# Monetary Policy with Supply Regimes*

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

[1]Banco de España, CEMFI, CEPR
[2]University of Lancaster
[3]University of Lausanne

August 1, 2025

**Abstract**

This paper studies monetary policy in a New Keynesian model featuring supply regimes, that is, periods of sustained increases in production costs arising from tariffs, wars, or geopolitical fragmentation. We compare optimal policy responses under commitment and discretion. Under commitment, the optimal strategy is to lean against the wind without compensating for past inflation, implying that "bygones are bygones." Under discretion, however, a persistent inflationary bias emerges in regimes with high production costs. Furthermore, we show that standard Taylor rules fail to stabilize long-term inflation due to endogenous shifts in the natural real interest rate.

*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 was visiting the BIS. We are grateful to Ryan Banerjee, Florin Bilbiie, Eduardo Davila, Davide Debortoli, Jose Elias Gallegos, Boris Hofmann, Joël Marbet, Leonardo Melosi, Dominik Thaler, Mathias Trabandt, and participants at different seminars and conferences for 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 project ID "New methods for asset pricing with frictions". The views expressed in this manuscript are those of the authors and do not necessarily represent the views of the Banco de España, or the Eurosystem.

*"Behold, there come seven years of great plenty throughout all the land of Egypt: And there shall arise after them seven years of famine; and all the plenty shall be forgotten in the land of Egypt; and the famine shall consume the land".*

---

*Genesis 41:29-30*

# 1 Introduction

The design of optimal monetary policy remains a central focus in macroeconomics. The recent inflationary period (e.g., post-2021) and the increase in tariffs have prompted academic researchers and central banks to re-evaluate the appropriate responses of monetary policy when supply-side factors drive inflation.[1] A critical question emerging from this debate is how monetary policy should be formulated in the face of supply regimes.[2] By supply regimes, we refer to sustained increases in production costs due to factors such as tariffs, wars, or geopolitical fragmentation, which can persist for years or even decades.[3] These supply regimes challenge the validity of traditional monetary policy prescriptions, typically derived under the assumption of small temporary shocks.

This paper investigates monetary policy responses to persistent supply disruptions by extending the standard New Keynesian framework (Clarida et al., 1999, Woodford, 2003) to incorporate supply regimes, that is, periods of sustained increases in production costs. In this economy, a representative household consumes a continuum of differentiated goods and supplies labor in a centralized, frictionless market. Each good is produced by a single firm using labor as

[1]See Bandera et al. (2023) or Mankiw (2024).

[2]See, for instance, Powell (2023), Schnabel (2024) or Maechler (2024).

[3]Federle et al. (2024) find that the macroeconomic effect of war on nearby countries' output remains substantial 8 years after the outbreak. Fernandez-Villaverde et al. (2024a) find that while geopolitical fragmentation decreased in the "globalization era"after the collapse of the Soviet Union, fragmentation increased substantially after the 2007-2008 financial crisis and remains at high values.

the sole input. The economy experiences aggregate productivity shocks, cost-push shocks, and government spending shocks, all modeled as standard zero-mean autoregressive processes. The model incorporates nominal price rigidities *à la* Calvo (Calvo, 1983), where firms adjust prices only with a certain probability; otherwise, prices remain unchanged. The central bank conducts monetary policy through adjustments in short-term nominal interest rates, while a labor subsidy offsets the markup distortion inherent in monopolistic competition. We analyze optimal monetary policy under both discretion and commitment, and also consider Taylor rules.

The novelty of our model is the introduction of supply regimes that shift 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 zero, and "bad times," when it is positive, reflecting a uniform increase in firms' production costs. This mechanism serves as a proxy for various cost-increasing shocks, such as tariffs.[4]

This feature complicates the model's solution, as it necessitates a global solution method that accounts for regime-switching dynamics.[5] We tackle this challenge by introducing a novel deep learning algorithm. Specifically, we extend the "deep equilibrium nets" methodology of Azinovic et al. (2022) to accommodate Markov-switching dynamics. This methodological innovation allows us to accurately approximate the model's nonlinear functions and efficiently compute its solution, thereby enabling analysis of optimal monetary policy under different regimes.

Our contributions are fourfold: First, we uncover a new precautionary-savings mechanism through which the supply regime affects the natural rate. During normal times, the economy is undistorted, and consumption coincides with the efficient allocation; hence, the output gap is zero. In contrast, during bad times, the average markup becomes suboptimal, reducing output and consumption and creating a negative output gap. Consequently, the economy exhibits two distinct stochastic steady states (SSSs), each corresponding to one of the regimes.[6] The difference

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

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

[6]An SSS refers to the equilibrium state when all shocks are zero, the regime is fixed, and agents correctly anticipate the stochastic process.

in consumption across the two regimes drives the dynamics of the natural interest rate.[7] In normal times, households anticipate a possible shift to a regime with lower average consumption and thus increase precautionary saving. Given a fixed supply of savings instruments, 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), allowing inflation to rise when the persistent shock occurs but then gradually tightening monetary policy so that inflation returns to zero after a few quarters. The central bank thereby partially cushions the shock'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. When the supply regime changes, however, no such deflationary episode is required, and the price level rises permanently—"bygones are bygones." This result shows that the price-level-targeting implications of standard policy prescriptions hold only under autoregressive shocks and do not extend to frameworks in which shocks follow a Markov chain.

Third, if the central bank cannot commit to future actions but understands how it can affect future decisions by changing the price distribution, we find that the optimal monetary policy under discretion displays a state-contingent inflationary bias (Kydland and Prescott, 1977, Barro and Gordon, 1983): while in normal times, long-term inflation is centered around zero, inflation surges when the persistent shock arrives. The persistent supply shock distorts the SSS, creating an incentive for the central bank to loosen monetary policy to reduce the average markup. Private agents anticipate this reaction, increasing their prices, which results in persistently positive inflation during bad times.

Finally, we show that traditional interest-rate rules cannot stabilize inflation in either regime because of shifts in the natural interest rate. In particular, a Taylor rule that targets the long-term

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

real rate at the average natural rate across regimes is excessively tight in normal times and overly loose in bad times. As a result, inflation systematically deviates from the target, highlighting the shortcomings of policy rules premised on a constant natural rate.

**Related literature.** This paper relates to five 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 (e.g., Albanesi et al., 2003; King and Wolman, 2004; Zandweghe and Wolman, 2019; Arellano et al., 2020; or Afrouzi et al., 2023). 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 persistent shocks.

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 stochastic steady states under optimal policy, with the economy transitioning among them in response to supply regimes.[8]

Third, while most of the literature has focused on how the natural rate depends on structural variables, such as total factor productivity (TFP) growth or demographics (Cesa-Bianchi et al., 2022; Gagnon et al., 2021; Del Negro et al., 2017),[9] 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, our paper extends the emerging literature on the use of deep learning to solve high-

[8]There are exceptions, such as Choi and Foerster (2021), that analyze regime-switching models globally, but to the best of our knowledge none has attempted to analyze optimal policy due to the curse of dimensionality.

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

dimensional general equilibrium models (e.g., Maliar et al., 2021; Han et al., 2021; Friedl et al., 2023; Gu et al., 2024).[10] Our paper extends 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.

Fifth, the 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 remainder of the paper is structured as follows. Section 2 introduces the model, and Section 4 defines the optimal monetary policy problems under discretion and commitment. Section 3 examines the efficient allocation, the flexible price equilibrium, and discusses why natural rates inherit the regime-switching nature of supply shocks. Section 5 details the calibration and our deep learning-based solution method. Section 6 discusses the main results regarding optimal monetary policy design, and Section 7 explores the case of Taylor rules. Finally, Section 8 concludes.

# 2 Model

In this section, we present the formal structure of our model. The model is a standard New Keynesian model (Galí, 2008) expanded to include regime-switching cost-push shocks. Time $t$ is discrete, and there are three types of agents: households, firms, and the central bank.

## 2.1 Households

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

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

subject to the budget constraint:

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

[10] See Fernandez-Villaverde et al. (2024b) for a recent review.

where $B_t$ are holdings of a nominal bond that pays interest $1 + i_t$, $w_t$ is the real wage, $p_t$ is the price level, and $T_t$ are the profits from monopolistic producers.

Under these assumptions, the optimal household choices of consumption, labor, and bonds, as well as the first-order conditions, are given by:

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

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

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

where $\lambda_t$ is the Lagrange multiplier associated with the budget constraint, 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 unit measure of differentiated retail goods indexed by $j$. Cost-minimization implies that the optimal demand for the $j$-th variety of good:

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

where $y_t$ is total output of the final good, $\epsilon$ is the elasticity of substitution, and $p_t$ is defined as

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

and $p_t(j)$ is the nominal price of $j$-th variety,

## 2.3 Retail good producers

There is a continuum of monopolistic firms, each of them producing a variety $j$. Each firm uses labor to produce the good according to the technology $y_t(j) = A_t h_t(j)$, where $A_t$ is the stochastic total factor productivity. There is a labor subsidy $\bar{\tau} = \frac{1}{\epsilon}$ to correct the distortions associated with monopolistic competition. Firms face a temporary cost-push shock $\xi_t$ as well as a regime-switching one $\eta_t$ which controls the supply regime, as explained below.

We define the labor wedge as

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

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

Each firm has monopoly power over its variety and takes the demand for that variety, $c_t(j)$, 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 $\Psi\left(y_{t+k}(j)\right) \equiv \left(1+\tau_{t+k}\right) w_{t+k} \left(\frac{y_{t+k}(j)}{A_{t+k}}\right)$ 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 $t+k$. The rest of the firms maintain prices constant, that is, $p_{t+k}(j) = p_t(j)$. The optimization problem can thus 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 optimality condition is given by:

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

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 firms are identical, 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 4, we define the optimal monetary policy problem both under discretion and commitment.

## 2.5 Market clearing conditions

The market-clearing condition for goods is given by:

$$y_t = c_t + g_t,$$

where $g_t$ is government spending, which follows a stochastic process.

Since all firms face the same probability $\theta$ of keeping prices fixed, the law of large numbers ensures that a fraction $\theta$ 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

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

The market-clearing condition for labor is given by:

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

We define price dispersion as:

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

such that the aggregate production function becomes:

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

## 2.6 Shocks and regimes

We define the temporary shocks and the regimes.

**Temporary shocks**. We consider TFP, government spending, and cost-push shocks, each following an AR(1) process. 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:[11]

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

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

**Regimes.** We assume that the shock $\eta_t$ evolves according to a two-state Markov chain. This is equivalent to assuming that the mean of the temporary AR(1) cost-push shock $\xi_t$ is regime-switching. We consider two supply regimes. In regime 1, "normal times", the value of $\eta_t$ is zero.

[11] The constant term is introduced to guarantee that the ergodic mean of the log-normal variable is one.

In regime 2, "bad times", its value is $\bar{\eta} > 0$. This implies that the shock is only active during bad times. The transition probabilities are $p_{12}$ from regime 1 to 2:

$$p_{12} = \mathbb{P}\left(\eta_t = \bar{\eta} \mid \eta_{t-1} = 0\right), \tag{11}$$

and $p_{21}$ from 2 to 1:

$$p_{21} = \mathbb{P}\left(\eta_t = 0 \mid \eta_{t-1} = \bar{\eta}\right). \tag{12}$$

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

# 3 Regime-based natural rates

## 3.1 Efficient allocation and flexible-price equilibrium

**Efficient allocation.** We begin by analyzing the efficient allocation in the model, which is the allocation produced by a social planner maximizing household welfare subject to 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, that is,

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

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

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

[12] We denote variables in the efficient allocation with a hat.

