Monetary Policy with Persistent Supply Shocks*

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

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

March 23, 2025

Abstract

This paper studies monetary policy in a New Keynesian model with persistent supply shocks, that is, sustained increases in production costs due to factors such as wars or geopolitical fragmentation. First, we demonstrate that Taylor rules fail to stabilize long-term inflation due to endogenous shifts in the natural interest rate. Second, we analyze optimal policy responses under discretion and commitment. Under discretion, a systematic inflationary bias emerges when the shock impacts the economy. Under commitment, the optimal policy adopts a lean-against-the-wind approach without compensating for past inflation, implying that "bygones are bygones". The addition of the zero lower bound reinforces this result.

JEL Classification: E32, E58, E63.
Keywords: Deep Learning, Markov switching model, cost-push shocks.

*This paper was partially completed while Galo was visiting the BIS. We are grateful to Ryan Banerjee, Jose Elias Gallegos, Boris Hofmann, Joël Marbet, 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, particularly in the wake of the recent inflationary period (e.g., post-2021). This episode has prompted academic researchers and central banks to re-evaluate the appropriate responses of monetary policy, especially when inflation is driven by supply-side factors.[1] A critical question emerging from this debate is how monetary policy should be formulated in the face of persistent supply shocks.[2] By persistent supply shocks, we refer to sustained increases in production costs due to factors such as wars or geopolitical fragmentation, which can persist over years or even decades.[3] The presence of these shocks challenges the validity of traditional monetary policy prescriptions, typically derived under the assumption of small temporary shocks.

This paper investigates monetary policy responses to persistent disruptions by developing a New Keynesian model that incorporates a novel, persistent cost-push shock, thereby extending the standard framework (Clarida et al., 1999, Woodford, 2003). In this model, a representative household consumes a continuum of differentiated goods and supplies labor in a centralized,

[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.

frictionless market. Each good is produced by a single firm using labor as 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.

The novelty of our model lies in the introduction of a persistent cost-push shock in addition to the standard autoregressive "temporary"shocks. The economy randomly switches between two regimes: "normal times", when the persistent shock does not affect the economy, and "bad times", when the shock increases production costs for all firms.[4] This addition complicates the model solution, as it requires solving the model globally with regime-switching dynamics.[5] We address this issue by proposing a novel algorithm based on deep learning. The algorithm extends the "deep equilibrium nets"methodology of Azinovic et al. (2022) to accommodate a Markov-switching environment. This methodological advancement enables us to approximate nonlinear functions and efficiently solve the complex model, enabling the analysis of optimal monetary policy responses under different regimes.

Our contributions are threefold: First, we demonstrate that traditional Taylor rules fail to stabilize inflation in both regimes due to shifts in the natural interest rate driven by precautionary savings behavior. Specifically, the Taylor rule sets the long-term real interest rate equal to the average natural rate across regimes, which proves too restrictive during normal times and too lax during bad times. Consequently, inflation systematically deviates from the target, underscoring the limitations of existing policy rules that assume a constant natural rate.

We uncover a new precautionary-savings mechanism through which the persistent cost-push shock affects the natural rate. During normal times, the economy is undistorted, and consumption

[4]This can be interpreted as a proxy for various types of shocks that increase costs. For example, Afrouzi et al. (2023) consider a labor wedge similar to the one used here as a proxy for changes in the labor market composition towards more regulated labor sources or a deceleration in globalization.

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

aligns with that of efficient allocation; thus, the output gap is zero. In contrast, during bad times the average markup becomes suboptimal, leading to a reduction in output and consumption which results in a negative output gap. Consequently, the economy features two distinct stochastic steady states (SSSs), each associated with a different regime.[6] The variation in consumption between the two regimes explains the dynamics of the natural interest rate.[7] In normal times, households anticipate a potential transition to a regime where their average consumption would decline, leading to a precautionary increase in the demand for savings. Given the fixed supply of savings instruments, this results in a decline in the natural rate. Conversely, when the economy transitions to the bad times regime, consumption falls, and the demand for savings decreases as households anticipate higher future consumption once the regime ends, causing the natural rate to rise.

