cost-push shock. This is evident from panel b: the SSS output gap under endogenous pricing is similar to that under flexible prices. As the central bank's ability to stimulate the economy diminishes, so does its incentive to do so. Private agents correctly anticipate that the central bank is less inclined to "run the economy hot" in response to a regime transition, and they adjust their inflation expectations accordingly. Consequently, SSS inflation increases proportionally less as the shock magnitude rises (panel a).

This mechanism complements the "strike while the iron is hot" effect identified by Karadi et al. (2024). In their framework, the central bank tightens monetary policy more aggressively under state-dependent pricing than under Calvo pricing in response to a large temporary cost-push shock, exploiting the temporary steepening of the Phillips curve as the repricing frequency increases. Here, this mechanism is also present (panel c), reinforcing the change in inflation expectations driven by differing central bank incentives.

5 Conclusions

This paper advances the understanding of optimal monetary policy in the presence of regime-switching supply shocks. We develop a New Keynesian model that captures the economy's transition between "normal times" and "bad times," with the latter characterized by persistently higher production costs, such as those arising from commodity supercycles,

tariffs, wars, or geopolitical fragmentation.

References

Adam, K., Weber, H., 2019. Optimal trend inflation. American Economic Review 109, 702–37.

Adjemian, S., Bastani, H., Juillard, M., Mihoubi, F., Perendia, G., Ratto, M., Villemot, S., 2011. Dynare: Reference manual, version 4 .

Afrouzi, H., Halac, M., Rogoff, K.S., Yared, P., 2023. Monetary Policy without Commitment. NBER Working Papers 31207. National Bureau of Economic Research, Inc.

Albanesi, S., Chari, V.V., Christiano, L.J., 2003. Expectation Traps and Monetary Policy. The Review of Economic Studies 70, 715–741.

Arellano, C., Bai, Y., Mihalache, G.P., 2020. Monetary Policy and Sovereign Risk in Emerging Economies (NK-Default). NBER Working Papers 26671. National Bureau of Economic Research, Inc.

Azinovic, M., Gaegauf, L., Scheidegger, S., 2022. Deep equilibrium nets. International Economic Review 63, 1471–1525.

Azinovic, M., Žemlička, J., 2024. Intergenerational consequences of rare disasters.

Barro, R.J., Gordon, D.B., 1983. Rules, discretion and reputation in a model of monetary policy. Journal of Monetary Economics 12, 101–121.

Bayer, C., Born, B., Luetticke, R., 2023. The liquidity channel of fiscal policy. Journal of Monetary Economics 134, 86–117.

Bellman, R., 1957. Dynamic Programming. Princeton University Press, Princeton, NJ. Introduction of the curse of dimensionality in optimization and decision processes.

Benigno, P., Rossi, L., 2021. Asymmetries in monetary policy. European Economic Review 140.

Benigno, P., Woodford, M., 2005. Inflation Stabilization And Welfare: The Case Of A Distorted Steady State. Journal of the European Economic Association 3, 1185–1236.

Bergin, P.R., Corsetti, G., 2023. The macroeconomic stabilization of tariff shocks: What is the optimal monetary response? Journal of International Economics 143, 103758.

Bianchi, F., Melosi, L., 2017. Escaping the great recession. American Economic Review 107, 1030–58.

Bianchi, F., Melosi, L., Rottner, M., 2021. Hitting the elusive inflation target. Journal of Monetary Economics 124, 107–122.

Bianchi, J., Coulibaly, L., 2025. The Optimal Monetary Policy Response to Tariffs. NBER Working Papers 33560. National Bureau of Economic Research, Inc.

Blake, A.P., Zampolli, F., 2011. Optimal policy in markov-switching rational expectations models. Journal of Economic Dynamics and Control 35, 1626–1651.

Blanchard, O., Galí, J., 2007. Real Wage Rigidities and the New Keynesian Model. Journal of Money, Credit and Banking 39, 35–65.

Blanco, A., 2021. Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound. American Economic Journal: Macroeconomics 13, 108–141.

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

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

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

Caratelli, D., Halperin, B., 2024. Optimal monetary policy under menu costs.

Cesa-Bianchi, A., Harrison, R., Sajedi, R., 2022. Decomposing the drivers of Global R. Bank of England working papers 990. Bank of England.

Chen, H., Didisheim, A., Scheidegger, S., 2026. Deep surrogates for finance: With an application to option pricing. Journal of Financial Economics 177, 104222. URL: https://www.sciencedirect.com/science/article/pii/S0304405X25002302, doi:https://doi.org/10.1016/j.jfineco.2025.104222.

Choi, J., Foerster, A., 2021. Optimal monetary policy regime switches. Review of Economic Dynamics 42, 333–346.

Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: A new keynesian perspective. Journal of Economic Literature 37, 1661–1707.

Coibion, O., Gorodnichenko, Y., Wieland, J., 2012. The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound? The Review of Economic Studies 79, 1371–1406.

Davig, T., 2016. Phillips curve instability and optimal monetary policy. Journal of Money, Credit and Banking 48, 233–246.

Davig, T., Doh, T., 2014. Monetary Policy Regime Shifts and Inflation Persistence. The Review of Economics and Statistics 96, 862–875.

Debortoli, D., Nunes, R., 2014. Monetary regime switches and central bank preferences. Journal of Money, Credit and Banking 46, 1591–1626.