Notice how the efficient allocation is influenced by TFP and government expenditure shocks, but not by cost-push shocks. The value of 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 in the baseline model and that in the efficient allocation:

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

**Flexible-price equilibrium.** We consider the counterfactual equilibrium with flexible prices, that is, when $\theta = 0$. The optimal relative reset prices and price dispersion remain at one, $p_t^* = \Delta_t = 1$, reflecting that individual prices are always optimal. This setup creates a potential wedge between the marginal rate of substitution between consumption and labor and the marginal rate of transformation, as follows:[13]

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

Here, $(1 + \tau_t\mathcal{M})$ has a mean of one in the normal-times regime, but a mean of $(1 + \bar{\eta}\mathcal{M})$ in the bad-times regime. Consumption in this case satisfies the equation

$$c_t^* + g_t = y_t^* = A_t h_t^* = A_t^{\frac{1}{\omega}+1}((1+\tau_t)\,\mathcal{M})^{-\frac{1}{\omega}} c_t^{*-\frac{\gamma}{\omega}}. \tag{13}$$

This can be rewritten as

$$\left(\frac{c_t^* + g_t}{A_t}\right)^{\omega} - \frac{A_t}{(1+\tau_t)\mathcal{M}c_t^{*\gamma}} = 0. \tag{14}$$

Notice that consumption now depends on the cost-push shock. We define a *stochastic steady state* (SSS) in this economy as a steady state when the innovations of the shocks are zero, and the temporary shocks remain at their mean values. 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 two different values, the economy exhibits two SSSs, one in each regime.

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

### 3.2 Natural rates

The real interest rate in the flexible-price economy 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

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

where the notation $z_{n,t}$ denotes variable $z$ at time $t$ and regime $n = \{1, 2\}$. The real rate in the SSS of regime 1 thus approximately satisfies

$$1 + r_{1,ss}^* \approx \frac{1}{\beta} \frac{{c_{2,ss}^*}^\gamma}{\left( p_{12} {c_{1,ss}^*}^\gamma + (1 - p_{12}) {c_{2,ss}^*}^\gamma \right)}, \tag{15}$$

reflecting that 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 rate in the flexible-price economy as the *natural rate*, similar to the definition in Obstfeld (2023): the "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 (13):

$$\left( c_{n,ss}^* + \bar{g} \right)^\omega - \frac{1}{(1 + \eta_n \mathcal{M}) {c_{n,ss}^*}^\gamma} = 0.$$

where $(1 + \eta_n \mathcal{M})$ equals one in regime 1, and $(1 + \bar{\eta} \mathcal{M})$ in regime 2. Due to the markup distortion in the bad state (state 2), we have $c_{2,ss}^* < c_{1,ss}^*$. Therefore, the denominator on the right-hand side of equation (15) exceeds the numerator $(c_{2,ss}^*)^\gamma$. This implies that the natural rate in regime 1, $r_{1,ss}^*$, is *lower* than that in the efficient allocation, $1/\beta$. Conversely, the natural rate in regime 2, $r_{2,ss}^*$, is *higher* than that in the efficient allocation. In our calibration, these values are 0 percent and 2.7 percent, respectively, compared to 1 percent in the efficient allocation, as discussed below.

The differences in the natural rate as a result of the cost-push shock regime are driven by a *precautionary-savings* motive by households. 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 given the fixed supply of government debt, the only asset in this economy, their increased demand for savings merely leads to a fall in the bond return, that is, in the natural rate. During bad times, households are forced to reduce their savings to smooth consumption, which results in a higher natural rate.

# 4 Optimal monetary policy: problem statement

We turn next to define the optimal monetary policy under both discretion and commitment.

## 4.1 Discretion

First, consider the case in which the central bank maximizes household welfare under discretion. The central bank value function is $V(\Delta, A, \tau, g, n)$, where price dispersion $\Delta_t$ is the endogenous state variable, $\{A_t, \tau_t, g_t\}$ is a vector of TFP, cost-push, and government spending shocks, and $n_t$ is the regime.[14] This problem, therefore, includes one endogenous state variable, three exogenous variables, and a regime change. The value function satisfies the Bellman equation, that is,

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

[14] We use $\tau_t$ instead of $\xi_t$ to simplify notation.

subject to the equilibrium conditions, that is,

$$c_t^{-\gamma} = h_t^{\omega}/w_t, \tag{16}$$

$$1 = \theta\left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}, \tag{17}$$

$$\Delta_t = \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}, \tag{18}$$

$$y_t = A_t h_t\left(\Delta_t\right)^{-1}, \tag{19}$$

$$y_t = c_t + g_t, \tag{20}$$

$$p_t^* = \mathcal{M}\frac{y_t w_t\left(1+\tau_t\right)\left(A_t\right)^{-1} + \mathbb{E}_t\left[\theta\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right]}{y_t + \mathbb{E}_t\left[\theta\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right]}. \tag{21}$$

Under discretion, the central bank cannot commit to future policy paths. 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.

## 4.2 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 (16)-(20) and the constraints

$$p_t^*\Xi_t^D = \mathcal{M}\Xi_t^N,$$

$$\Xi_t^N = (c_t+g_t)^{1+\omega}\left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^{\gamma}\left(1+\tau_t\right)\left(A_t\right)^{-1} + \mathbb{E}_t\left[\beta\theta c_t^{\gamma}c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right], \tag{22}$$

$$\Xi_t^D = (c_t+g_t) + \mathbb{E}_t\left[\beta\theta c_t^{\gamma}c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right]. \tag{23}$$

The first-order conditions can be found in Appendix A.3. There, it can be seen that a recursive solution to the Ramsey problem is a function of seven states: the five states already present in the problem under discretion, plus the two backward-looking Lagrange multipliers associated with equations (22)-(23).

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 (22) and (23), evolve around their SSS values.[15]

# 5 Calibration and numerical method

## 5.1 Calibration

The model outlined in Section 2 is calibrated at a quarterly frequency, and the parameters are reported in Table 1. The calibration relies as much as possible on standard values from the literature. Regarding preferences, the quarterly discount factor is 0.9975, implying a real interest rate of 1 percent in the deterministic steady state. The elasticity of substitution across products is $\epsilon = 7$, resulting in a frictionless net markup of 1/6. The inverse of the intertemporal elasticity of substitution, $\gamma$, is set to 2, and the inverse of the Frisch elasticity, $\omega$, is set to 1. The long-run productivity level, $A$, is normalized to one, and the government spending constant, $\bar{g}$, is set to 20 percent.

The parameters of the TFP, government spending, and cost-push shocks are taken from Coibion et al. (2012). The value of the Markov-switching cost-push shock in bad times is set to $\bar{\eta} = 1/\epsilon$, so that it fully offsets the optimal labor subsidy. The average duration of regime 1 ("normal times") is 48 quarters (12 years), to capture a period encompassing a business cycle, giving $p_{12} = 1/48$, while the average duration of regime 2 ("bad times") is set to half of that of normal times, 24 quarters (6 years), so $p_{21} = 1/24$. We later explore the sensitivity to alternative durations.

[15]This 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.

| Parameter | | Value |
|---|---|---|
| Long-run productivity level | $A$ | 1 |
| Inverse Frisch elasticity | $\omega$ | 1 |
| Inverse of intertemporal elasticity of substitution | $\gamma$ | 2 |
| Discount factor | $\beta$ | 0.9975 |
| Elasticity of substitution among varieties | $\epsilon$ | 7 |
| Government spending constant | $\bar{g}$ | 0.2 |
| Calvo constant | $\theta$ | 0.75 |
| Labor subsidy | $\bar{\tau}$ | $\frac{1}{\epsilon}$ |
| | | |
| Mean of cost-push shock in bad times | $\bar{\eta}$ | $\frac{1}{\epsilon}$ |
| Transition probability from normal to bad times | $p_{12}$ | 1/48 |
| Transition probability from bad to normal times | $p_{21}$ | 1/24 |
| Persistence of TFP shock | $\rho^A$ | 0.99 |
| Persistence of cost-push shock | $\rho^\tau$ | 0.90 |
| Persistence of government spending shock | $\rho^g$ | 0.97 |
| Standard deviation of TFP shock | $\sigma^A$ | 0.009 |
| Standard deviation of cost-push shock | $\sigma^\tau$ | 0.0014 |
| Standard deviation of government spending shock | $\sigma^g$ | 0.0052 |

Table 1: Key parameters of the model.

## 5.2 Numerical method: deep equilibrium nets

**Motivation**. The problem above cannot be solved using standard perturbation methods. As the model has two distinct stochastic steady states, perturbing around one of them, or around the deterministic steady state (DSS), would leave significant approximation errors when the economy is in the region around the other(s) steady states. This is distinct to several previous analyses on regime-switching models (e.g., Schorfheide, 2005; Davig and Doh, 2014; Blake and Zampolli, 2011; Bianchi and Melosi, 2017) where the dynamics are local around the deterministic steady state.

**Deep equilibrium nets**. Solving this model this requires a global solution. This involves solving a nonlinear stochastic dynamic general equilibrium model with three exogenous autoregressive shocks, one endogenous state variable, and two distinct regimes. This complexity intensifies when analyzing optimal monetary policy under commitment, which introduces two additional endogenous state variables corresponding to the Lagrange multipliers associated with forward-looking equations. The high dimensionality resulting from these state variables exceeds the capabilities of most numerical methods due to the "curse of dimensionality," where computational demands scale exponentially with each added dimension. To address this challenge, we significantly modify the deep equilibrium nets (DEQNs) approach of Azinovic et al. (2022), adapting it to effectively handle Markov-switching models. In Appendix B, we outline the basic DEQN algorithm, closely following the exposition by Azinovic et al. (2022) and Friedl et al. (2023), before detailing the necessary enhancements required to extend it to Markov-switching models.[16]

# 6 Optimal policy analysis

## 6.1 Commitment

**Ergodic distribution and dynamics**. We first analyze the ergodic distribution of the main macroeconomic variables under commitment. The distribution in Figure 1 is obtained by simulating the economy for a large number of periods. The blue bars represent the share of the ergodic distribution that happens during the normal times regime, whereas the orange bars correspond to those periods in bad times.

Several results emerge. First, the output gap (panel b) and the real interest rate (panel c) exhibit *bimodal* distributions,: the distribution of realizations clusters around two distinct points. These two points roughly correspond to the SSS of each variable, as reported in Table 2.[17] As discussed in Section 3, the bimodality in output gap and interest rates is directly related to the distortion created by the regime-switching supply shock. When in bad times, the supply shock offsets the optimal labor subsidy, leading to a distorted steady-state markup. This reduces consumption in bad times compared to the efficient allocation, explaining the abrupt fall in the output gap. Similarly, the anticipation of these changes in consumption produces an important precautionary savings motive, which accounts for the jumps in real rates.

Second, inflation clusters around zero in both regimes (panel a). This is a well-known result in monetary economics: the steady-state value of the optimal Ramsey problem under commitment

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

[17] Due to the nonlinearity of the model, the distributions are not symmetrical around the mean, and the latter does not always correspond to the SSS. We return to this point below when discussing the role of nonlinearity on our results.