Second, we analyze optimal monetary policy under both discretion and commitment. Under discretion, the central bank cannot commit to future actions but understands how it can affect future decisions by changing price dispersion. 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.

Under commitment, the central bank can credibly commit to a state-contingent plan. The optimal policy "leans against the wind" (Galí, 2008), allowing inflation to increase once the persistent shock arrives but progressively tightening monetary policy so that inflation returns to zero a few quarters later. The central bank thus partially cushions the negative consequences of the shock on the output gap, which converges progressively to the new, lower SSS.

The optimal response to the persistent cost-push shock markedly differs from the response to a

[6]An SSS refers to the equilibrium state when shocks are zero, and the economy remains within its current regime, with agents anticipating the stochastic processes.

[7]The natural interest rate is defined as the real interest rate in each SSS corresponding to the long-term real interest rate when temporary shocks dissipate.

temporary cost-push shock. In the latter case, the central bank commits to a deflationary period after the inflation spike to return the aggregate price level to its pre-shock value. With a persistent shock, no such deflationary period exists, and the price level permanently increases—"bygones are bygones". This result shows that the price-level targeting features of standard monetary policy prescriptions only hold in the case of autoregressive shocks, but they do not carry over to a more general framework in which shocks are driven by Markov-chain processes.[8]

Third, we extend the model to include the zero lower bound (ZLB) on nominal interest rates and show that it widens the gap between natural rates in the two regimes when the central bank implements the optimal policy under commitment. The ZLB prevents the central bank from providing the necessary macro stabilization, which increases precautionary savings in normal times.

The optimal response to the arrival of a persistent shock implies a higher increase in inflation and a more gradual policy tightening than in the counterfactual case without the ZLB. Thus, the optimal response inherits some of the "low for longer" features that characterize the optimal response to demand shocks at the ZLB (Eggertsson et al., 2003).

Our findings underscore the crucial importance of accounting for persistent supply shocks in monetary policy design, as they significantly influence the natural rate and require tailored interventions, especially when constrained by the ZLB. This study highlights the limitations of traditional policy rules and offers novel insights into optimal policy strategies across different economic regimes.

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.

[8]As temporary autoregressive shocks become more persistent, the deflationary period increases, but the level of negative inflation diminishes, ultimately vanishing in the limit.

Second, the paper relates to the extensive literature analyzing monetary policy in regime-switching models (e.g., Schorfheide, 2005; Davig and Doh, 2014; Blake and Zampolli, 2011; Bianchi and Melosi, 2017). These models 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, with the economy transitioning among them in response to permanent cost-push shocks.

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: persistent supply shocks may strongly affect natural rates through precautionary motives.

Fourth, this paper extends the literature analyzing optimal monetary policy at the ZLB (e.g., Eggertsson et al., 2003; Adam and Billi, 2006; Adam and Billi, 2007; Nakov et al., 2008; Daudignon and Tristani, 2023). First, we analyze optimal policy at the ZLB globally, which allows us to obtain the optimal inflation in both regimes. This complements the existing literature, such as Coibion et al. (2012) or Andrade et al. (2019), which analyze the optimal inflation target of a central bank following a Taylor rule. Second, our focus is on how the ZLB interacts with cost-push shocks, instead of demand shocks.

Fifth, our paper extends the emerging literature on the use of deep learning to solve high-dimensional general equilibrium models (e.g., Maliar et al., 2021; Han et al., 2021; Friedl et al., 2023; Gu et al., 2024).[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, providing new insights into how the ZLB interacts with persistent supply shocks —a topic not extensively explored in the existing literature.

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

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

The remainder of the paper is structured as follows. Section 2 introduces the model, while Section 3 details the calibration and our deep learning-based solution method. Section 4 examines the efficient allocation, the flexible price equilibrium, and the dynamics under Taylor rules. Section 5 analyzes optimal monetary policy, and Section 6 explores the implications of the ZLB. Finally, Section 7 concludes.

2 Model

In this section, we present the formal structure of our model. Time t is discrete, and there are three types of agents: households, firms, and the central bank. We begin by describing the households' problem in Section 2.1, followed by the firms' problem in Section 2.2, and finally characterize the central bank in Section 2.3. Details on the specification of the regimes, as well as the autoregressive processes driving the shocks in the model, are provided in Sections 2.4 and 2.5.