Del Negro, M., Giannone, D., Giannoni, M.P., Tambalotti, A., 2017. Safety, Liquidity, and the Natural Rate of Interest. Brookings Papers on Economic Activity 48, 235–316.

Federle, J., Meier, A., Müller, G., Schularick, M., 2024. The Price of War. CEPR Discussion Papers 18834. C.E.P.R. Discussion Papers.

Fernandez, A., Schmitt-Grohe, S., Uribe, M., 2020. Does the Commodity Super Cycle Matter? NBER Working Papers 27589. National Bureau of Economic Research, Inc.

Fernandez-Villaverde, J., Mineyama, T., Song, D., 2024a. Are We Fragmented Yet? Measuring Geopolitical Fragmentation and Its Causal Effects. CEPR Discussion Papers 19184. C.E.P.R. Discussion Papers.

Fernandez-Villaverde, J., Nuno, G., Perla, J., 2024b. Taming the Curse of Dimensionality: Quantitative Economics with Machine Learning. Papers.

Fernández-Villaverde, J., Marbet, J., Nuño, G., Rachedi, O., 2024. Inequality and the zero lower bound. Journal of Econometrics , 105819.

Friedl, A., Kübler, F., Scheidegger, S., Usui, T., 2023. Deep uncertainty quantification: With an application to integrated assessment models. Working Paper.

Gagnon, E., Johannsen, B.K., López-Salido, D., 2021. Understanding the New Normal: The Role of Demographics. IMF Economic Review 69, 357–390.

Galí, J., 2008. Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework. Princeton University Press.

Gu, Z., Lauriere, M., Merkel, S., Payne, J., 2024. Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models. Papers 2406.13726. arXiv.org.

Han, J., Yang, Y., E, W., 2021. Deepham: A global solution method for heterogeneous agent models with aggregate shocks. arXiv preprint arXiv:2112.14377 .

Kahou, M.E., Fernandez-Villaverde, J., Gomez-Cardona, S., Perla, J., Rosa, J., 2024. Spooky Boundaries at a Distance: Inductive Bias, Dynamic Models, and Behavioral Macro. CESifo Working Paper Series 11292. CESifo. doi:None.

Kaplan, G., Nikolakoudis, G., Violante, G.L., 2023. Price Level and Inflation Dynamics in Heterogeneous Agent Economies. Technical Report. Princeton.

Karadi, P., Nakov, A., Nuno, G., Pasten, E., Thaler, D., 2024. Strike while the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve. CEPR Discussion Papers 19339. C.E.P.R. Discussion Papers.

Kase, H., Melosi, L., Rottner, M., 2025. Estimating nonlinear heterogeneous agent models with neural networks. BIS Working Papers 1241. Bank for International Settlements. doi:None.

King, R.G., Wolman, A.L., 2004. Monetary Discretion, Pricing Complementarity, and Dynamic Multiple Equilibria. The Quarterly Journal of Economics 119, 1513–1553.

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

Kydland, F.E., Prescott, E.C., 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. Journal of Political Economy 85, 473–491.

Maliar, L., Maliar, S., Winant, P., 2021. Deep learning for solving dynamic economic models. Journal of Monetary Economics 122, 76–101.

Mian, A., Straub, L., Sufi, A., 2021. Indebted demand Quarterly Journal of Economics, 136 (4) 2021: 2243-2307.

Monacelli, T., 2025. Tariffs and Monetary Policy. CEPR Working Papers 20142. CEPR.

Nakov, A., Thomas, C., 2014. Optimal Monetary Policy with State-Dependent Pricing. International Journal of Central Banking 36.

729 Obstfeld, M., 2023. Natural and Neutral Real Interest Rates: Past and Future. NBER
730 Working Papers 31949. National Bureau of Economic Research, Inc.

731 Rachel, L., Summers, L.H., 2019. On Secular Stagnation in the Industrialized World.
732 Brookings Papers on Economic Activity 50, 1–76.

733 Rios-Rull, J.V., Ritto, J., Elkina, M., 2026. The Reach of Monetary Policy and the Inflation
734 Bias. Mimeo.

735 Romei, F., Cesa-Bianchi, A., de Ferra, S., Ferrero, A., Kohlhas, A., McMahon, M., Rosso, G.,
736 2025. Monopsony, Income Risk and R* Multiplicity. CEPR Working Papers 20582. CEPR.

737 Sahuc, J.G., Smets, F., Vermandel, G., 2023. The new keynesian climate model .

738 Scheidegger, S., 2026. Deep learning for solving and estimating dynamic models in economics
739 and finance. `arXiv:2605.14493`.

740 Schorfheide, F., 2005. Learning and Monetary Policy Shifts. Review of Economic Dynamics
741 8, 392–419.

742 Werning, I., Lorenzoni, G., Guerrieri, V., 2025. Tariffs as Cost-Push Shocks: Implications for
743 Optimal Monetary Policy. NBER Working Papers 33772. National Bureau of Economic
744 Research, Inc.

745 Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy.
746 Princeton University Press.

747 Yun, T., 2005. Optimal Monetary Policy with Relative Price Distortions. American Economic
748 Review 95, 89–109.

749 Zandweghe, W.V., Wolman, A.L., 2019. Discretionary monetary policy in the Calvo model.
750 Quantitative Economics 10, 387–418.