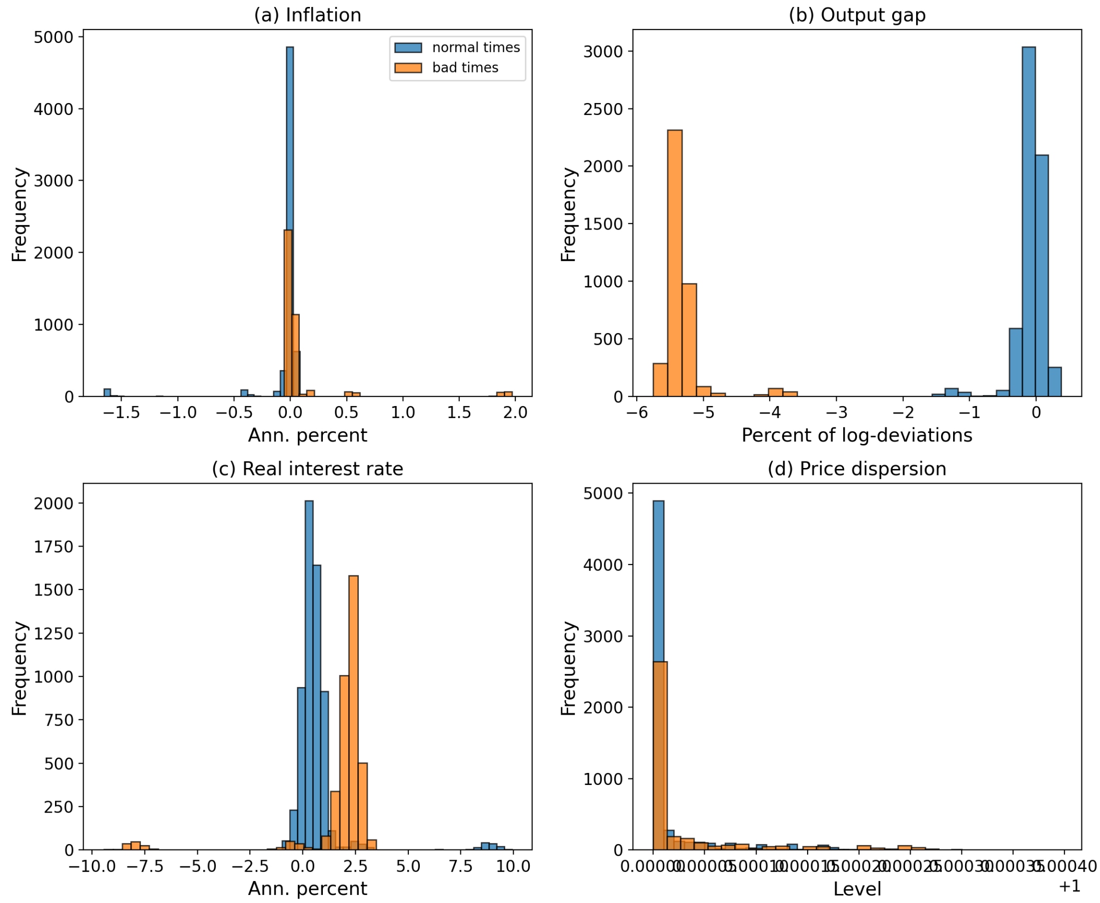


Figure 1: Ergodic distribution: commitment.
*Note*: The figure displays the ergodic distribution in the model under commitment. Colors distinguish the two regimes: blue denotes the samples corresponding to normal times, and orange denotes bad times. The figure is produced by simulating the model for a large number of periods.

is zero, irrespective of whether the steady state is distorted or efficient (Benigno and Woodford, 2005). In our case, the fact that the economy transitions from a distorted to an undistorted SSS does not modify this finding. In response to government spending and TFP shocks, the optimal policy is to keep inflation at zero at all times, as that also guarantees a zero output gap (a result commonly known as “divine coincidence”, Blanchard and Galí, 2007). In contrast, in response to cost-push shocks, the optimal policy “leans against the wind” (Galí, 2008), so that the central bank lets inflation move in the opposite direction to the output gap. These optimal responses hold in both regimes.

The comparison with the counterfactual economy with flexible prices is quite informative (see Table 2). In both cases, steady-state inflation is zero, but with nominal rigidities, it deviates

| | 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 | 0.00% | -0.04% | 0.24% | -6.12 | 0.04% | 0.02% | 0.07% | -0.79 |
| bad times | 0.00% | 0.00% | 0.00% | 0.00 | 0.00% | 0.09% | 0.35% | 4.61 | 2.84% | 2.86% | 0.07% | 0.74 |
| Output gap | | | | | | | | | | | | |
| normal times | 0.00% | 0.00% | 0.14% | 0.05 | -0.19% | -0.08% | 0.22% | -3.29 | -0.24% | -0.11% | 0.12% | -0.07 |
| bad times | -5.47% | -5.46% | 0.11% | -0.11 | -5.55% | -5.32% | 0.30% | 3.86 | -5.62% | -5.46% | 0.10% | 0.05 |
| Real rates | | | | | | | | | | | | |
| normal times | -0.01% | 0.06% | 0.39% | -0.17 | 0.35% | 0.64% | 1.29% | 5.28 | 0.12% | 0.21% | 0.40% | -0.21 |
| bad times | 2.69% | 2.74% | 0.38% | -0.35 | 2.24% | 1.80% | 1.94% | -4.42 | 2.53% | 2.54% | 0.39$ | -0.42 |
| Nominal rates | | | | | | | | | | | | |
| normal times | -0.01% | 0.06% | 0.39% | -0.17 | 0.35% | 0.60% | 1.07% | 4.92 | 0.16% | 0.23% | 0.41% | -0.23 |
| bad times | 2.69% | 2.74% | 0.38% | -0.35 | 2.24% | 1.89% | 1.59% | -4.31 | 5.38% | 5.40% | 0.37% | -0.42 |

Table 2: Key model moments.

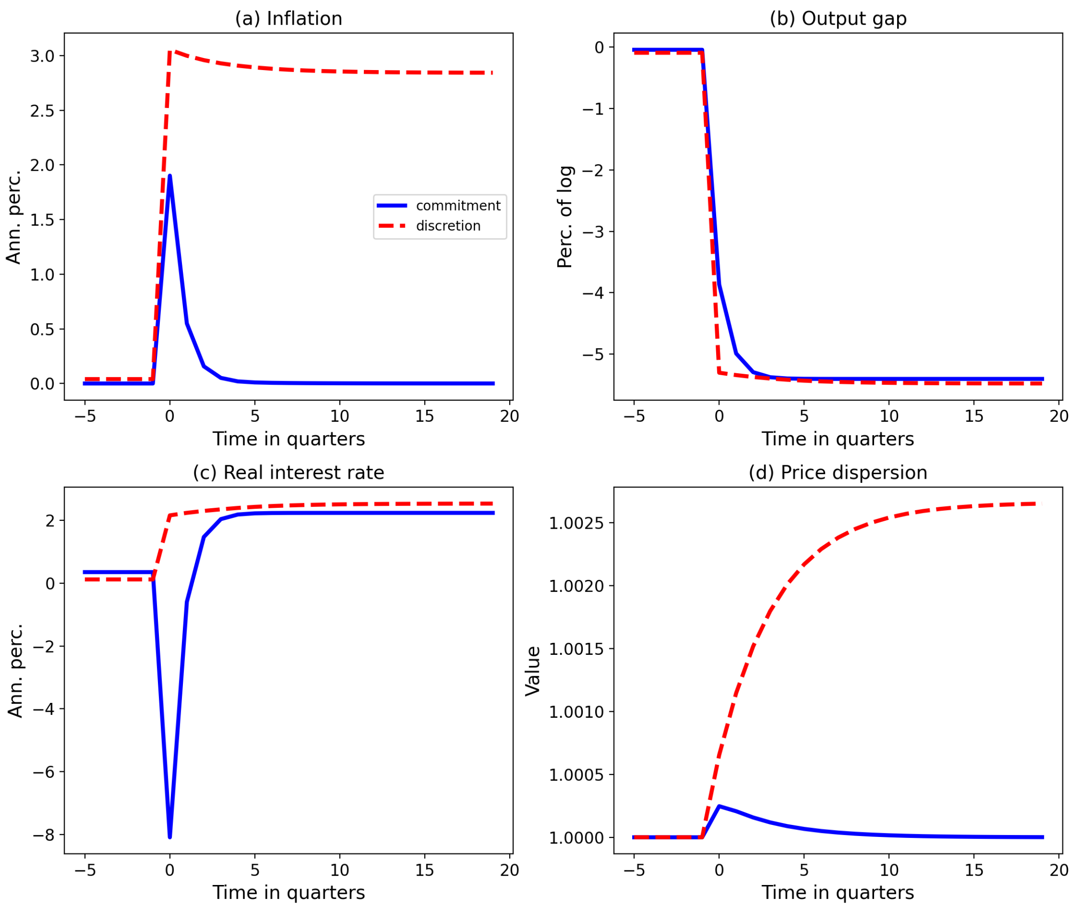


Figure 2: Response to a regime change: commitment versus discretion.
*Note*: The figure displays the transition from a normal times regime to a bad times regime in the case of the optimal policy under commitment (solid blue line) and discretion (dashed red line). We set the innovations to the temporary shocks at zero. The economy starts at the SSS of the normal times regime.

from zero in response to cost-push shocks. This explains why the distribution of inflation under commitment has positive conditional standard deviation in both normal and bad times, whereas they are zero with flexible prices.

Third, 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 (-5.55 percent versus -5.47 percent), the conditional mean is higher (-5.32 percent versus -5.47 percent). This implies that the average consumption loss in the economy under nominal rigidities is lower than under flexible prices during bad times. To understand this result, we display the transitional dynamics after a regime change in Figure 2 (solid blue line). It assumes that the economy starts at the SSS of the normal-times regime and, at time zero, a transition happens to the bad times regime, leaving the realization of all shocks at zero. In response to a transition from normal to bad times, the central bank tolerates a temporary increase in inflation (panel a), reducing the real interest rate (panel c). This leads to a smoother transition to the new SSS output gap (panel b). This is reflected by the mass of probability around -4 percent output gap panel b of Figure 1. The smoothing of the transition of the output gap when the central bank may influence it due to nominal rigidities contrasts with the case under flexible prices, when the output gap just "jumps" directly to the new SSS, as discussed in equation (14) above.

Another important difference between the baseline model under commitment and the flex-price economy is the fact that the gap between the two real rate regimes is narrower under commitment. While under flexible prices, the two real rate means are 0.06 and 2.74 percent, under commitment, they are 0.64 and 1.8 percent. The reason is that precautionary motives are lower under commitment, as the central bank provides some consumption insurance to households via its monetary policy, as discussed in the previous paragraph.

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

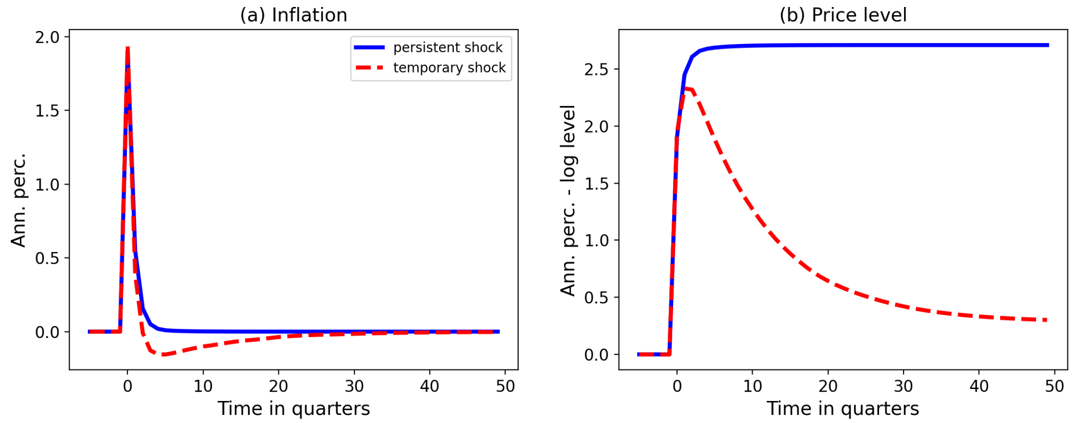


Figure 3: Comparison between a persistent (Markov switching) and temporary shocks.
*Note*: The figure displays the transition from a normal times regime to a bad times regime in the case of the optimal policy under commitment in the baseline (solid blue line) and a temporary cost-push shock calibrated to produce the same level of inflation on impact as the permanent shock (dashed red line). We set the innovations to the temporary shocks at zero. The economy starts at the SSS of the normal times regime in each case.