2.1 Households

Households consume goods c_t, and supply labor h_t to firms and 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:

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

where B_t are holdings of a nominal bond which 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 domestic 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 λ_t is the Lagrange multiplier associated with the budget constraint, and inflation is $\pi_t \equiv \frac{p_t}{p_{t-1}} - 1$.

We assume that the good is a basket of varieties indexed by j. Thus, the household also wants to choose the consumption of each variety $c_t(j)$ to minimize:

$$\min_{c_t(j)} \int_0^1 p_t(j) c_t(j) dj, \tag{4}$$

subject to:

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

where ϵ is the elasticity of substitution among varieties. The outcome of the optimization problem (4) is the optimal demand for j-th variety of good:

$$c_t(j) = \left(\frac{p_t(j)}{p_t} \right)^{-\epsilon} c_t. \tag{5}$$

Note that the optimization problem (4) implies that the price level is given by:

$$p_t = \left\{ \int_0^1 [p_t(j)]^{1-\epsilon} dj \right\}^{\frac{1}{1-\epsilon}}, \tag{6}$$

where $p_t(j)$ is the nominal price of j-th variety and p_t is the nominal price index for the basket.

2.2 Firms

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 temporary and persistent cost-push shocks, denoted by ξ_t and η_t, respectively. The firm's total cost is $(1-\bar{\tau})\,\eta_t w_t h_t\,(j)$. We define the labor wedge as

$$(1+\tau_t) \equiv (1-\bar{\tau})\,\xi_t\eta_t, \tag{7}$$

which lumps together the labor subsidy and the cost-push shocks.

Each of these firms has monopoly power over their respective variety and takes the demand for its 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 (5), where $\Psi\left(y_{t+k}(j)\right) \equiv (1+\tau_{t+k})\,w_{t+k}\left(\frac{y_{t+k}(j)}{A_{t+k}}\right)$ are total costs, and $\Lambda_{t,t+k} \equiv \beta^k \frac{\lambda_{t+k}}{\lambda_t}$ is the stochastic discount factor for payments between periods t and $t+k$. 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:

$$\frac{P_t^*}{p_t} = p_t^* = \mathcal{M} \frac{\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} y_{t+k} \left(p_{t+k}/p_t\right)^{\epsilon} \Psi'\left(y_{t+k}\right)}{\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \Lambda_{t,t+k} y_{t+k} \left(p_{t+k}/p_t\right)^{\epsilon-1}}. \tag{8}$$

Finally, we can express the numerator and denominator in (8) as

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

$$\Xi_t^D = y_t + \mathbb{E}_t \left[\theta \Lambda_{t,t+1} \left(1+\pi_{t+1}\right)^{\epsilon-1} \Xi_{t+1}^D\right]. \tag{10}$$

2.3 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 analyze monetary policy under a Taylor rule of the following form:

$$i_t = \frac{(1+\bar{\pi})}{\beta} - 1 + \psi \left(\pi_t - \bar{\pi}\right), \tag{11}$$

where $\bar{\pi}$ is the inflation target and ψ is the slope of the Taylor rule to deviations in inflation. In Sections 5 and 6, we consider instead optimal policy both under discretion and commitment.

2.4 Market clearing conditions

The market clearing conditions for goods are given by:

$$y_t(j) = c_t(j) + g_t,$$

where g_t is a government spending shock. By aggregating, we obtain:

$$y_t = \int y_t(j)\, dj = \int \left(c_t(j) + g_t\right) dj = c_t + g_t.$$

Since all firms face the same probability, θ, of keeping prices fixed, the law of large numbers ensures that a fraction θ of firms will keep their prices fixed, while the remaining fraction, $(1-\theta)$, will optimally reset their prices. As a result, equation (6) 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

$$1 = \theta \left(1 + \pi_t\right)^{\epsilon - 1} + \left(1 - \theta\right) \left(p_t^*\right)^{1-\epsilon}. \tag{12}$$