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 shares some parallelisms with the response to a regime change in the mean of this shock, already analyzed in the blue lines of Figure 2. In both cases, the central bank tries to cushion the effects on the output gap by letting inflation deviate from zero. The key difference between the two responses is that, in the case of a regime-switching shock, there is no deflationary period, so that the price level permanently increases after the shock. In this case, "bygones are bygones" and the central bank does not follow any price-level targeting strategy, as shown in the comparison between the two price levels as seen in panel (b) of Figure 3.

To provide some intuition for this result, Figure 4 also plots the optimal response under commitment to an autoregressive cost-push shocks featuring an autoregressive coefficient of 0.99, that is, much more persistent than in the baseline. The figure shows that, as temporary cost-push shocks become more persistent, the optimal policy results in a smaller but longer deflationary period, leading to prolonged deviations in the price level relative to its pre-shock value. This is consistent with the theoretical results derived from the log-linear approximation around the non-stochastic steady state found in textbooks (e.g., Galí, 2008). In that framework, the optimal response under commitment is given by:

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

This implies that inflation reacts in the opposite direction to the change in the output gap. As the shock becomes more persistent, the change in the output gap diminishes. In the limit, $x_t$ jumps down on impact by $-\Delta$ and then remains constant, which implies that inflation increases on impact $\frac{\Delta}{\epsilon}$ and then remains constant at zero. The logic above is local in nature; nonetheless, it provides a reasonable approximation to the global dynamics in response to a regime change, as shown in Figure 2.

**Persistent shocks or regimes?** Does the fact that a combination of two autoregressive shocks, one being extremely persistent, mimics the dynamics of the Markov-switching model? The short answer is no. The model without regimes fails to replicate the bimodality in output gap and real rates, as well as other features such as the strong skewness in price dispersion associated with shocks during bad times. This is shown in Figure 5, which compares the ergodic distribution in the baseline model (blue bars) with that of a counterfactual model that replaces the Markov transition by a highly persistent AR(1) shock with the same mean and variance. It is clearly seen how the model behaves differently. Inflation (panel a), for instance, now becomes much more volatile around zero, and the same can be said of the output gap (panel b). That completely eliminates the left skewness introduced by the regime-switching structure. More interestingly, the bimodality of real interest rates entirely disappears with very persistent supply shocks. The conclusion is that, even if very persistent supply shocks may help us rationalize the "bygones are bygones" feature of optimal policy under commitment, they fall short or reproducing the richness of new dynamics associated with regimes.

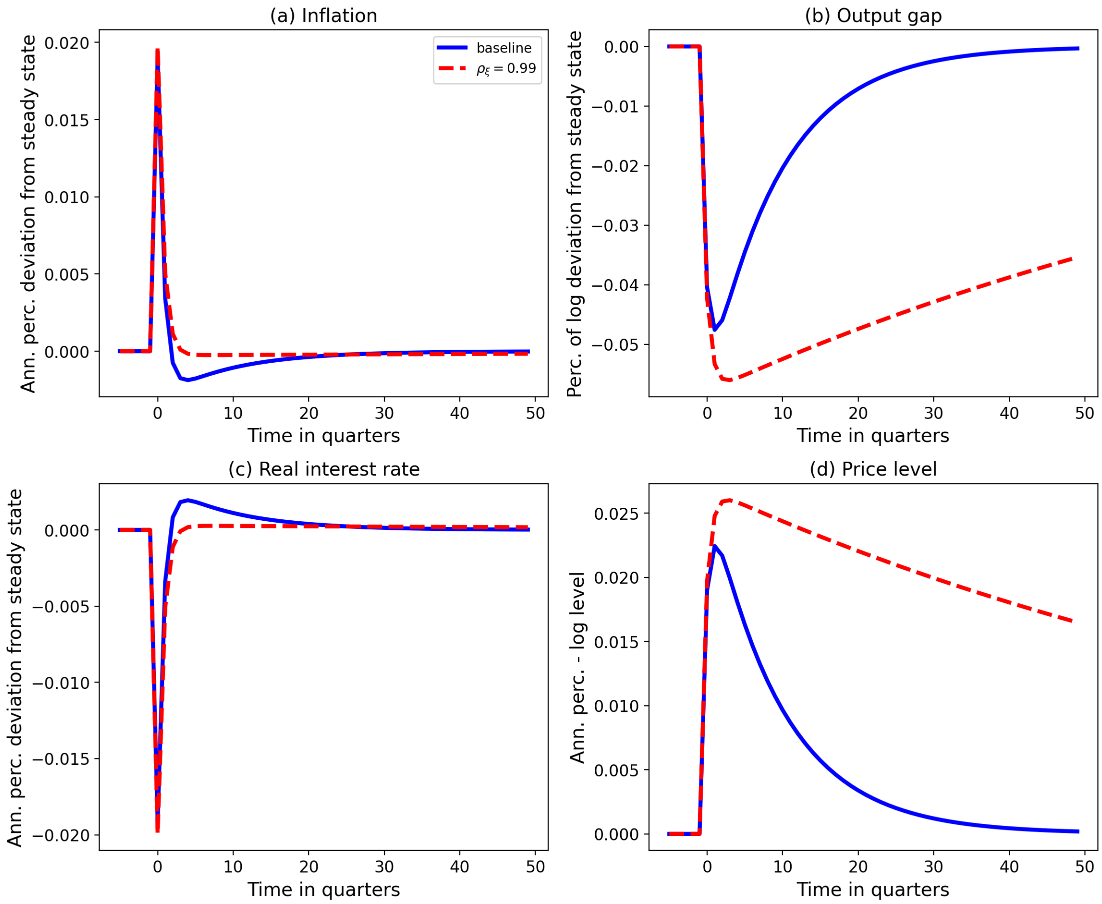


Figure 4: Impulse response to a transitory cost-push shock with different levels of persistence. *Note*: The figure displays the impulse responses to a temporary cost-push shock under the optimal policy with commitment. We consider two shocks of different autoregressive coefficients in a model without persistent cost-push shocks. The solid blue line is the baseline with $\rho^\tau = 0.90$ and the dashed red line is more persistent $\rho^\tau = 0.99$.

**On the role of non-linearities**. The standard New Keynesian model includes some nonlinearities, more relevant of all the fact that the profit function $\frac{P_t^*(j)}{p_{t+k}} y_{t+k}(P_t^*(j))$ is asymmetric for positive and negative price changes. Notwithstanding, from a quantitative perspective, nonlinearities are small (Karadi et al., 2024). The question is whether the introduction of the regime-switching shock introduces meaningful nonlinearities.

First, we start with the optimal transitions between regimes under commitment. Figure 6 compares the transitions from normal to bad times and vice-versa (with the sign reversed). It shows how the inflationary response (solid red line) discussed above are more pronounced than the deflationary responses (dashed blue line). This is related to the asymmetries in the profit function discussed above, which play a more prominent role in the case of large jumps in markups.

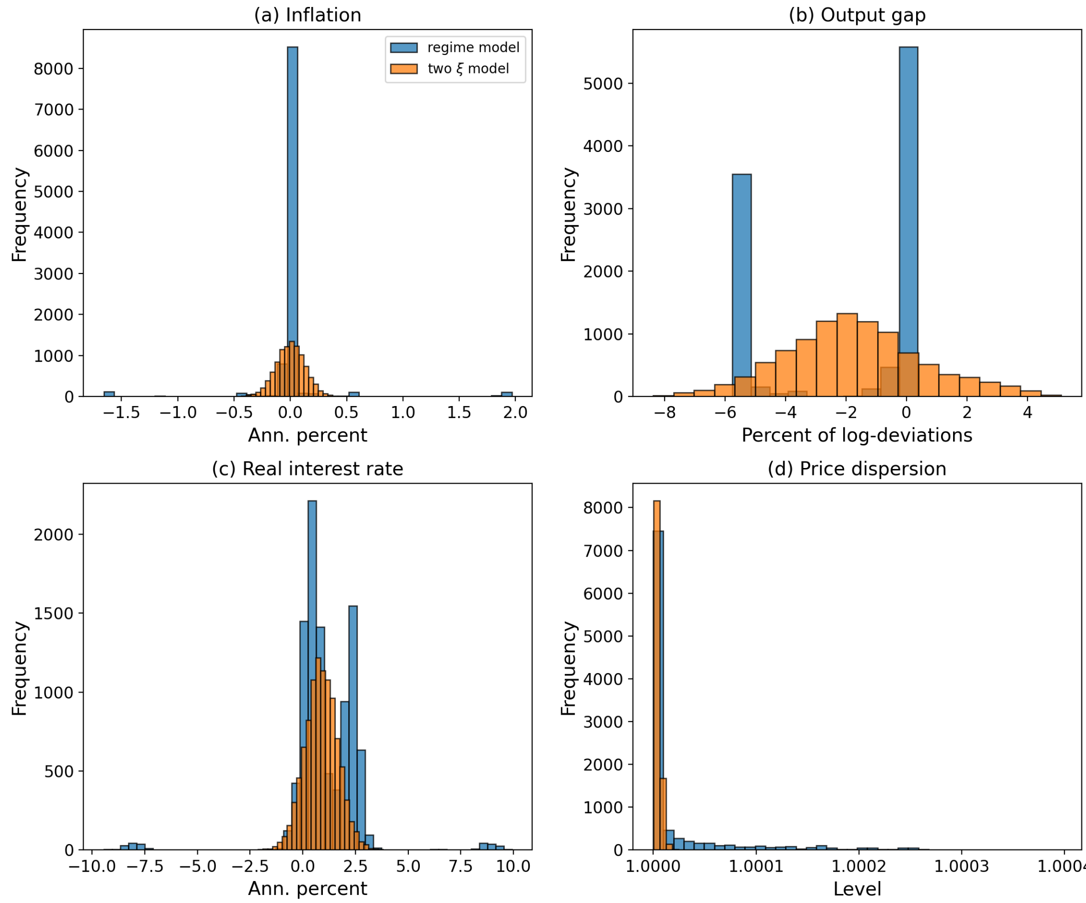


Figure 5: Histograms of two temporary ($\xi$) autoregressive shocks (orange), compared to the baseline model under commitment (blue).

*Note*: The figure displays the ergodic distribution in the model two AR(1) shocks and under commitment. Colors distinguish the two models: blue denotes the samples corresponding to commitment, and orange to the standard New Keynesian model calibrated with the optimal labor subsidy and two cost-push shocks of different autoregressive coefficients: one with $\rho^{\tau} = 0.90$ and the other $\rho^{\tau} = 0.99$.

Benigno and Rossi (2021) discuss how higher-order optimal policy exhibits an "expansionary bias", as the central bank responds asymmetrically to positive and negative cost-push shocks. The same happens here, but at a much larger scale due to the regimes.

The consequence of this asymmetry is a significant positive drift in the price level. As the central bank inflates more than it deflates in the transitional dynamics, a systematic increase in the price level happens. Figure 7 displays the price level in the standard New Keynesian model without regimes under commitment in panel (a), and in the regime-switching model in panel (b). The asymmetries embedded in the standard New Keynesian model give rise to a small positive

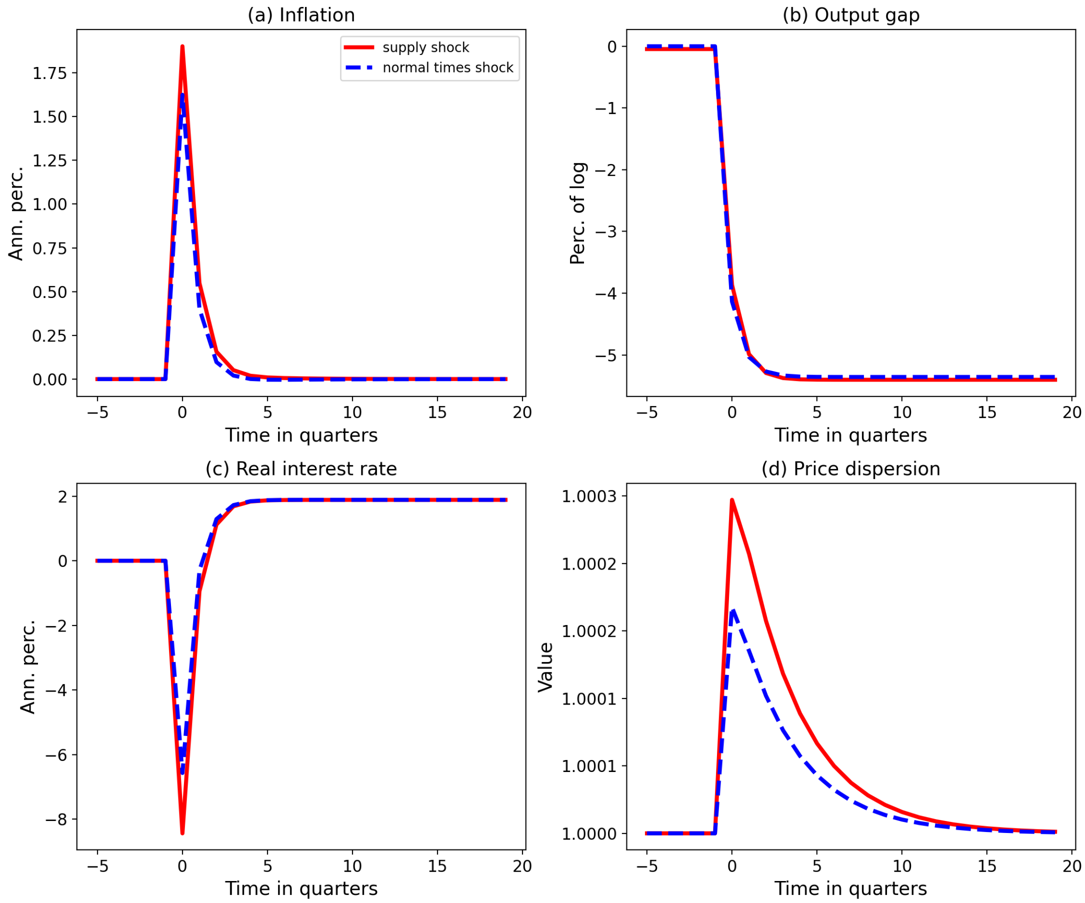


Figure 6: Transitions from and to bad timed under commitment.
*Note*: The figure displays the transition from a normal times regime to a bad times regime (solid red) and vice versa (dashed blue) in the case of the optimal policy under commitment. The responses from bad to normal times are multiplied by minus one. We set the innovations to the temporary shocks at zero. The economy starts at the SSS of teach regime.

trend (panel a): prices increase 15 times in 10,000 quarters. This result can only be obtained through a global solution, or a higher-order perturbation method, as linear approximations yield price stability. In the case of the regime-switching model, nonlinearities play a more relevant role, leading to price levels almost 100 times larger in the same period.

Second, we move to changes in shock size. Figure 12 in Appendix C analyzes the optimal response under commitment to a larger regime change: instead of a mean of $\bar{\eta}$, we consider a mean of $2\bar{\eta}$. The new mean implies a greater change in both the output gap and the real interest rate between the two regimes. This occurs because the change in SSS consumption between the two regimes is now larger, and the level of distortions in the bad times regime is also higher than in

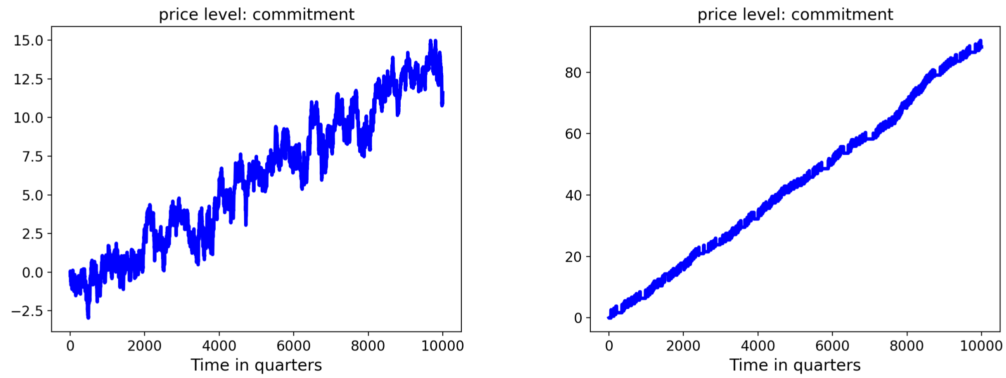


Figure 7: Price level dynamics.
*Note*: The figure displays the simulation of 10,000 quarters in the model under commitment. Panel a considers an additional AR(1) cost push shocks calibrated to replicate the mean and variance of the regime-switching shock in the baseline. Panel b is the baseline.

the baseline. The temporary increase in inflation is also more pronounced in this case, leading to a larger increase in price dispersion. Notwithstanding, proportional to the shock size, the effect is small. This is because the effect on price dispersion is proportionally larger, and the central bank weighs the disutility of misallocation due to higher price (and markup) dispersion against that of smoothing out the path of the average markup. In other words, the sacrifice ratio of the model is slightly nonlinear with the size of the regime shock.

## 6.2 Discretion

**Ergodic distribution and dynamics**. We move next to the case with discretion. Figure 8 displays the ergodic distributions. First, inflation clusters now around two points: zero inflation and 2.8 percent inflation (panel a). In normal times, the SSS of the economy is undistorted and the central bank has no incentive to create any inflation (see Table 2 above). This contrasts with bad times, when inflation centers around 2.8 percent. This difference is due to the “inflationary bias” originally highlighted by Kydland and Prescott (1977) and Barro and Gordon (1983). This bias emerges due to the distortion created by the regime-switching supply shock, as it offsets the optimal labor subsidy, leading to a distorted steady-state markup. When in bad times, the central bank has an incentive to surprise private agents by loosening monetary policy and creating inflation, in an attempt to reduce the markup. However, the private sector anticipates this incentive and

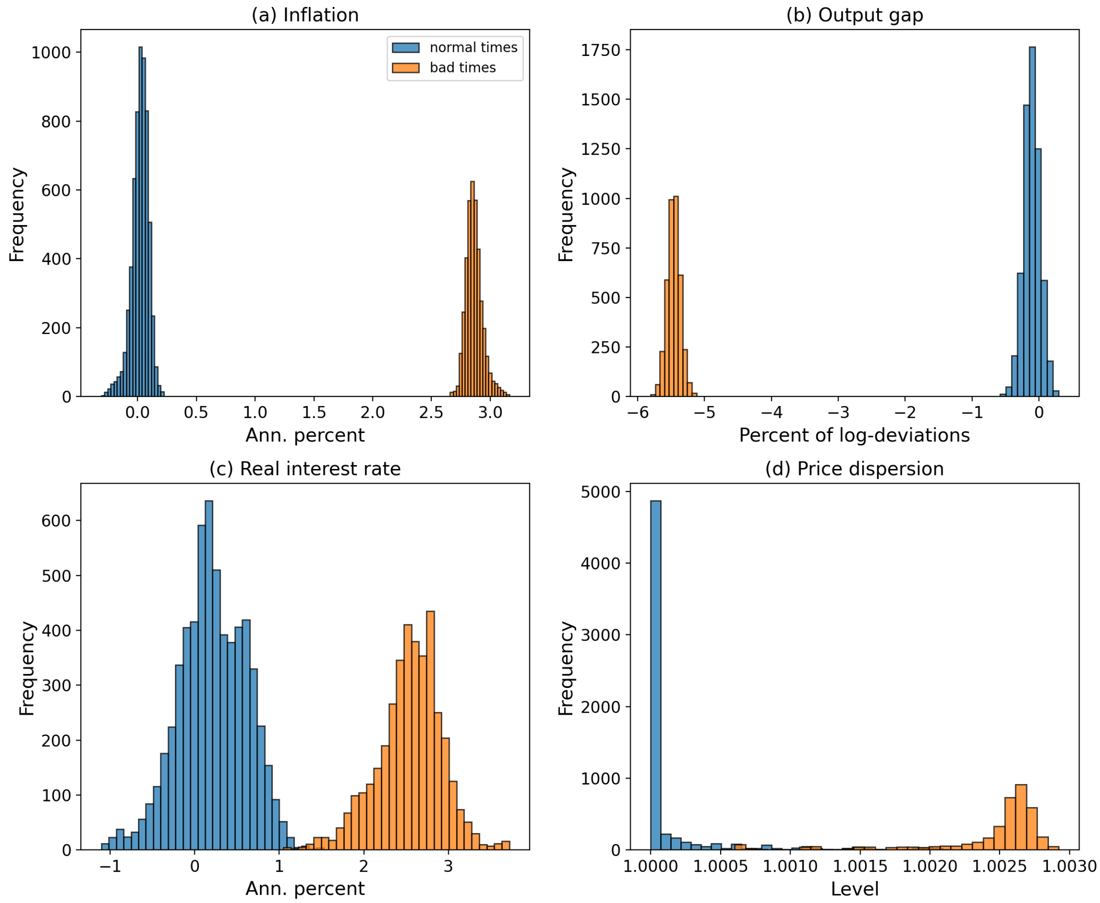


Figure 8: Ergodic distribution: discretion.

*Note*: The figure displays the ergodic distribution in the model under discretion. Colors distinguish the two regimes: blue denotes the samples corresponding to normal times, and orange to bad times. The figure is produced by simulating the model for a large number of periods.

incorporates it into their expectations, leading to positive inflation. The outcome is that the economy ends up with a systematically higher level of inflation due to the central bank's lack of commitment.

Compared to the case under commitment, a discretionary central bank provides less macro insurance to households. This can be seen in Figure 2 above, where the dashed red line displays the transition from normal to bad times in the model under discretion. When the shock arrives, inflation jumps, overshooting 2.8 percent and then declining to the latter value. The real interest rate does not decline, in contrast to the case under commitment. This reflects the fact that consumption and the output gap fall more abruptly than under commitment. By looking at Table 2, we see that the mean output gap in bad times is similar to that in the flex-price economy.

Households are therefore worse off, as they suffer higher markup (and labor) misallocation due to price dispersion while not enjoying significantly better macro insurance.

Afrouzi et al. (2023) analyze the transitional dynamics under discretion in a similar model. They consider a shock that shifts the economy from a deterministic steady state with low markups to another with high markups, equivalent to a transition from normal to bad times. This is a one-off, unanticipated ("MIT") shock under perfect foresight, instead of the fully stochastic environment that we consider here. They similarly find that inflation overshoots before converging to the high-inflation steady state. The key difference between the two papers is the evolution of the real rate, linked to the different assumptions regarding the nature of risk. In Afrouzi et al. (2023), the real interest rates jump down on impact, as the central bank initially stimulates the economy to weather the shock, and then gradually returns to its initial value, which corresponds to the deterministic steady state. In our economy, instead, the final value is above the initial one, as the precautionary motives lead real rates to jump (see the discussion in Section 3). The real rates do not jump down, *but up*, converging to their new value.

**A zero-mean regime**. We consider an alternative economy where the mean of the cost-push shock in normal times is such that the mean of the regime-change is zero. In this case, the SSSs are distorted in both regimes, as the average markup is now too high or too low compared to the efficient one, depending on the regime.

The inflationary bias of the economy disappears: while inflation is still positive (2.6 percent) in bad times, it is negative (-2.77 percent) in normal times. This reflects the fact that the central bank has the opposite incentives when labor is excessively subsidized. Similarly, real rates are 1.03 percent in normal times and 4.28 percent in bad times, reflecting the widening gap in consumption inequality across regimes, and the need to self-insure against it.

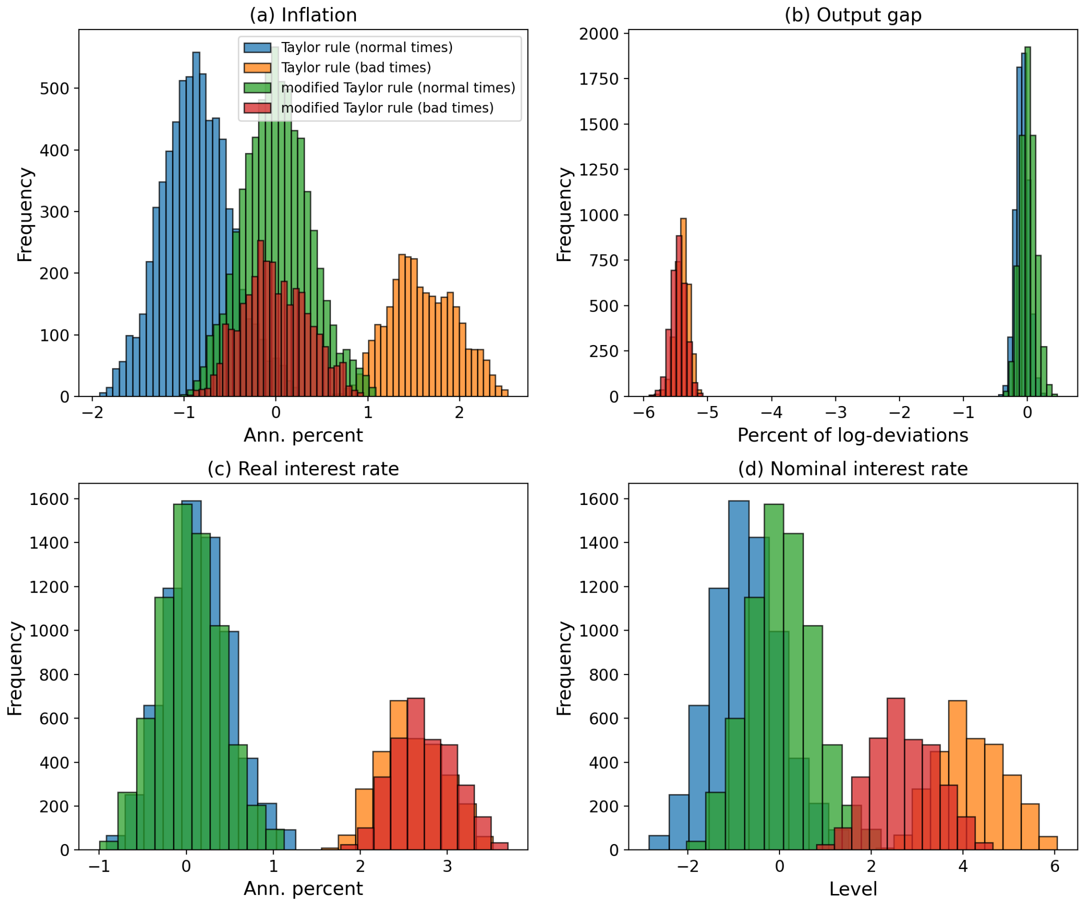


Figure 9: Ergodic distribution
*Note*: The figure displays the ergodic distribution in the model under a standard Taylor rule and a modified one. Colors distinguish the two regimes: blue denotes the samples corresponding to the standard Taylor rule in normal times, and orange in bad times. Green is the modified Taylor rule in normal times, and red in bad times. The figure is produced by simulating the model for a large number of periods.

# 7 Monetary policy rules

**Taylor rule**. Consider now a Taylor rule of the following form:

$$i_t = \frac{(1+\bar{\pi})}{\beta} - 1 + \psi\left(\pi_t - \bar{\pi}\right), \tag{24}$$

where $\bar{\pi} = 0$ is the inflation target and $\psi = 2$ is the slope of the Taylor rule to deviations in inflation.

Figure 9 shows the ergodic distribution in this case. While results share the bimodality of output gap and the real rate with optimal policy above, long-term inflation in this model is not zero (panel a). In the normal times regime (blue bars), it is negative (centered around -0.9 percent) and in the bad times (orange bars) regime, it is positive (centered around 1.6 percent). Such deviations from the central bank target $\bar{\pi} = 0$ are the result of the Taylor rule not targeting the adequate natural rate. If we evaluate the Taylor rule in a SSS $n = 1, 2$, we obtain

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

where $\bar{r}$ is the real rate target of the central bank, which, under the Taylor rule (24), coincides with the real rate in the non-stochastic steady state of the efficient allocation, $\bar{r} = \hat{r}$. Replacing the nominal rate using the Fisher equation $i_{n,ss} = r_{n,ss} + \pi_{n,sss}$, we get

$$\pi_{n,sss} \simeq \bar{\pi} + \frac{r_{n,ss} - \bar{r}}{\psi - 1}. \tag{25}$$

Equation (25) illustrates how long-run inflation deviates from the central bank's target if the monetary policy rule targets an incorrect long-term real rate. This equation was first proposed by Campos et al. (2024) in the context of HANK models, where the natural rate is endogenous to fiscal policy. In contrast, in our model, it is the regime-switching nature of the cost-push shock that drives changes in the natural rate. In our model, there is a significant gap between the natural rate in each regime and the central bank's target rate: in normal times, the central bank targets a natural rate that is too high, which tightens monetary policy excessively and explains why inflation is consistently below target: $\pi_{1,sss} \simeq 0.1\% - 1\% = -0.9\%$. Conversely, in bad times, the central bank sets nominal rates too low, explaining why inflation is above target: $\pi_{2,sss} \simeq 2.6\% - 1\% = 1.6\%$.

Despite the central bank's failure to stabilize inflation in this economy, it meets its price stability mandate *on average*. Average inflation in the ergodic distribution is -0.1 percent, and the average real interest rate is 0.9 percent, satisfying equation (25) on average: $0.9\% - 1\% = -0.1\%$.

The gap between the two real rate regions is also lower in this economy than under flexible prices, the two SSS rates are 0.11 percent in normal times and 2.59 percent in bad times, reflecting again the insurance provided by the central bank.

These results provide a rationale for the joint dynamics of inflation and long-term interest rates before and after the COVID pandemic (Benigno et al., 2024). Before COVID, long-term inflation expectations in advanced economies, such as the Euro area, were below target, while natural rate estimates were close to zero or even negative. Since the pandemic, both inflation expectations and natural rate estimates have increased abruptly in line with negative supply shocks, such as the war in Ukraine or an incipient de-globalization process.

**Modified Taylor rule**. We consider an alternative monetary policy rule that is regime-contingent. The new Taylor rule is

$$i_t = \bar{r}_t + \psi \left( \pi_t - \bar{\pi} \right), \tag{26}$$

where the Taylor rule intercept equals the natural rate in each regime

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

Here $r_n^*$ is the natural rate in regime $n$, that is, the SSS real rate in the counterfactual flex-price allocation. In this case, it is easy to see that equation (25) is compatible with a zero inflation target $\bar{\pi} = 0$. The red and green bars in Figure 9 display the ergodic distribution under this modified Taylor rule. The inflation distribution is centered on zero inflation. The variances of the different variables are similar under the original and the modified Taylor rules.

The policy prescription is clear: in the presence of supply regimes the central bank should endogenously adapt its interest rate target to track the natural rate, which becomes a regime-contingent object.

**Sensitivity to regime length**. Finally, Figure 10 displays the sensitivity of SSS inflation and real rates to alternative average lengths of the persistent supply shock period under the Taylor rule (24). The blue solid line represents the normal-times SSS baseline model. As discussed above, the parameter $p_{21}$ controls the probability of transitioning from regime 2 (bad times) to 1 (normal times). This is the inverse of the duration of the bad times. The baseline calibration implies an average duration of 24

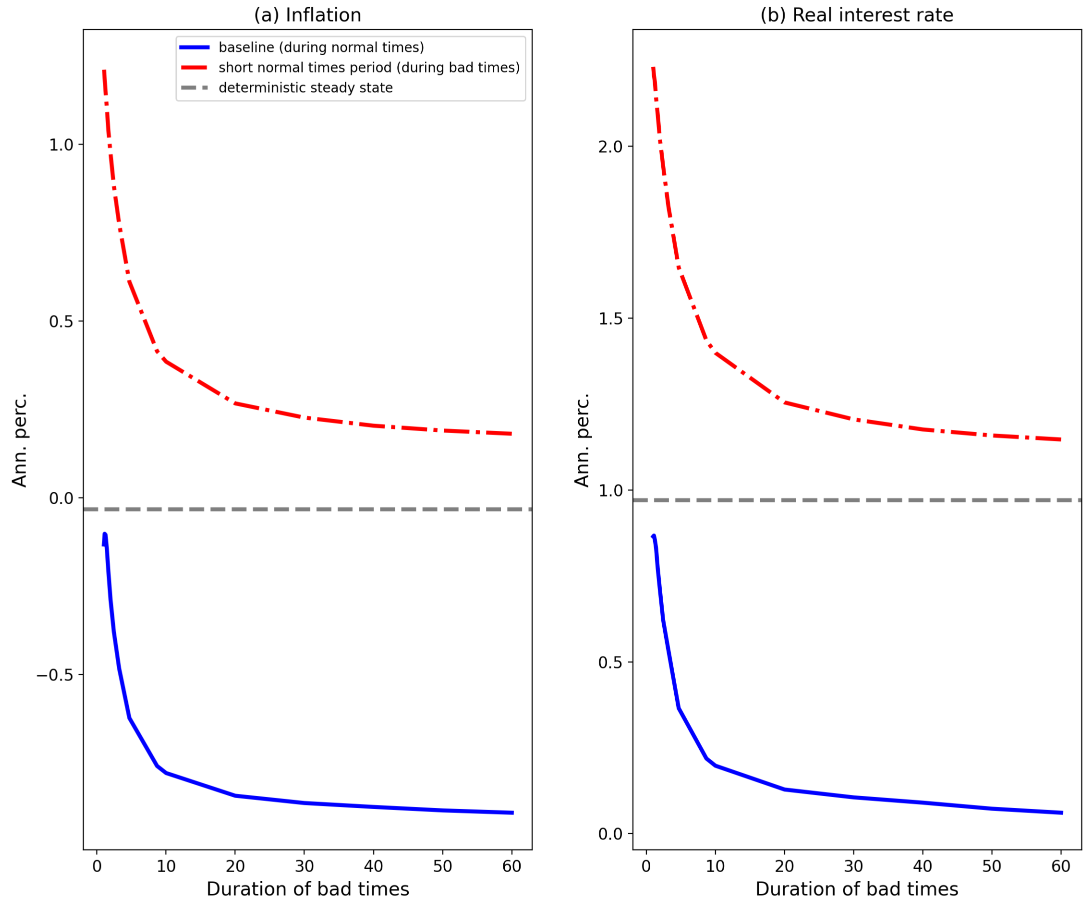


Figure 10: Sensitivity to regime length (Taylor rule).
*Note*: The figure displays the SSS values of inflation and the real interest rate for alternative values of the average duration for the bad times regime. The blue line displays the normal-times SSS values in the baseline model, and the dashed red line shows the bad-times SSS in a counterfactual case in which $p_{12} \to 1$. The dotted gray line indicates the values in the efficient allocation.

quarters.

The SSS values of inflation and the real interest rate are very nonlinear with respect to changes in the average ratio of the bad times regime. If the duration is reduced, making it less persistent, the SSS value of inflation and real rates during normal times converges to the one in the efficient allocation (zero inflation and 1 percent rates). This is precisely because, as $p_{21} \to 1$, agents understand that the regime change is equivalent to a one-period iid shock, and thus the SSS converges to the one in a model without regimes. If instead the duration increases, the values of inflation and real rates eventually converge to values close to those in the baseline calibration.

The dashed red line in Figure 10 shows the bad-times SSS in a counterfactual case in which

$p_{12} \to 1$, that is, in which the average duration of the normal times is one period. In this case, the convergence to the SSS of the efficient allocation happens if the duration of the persistent shock regime *increases*. The logic is reversed compared to the previous case. In this case, the economy spends most of the time in the bad-times regime, but it experiences very short transitions to the normal-times regime, equivalent to positive iid cost-push shocks, with temporarily reduced markups.

# 8 Conclusions

This paper advances the understanding of optimal monetary policy in the presence of regime-switching supply shocks. We introduce a New Keynesian model that captures the economy's transition between "normal times" and "bad times", characterized by sustained increases in production costs, such as with tariffs.

This extension challenges traditional monetary policy prescriptions derived under the assumption of temporary shocks, particularly the price-level targeting strategies that may not be applicable in a regime-switching environment. First, under commitment, the central bank cushions the persistent increase in the mean of supply shocks by allowing for a temporary increase in inflation, followed by a gradual return to the target without necessitating a deflationary period. This contrasts with the standard prescription under temporary shocks, where future deflation is used to offset current inflation, reaffirming that "bygones are bygones" in the context of supply regimes. Second, an inflationary bias emerges during bad times when analyzing optimal policy under discretion because the central bank cannot commit to future policies. Third, we demonstrate that traditional Taylor rules fail to stabilize inflation across different regimes due to endogenous shifts in the natural interest rate.

Methodologically, our paper is, to the best of our knowledge, the first to solve for the optimal policy in a New Keynesian model globally with a regime-switching cost-push shock, utilizing advanced computational techniques based on deep learning. This allows us to capture the global dynamics of the model accurately and efficiently.

## References


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

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

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

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

Azinovic, M. and J. Žemlička (2024): "Intergenerational Consequences of Rare Disasters," .

Bandera, N., L. Barnes, M. Chavaz, Silvana, Tenreyro, and L. von dem Berge (2023): "Monetary policy in the face of supply shocks: the role of inflation expectations," Tech. rep., European Central Bank.

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

Bayer, C., B. Born, and R. Luetticke (2023): "The Liquidity Channel of Fiscal Policy," *Journal of Monetary Economics*, 134, 86–117.

Bellman, R. (1957): *Dynamic Programming*, Princeton, NJ: Princeton University Press, introduction of the curse of dimensionality in optimization and decision processes.

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*.

Benigno, P. and L. Rossi (2021): "Asymmetries in monetary policy," *European Economic Review*, 140.

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

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

Bergstra, J. S., R. Bardenet, Y. Bengio, and B. Kégl (2011): "Algorithms for Hyper-Parameter Optimization," in *Advances in neural information processing systems*, 2546–2554.

Bianchi, F. and L. Melosi (2017): "Escaping the Great Recession," *American Economic Review*, 107, 1030–58.

Bianchi, F., L. Melosi, and M. Rottner (2021): "Hitting the elusive inflation target," *Journal of Monetary Economics*, 124, 107–122.

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

Blake, A. P. and F. Zampolli (2011): "Optimal policy in Markov-switching rational expectations models," *Journal of Economic Dynamics and Control*, 35, 1626–1651.

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

Calvo, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383 – 398.

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.

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

Choi, J. and A. Foerster (2021): "Optimal Monetary Policy Regime Switches," *Review of Economic Dynamics*, 42, 333–346.

Clarida, R., J. Gali, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Coibion, O., Y. Gorodnichenko, and J. Wieland (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. and T. Doh (2014): "Monetary Policy Regime Shifts and Inflation Persistence," *The Review of Economics and Statistics*, 96, 862–875.

Debortoli, D. and R. Nunes (2014): "Monetary Regime Switches and Central Bank Preferences," *Journal of Money, Credit and Banking*, 46, 1591–1626.

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

Federle, J., A. Meier, G. Müller, and M. Schularick (2024): "THE PRICE OF WAR," CEPR Discussion Papers 18834, C.E.P.R. Discussion Papers.

Fernandez-Villaverde, J., T. Mineyama, and D. Song (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., G. Nuno, and J. Perla (2024b): "Taming the Curse of Dimensionality: Quantitative Economics with Machine Learning," Papers.

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

Friedl, A., F. Kübler, S. Scheidegger, and T. Usui (2023): "Deep uncertainty quantification: With an application to integrated assessment models," Working paper.

Gagnon, E., B. K. Johannsen, and D. López-Salido (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.

Goodfellow, I., Y. Bengio, and A. Courville (2016): *Deep Learning*, MIT Press, `http://www.deeplearningbook.org`.

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

Han, J., Y. Yang, and W. E (2021): "DeepHAM: A global solution method for heterogeneous agent models with aggregate shocks," *arXiv preprint arXiv:2112.14377*.

Hornik, K., M. Stinchcombe, and H. White (1989): "Multilayer feedforward networks are universal approximators," *Neural Networks*, 2, 359–366.

Kaplan, G., G. Nikolakoudis, and G. L. Violante (2023): "Price Level and Inflation Dynamics in Heterogeneous Agent Economies," Tech. rep., Princeton.

Karadi, P., A. Nakov, G. Nuno, E. Pasten, and D. Thaler (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.

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

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

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

Maechler, A. (2024): "Monetary policy in an era of supply headwinds – do the old principles still stand?" Speech by Ms Andréa M Maechler, Deputy General Manager, Bank for International Settlements, at the London School of Economics, London.

Maliar, L., S. Maliar, and P. Winant (2021): "Deep learning for solving dynamic economic models." *Journal of Monetary Economics*, 122, 76–101.

Mankiw, N. G. (2024): "Six beliefs I have about inflation: Remarks prepared for NBER conference on "Inflation in the Covid era and beyond"," *Journal of Monetary Economics*, 103631.

Mian, A., L. Straub, and A. Sufi (2021): "Indebted Demand," Youtube video hereQuarterly Journal of Economics, 136 (4) 2021: 2243-2307.

Monacelli, T. (2025): "Tariffs and Monetary Policy," CEPR Working Papers 20142, CEPR.

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

Powell, J. H. (2023): "Monetary Policy Challenges in a Global Economy," Remarks by Chairman Jerome Powell at the 24th Jacques Polak Annual Research Conference, hosted by the International Monetary Fund, Washington, D.C.

Rachel, L. and L. H. Summers (2019): "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity*, 50, 1–76.

Sahuc, J.-G., F. Smets, and G. Vermandel (2023): "The New Keynesian Climate Model," .

Schnabel, I. (2024): Interview with Frankfurter Allgemeine Zeitung.

Schorfheide, F. (2005): "Learning and Monetary Policy Shifts," *Review of Economic Dynamics*, 8, 392–419.

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

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Yun, T. (2005): "Optimal Monetary Policy with Relative Price Distortions," *American Economic Review*, 95, 89–109.

Zandweghe, W. V. and A. L. Wolman (2019): "Discretionary monetary policy in the Calvo model," *Quantitative Economics*, 10, 387–418.

# Online appendix

## A Optimal policy derivations

This Appendix first reproduces the private equilibrium conditions and then derives the first-order conditions of the optimal monetary policy problem under discretion and commitment that we introduced in Section 4.

### A.1 Equilibrium conditions

The equilibrium conditions are given by:

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

The shocks are:

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

where $\varepsilon_t^A \sim N\left(0, \sigma^A\right), \varepsilon_t^g \sim N\left(0, \sigma^g\right)$, $\varepsilon_t^\tau \sim N\left(0, \sigma^\tau\right)$, and $\varepsilon_t^\eta \sim [[1-p_{12}, p_{12}], [p_{21}, 1-p_{21}]]$.

## A.2 Optimal policy under discretion

Next, we derive the first-order conditions of the problem under discretion. In the case in which the central bank cannot commit to future policies, we express the problem as follows:

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

subject to

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

where we have defined

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

The Lagrangian is given by:

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

The first-order condition with respect to consumption reads:

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

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

$$0 = \rho_t \theta \left(\epsilon - 1\right) \left(1 + \pi_t\right)^{\epsilon - 2} + \zeta_t \theta \epsilon \left(1 + \pi_t\right)^{\epsilon - 1} \Delta_{t-1}. \tag{27}$$

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

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

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

$$\begin{aligned} 0 =& - \frac{\left(\frac{(c_t + g_t)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} + \beta \mathbb{E}_t \left[\frac{\partial V}{\partial \Delta_t}\right] - \zeta_t + \\ & \mu_t \left[ p_t^* c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial F}{\partial \Delta_t}\right] - \mathcal{M} \left( \left(c_t + g_t\right)^{1+\omega} c_t^{\gamma} \frac{\omega}{\Delta_t} \left(\frac{\Delta_t}{A_t}\right)^{\omega} \left(1 + \tau_t\right) \left(A_t\right)^{-1} + c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_t}\right] \right) \right]. \end{aligned} \tag{28}$$

According to the envelope theorem, the following holds at the optimum:

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

Therefore we can substitute it in equation (28):

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

The full set of equations is then given by:

$$
\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}\right)\right]\right) - \right. \\
& \left. \mathcal{M} \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} + \gamma c_t^{\gamma-1} \mathbb{E}_t \left[G\left(\Delta_t, A_{t+1}, \tau_{t+1}\right)\right]\right)\right], \\
0 =& \rho_t \theta (\epsilon - 1) (1+\pi_t)^{\epsilon-2} + \zeta_t \theta \epsilon (1+\pi_t)^{\epsilon-1} \Delta_{t-1}, \\
0 =& \mu_t \left[(c_t + g_t) + c_t^{\gamma} \mathbb{E}_t \left[F\left(\Delta_t, A_{t+1}, \tau_{t+1}\right)\right]\right] + \rho_t (1-\theta)(1-\epsilon)(p_t^*)^{-\epsilon} - \zeta_t (1-\theta) \epsilon (p_t^*)^{-\epsilon-1}, \\
0 =& \mu_t \Xi_t^D + \rho_t (1-\theta)(1-\epsilon)(p_t^*)^{-\epsilon} - \zeta_t (1-\theta) \epsilon (p_t^*)^{-\epsilon-1}, \\
0 =& -\frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} + \beta \mathbb{E}_t \left[\theta (1+\pi_{t+1})^{\varepsilon} \zeta_{t+1}\right] - \zeta_t + \\
& \mu_t \left[p_t^* c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial F}{\partial \Delta_{t-1}}\right] - \mathcal{M} \left((c_t + g_t)^{1+\omega} c_t^{\gamma} \frac{\omega}{\Delta_t} \left(\frac{\Delta_t}{A_t}\right)^{\omega} (1+\tau_t) (A_t)^{-1} + c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_{t-1}}\right]\right)\right], \\
c_t^{-\gamma} =& \lambda_t, \\
h_t^{\omega} =& w_t \lambda_t, \\
\Xi_t^N =& y_t w_t (1+\tau_t) (A_t)^{-1} + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} (1+\pi_{t+1})^{\epsilon} \Xi_{t+1}^N\right], \\
\Xi_t^D =& y_t + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} (1+\pi_{t+1})^{\epsilon-1} \Xi_{t+1}^D\right], \\
p_t^* =& \mathcal{M} \frac{\Xi_t^N}{\Xi_t^D}, \\
1 =& \theta (1+\pi_t)^{\epsilon-1} + (1-\theta)(p_t^*)^{1-\epsilon}, \\
\Delta_t =& \theta (1+\pi_t)^{\epsilon} \Delta_{t-1} + (1-\theta)(p_t^*)^{-\epsilon}.
\end{aligned}
$$

## A.3 Optimal policy under commitment

Finally, we analyze the optimal policy under commitment. The Lagrangian is given by:

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

The first-order condition with respect to consumption reads as:

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

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

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

The first-order condition for inflation reads as:

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

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 $\Xi_t^N$ 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 $\Xi_t^D$ reads as:

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

The full set of equations is then given by:

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

# B DEQN algorithm

The DEQN algorithm is a simulation-based solution method that utilizes deep neural networks to compute an approximation of the *optimal policy function* $\mathbf{p} : X \to Y \subset \mathbb{R}^M$ for a dynamic model, assuming that the underlying economy is characterized by discrete-time first-order equilibrium conditions, expressed as:

$$\mathbf{G}(\mathbf{x}, \mathbf{p}) = \mathbf{0}, \ \forall \mathbf{x} \in X \subset \mathbb{R}^d. \tag{29}$$

Intuitively, DEQNs operate as follows: An unknown policy function is approximated using a neural network, denoted as $\mathbf{p}(\mathbf{x}) \approx \mathcal{N}(\mathbf{x})$, with trainable parameters $\nu$. These parameters are initially unknown and must be determined based on an appropriate loss function that measures the accuracy of the approximation at a given state of the economy.

Although there are several types of deep neural networks, in this paper, we utilize densely-connected feedforward neural networks (FNNs).[18] Following the literature, we define an $L$-layer FNN as a function $\mathcal{N}^L(\mathbf{x}) : \mathbb{R}^{d_{\text{input}}} \to \mathbb{R}^{d_{\text{output}}}$ and specify that there are $L-1$ hidden layers, with the $\ell$-th layer containing $N_\ell$ neurons. In our case, $N_0 = d_{\text{input}}$ and $N_L = d_{\text{output}}$.[19] Furthermore, for each $1 \leq \ell \leq L$, we define a weight matrix $\mathbf{W}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and a bias vector $\mathbf{b}^\ell \in \mathbb{R}^{N_\ell}$. Letting $A^\ell(\mathbf{x}) = \mathbf{W}^\ell \mathbf{x} + \mathbf{b}^\ell$ be the affine transformation in the $\ell$-th layer, for some non-linear activation function $\sigma(\cdot)$ (such as ReLU, Swish, or SELU), an FNN is expressed as:

$$\mathbf{p}\left(\mathbf{x}\right) \approx \mathcal{N}\left(\mathbf{x}\right) = \mathcal{N}^L\left(\mathbf{x}\right) = A^L \circ \sigma_{L-1} \circ A^{L-1} \circ \cdots \circ \sigma_1 \circ A^1\left(\mathbf{x}\right). \tag{30}$$

Figure 11 in Appendix C illustrates a simple neural network with two hidden layers.

The selection of hyperparameters $\left\{L, \{N_\ell\}_{\ell=1}^L, \{\sigma_\ell(\cdot)\}_{\ell=1}^L\right\}$ is known as the architecture selection. Approaches to determine these hyperparameters include using prior experience, manual tuning, random or grid search, as well as more sophisticated methods such as Bayesian optimization (see, e.g., Bergstra et al., 2011).

The DEQN algorithm to determine $\mathbf{p}\left(\mathbf{x}\right)$ begins by randomly initializing the parameters $\nu$, representing an initial guess for the unknown approximate policy function. Next, a sequence of $N_{\text{path length}}$ states are simulated. Starting from a given state $\mathbf{x}_t$, the next state $\mathbf{x}_{t+1}$ is determined by the policies encoded in the neural network, $\mathcal{N}\left(\mathbf{x}\right)$, combined with the remaining model-implied dynamics.

If the (approximate) policy function satisfying the equilibrium conditions were known, equation (29) would hold along the simulated path. However, since the neural network is initialized with random coefficients, $\mathbf{G}\left(\mathbf{x}_t, \mathcal{N}\left(\mathbf{x}_t\right)\right) \neq 0$ along the simulated path of length $N_{\text{path length}}$. This discrepancy is leveraged to improve the guessed policy function. Specifically, DEQNs use a loss function based on the error in the equilibrium conditions:

$$\ell_\nu = \frac{1}{N_{\text{path length}}} \sum_{\mathbf{x}_t \text{ on sim. path}} \sum_{m=1}^{N_{\text{eq}}} \left(\mathbf{G}_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t))\right)^2, \tag{31}$$

where $\mathbf{G}_m\left(\mathbf{x}_t, \mathcal{N}\left(\mathbf{x}_t\right)\right)$ represents each of the $N_{\text{eq}}$ first-order equilibrium conditions of the model,

[18] Neural networks are universal function approximators (Hornik et al., 1989) capable of resolving highly non-linear features and handling high-dimensional input data. See, for example, Goodfellow et al. (2016) for a general introduction to deep learning.

[19] Our models have at least $d_{\text{input}} = d = 5$ input dimensions and $d_{\text{output}} = 3$ output dimensions (see, for example, the discretion model in Section 4).

i.e., $\mathbf{G}(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)) = \sum_{m=1}^{N_{\text{eq}}} (\mathbf{G}_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t))$. Equation (31) is then used to update the network weights using any variant of (stochastic) gradient descent,[20] specifically,

$$\nu'_k = \nu_k - \alpha^{\text{learn}} \frac{\partial \ell(\nu)}{\partial \nu_k}, \tag{32}$$

where $\nu'_k$ denotes the updated $k$-th weight of the neural network, and $\alpha^{\text{learn}} \in \mathbb{R}$ is the learning rate. The updated neural network-based policy is subsequently used to simulate a new sequence of length $N_{\text{path length}}$, during which the loss function is recorded and again used to update the network parameters. This iterative process continues until $\ell_\nu < \epsilon \in \mathbb{R}$, indicating that an approximate equilibrium policy has been found.

To manage Markov-switching models, several modifications to the baseline DEQN method are necessary. In this context, the transition probabilities depend on the current Markov state, requiring an adjustment to our numerical integration routine. To illustrate, consider a model with one continuous shock $a_t$ and a two-state Markov chain $s_t \in \{0, 1\}$ with transition probabilities $\pi(i|j)$. We first employ quadrature to approximate the expectation computation with respect to the continuous shock. Let $w_i$ denote the weights and $x_i$ the nodes for the quadrature. For a function $f(a_t, s_t)$, the expectation is then computed as:

$$\mathbb{E}\left[f(a_{t+1}, s_{t+1})|s_t\right] \approx (1 - s_t) \sum_{l\in\{0,1\}} \sum_i \pi(l|0) w_i f(x_i, l) + s_t \sum_{l\in\{0,1\}} \sum_i \pi(l|1) w_i f(x_i, l).$$

Note that while the nodes are fixed, the weights depend on the state. This allows the expectation operation to be computed in batches, thereby reducing runtime.

In all our numerical experiments, we use a neural network architecture consisting of two hidden layers, each with 512 nodes, and SELU activation functions. Only the input and output dimensions vary depending on the model, as detailed in the respective sections below. Training is performed using the Adam optimizer with an initial learning rate of $1 \times 10^{-5}$, initially in single precision, and then resumed in double precision.

# C Additional figures

In this Appendix, we provide additional illustrations of the paper's main results. Figure 11 depicts a stylized neural network. Figure 12 shows the differences in the response to a larger persistent shock.

[20] In our applications, we employ the "Adam" optimizer (Kingma and Ba, 2014).

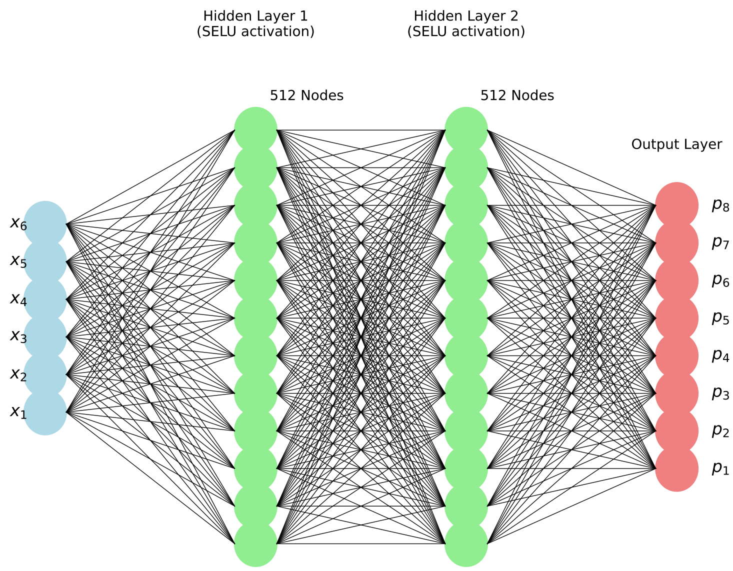


Figure 11: The figure above depicts a stylized neural network with an input $\mathbf{x} \in \mathbb{R}^6$. It consists of two hidden layers, each containing 512 neurons, using SELU activation functions, and produces an output $\mathbf{p}(\mathbf{x}) \in \mathbb{R}^8$.

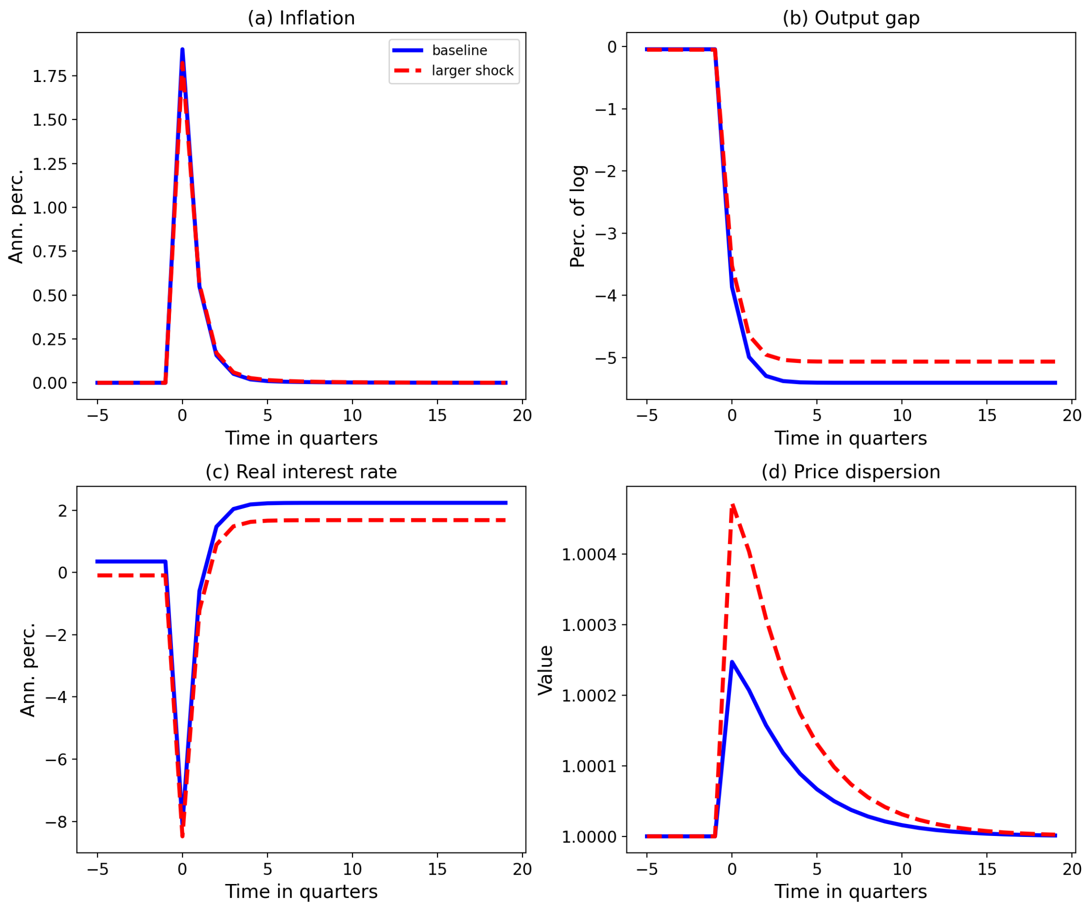


Figure 12: Response to a regime change: shock size.

*Note*: The figure displays the transition from a normal times regime to a bad times regime in the case of the optimal policy under commitment in the baseline (solid blue line) and with a larger shock (dashed red line). We set the innovations to the temporary shocks at zero. The economy starts at the SSS of the normal times regime in each case.