The market clearing condition for labor is given by:

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

We define price dispersion as:

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

such that the aggregate production function becomes:

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

2.5 Shocks and regimes

We define the temporary and persistent shocks. **Temporary shocks**. We consider TFP, government spending, and cost-push shocks, each following an AR(1) process in logs. First, we define:

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

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

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

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

and:

$$\log(\xi_t) = \left(1 - \rho^\tau\right) \left(-\frac{\left(\sigma^\tau\right)^2}{2}\right) + \rho^\tau \log(\xi_{t-1}) + \varepsilon_t^\tau,$$

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

Regimes. We assume that the persistent cost-push shock, η_t, evolves according to a two-state Markov chain. We consider two regimes. In regime 1, "good times", the value of η_t is one. In regime 2, "bad times", its value is $\bar{\eta} > 1$. This implies that the shock is only active during bad times. The transition probabilities are p_{12} from regime 1 to 2:

$$p_{12} = \mathbb{P}\left(\eta_t = \bar{\eta} \mid \eta_{t-1} = 1\right), \tag{13}$$

and p_{21} from 2 to 1

$$p_{21} = \mathbb{P}\left(\eta_t = 1 \mid \eta_{t-1} = \bar{\eta}\right). \tag{14}$$

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

3 Calibration and numerical method

3.1 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, γ, is set to 2, and the inverse of the Frisch elasticity, ω, 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.

If the central bank follows a Taylor rule, the inflation coefficient ψ is set to 2, and the inflation target $\bar{\pi}$ is zero. This makes the results comparable with the optimal monetary policy explored below.

The parameters of the TFP, government spending, and cost-push shocks are taken from Coibion et al. (2012). The value of the persistent cost-push shock is set to $\bar{\eta} = \mathcal{M}$, so that it fully offsets the optimal labor subsidy. The average duration of regime 1 ("normal times") is 48 quarters (12

Parameter		Value
Long-run productivity level	A	1
Inverse Frisch elasticity	ω	1
Inverse of intertemporal elasticity of substitution	γ	2
Discount factor	β	0.9975
Elasticity of substitution among varieties	ϵ	7
Government spending constant	$\bar{g}$	0.2
Calvo constant	θ	0.75
Taylor rule slope	ψ	2
Inflation target	$\bar{\pi}$	0
Labor subsidy	$\bar{\tau}$	$\frac{1}{\epsilon}$
Mean of cost-push shock during persistent supply shock	$\bar{\eta}$	$\mathcal{M} = \frac{\epsilon}{\epsilon-1}$
Transition probability from normal to negative supply times	p_{12}	1/48
Transition probability from negative supply to normal times	p_{21}	1/24
Persistence of TFP shock	ρ^A	0.99
Persistence of cost-push shock	ρ^τ	0.90
Persistence of government spending shock	ρ^g	0.97
Standard deviation of TFP shock	σ^A	0.009
Standard deviation of cost-push shock	σ^τ	0.0014
Standard deviation of government spending shock	σ^g	0.0052

Table 1: Key parameters of the model.

years), to capture a period encompassing a business cycle, giving $p_{12} = 1/48$, while the average duration of regime 2 ("persistent supply shock") 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.

3.2 Deep Equilibrium Nets

A global solution of our model is crucial, as it 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 networks (DEQNs) approach of Azinovic et al. (2022), adapting it to effectively handle Markov-switching models. In what follows, 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.[11]

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:

$$\mathbf{G}(\mathbf{x}, \mathbf{p}) = \mathbf{0}, \ \forall \mathbf{x} \in X \subset \mathbb{R}^d. \tag{15}$$

Intuitively, DEQNs operate as follows: An unknown policy function is approximated using a neural network, denoted as $\mathbf{p}(\mathbf{x}) \approx \mathcal{F}_\nu(\mathbf{x})$, with trainable parameters ν. 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).[12] 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 ℓ-th layer containing N_ℓ neurons. In our case, $N_0 = d_{\text{input}}$ and $N_L = d_{\text{output}}$.[13] 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 ℓ-th layer, for some non-linear activation function $\sigma(\cdot)$ (such as ReLU, Swish, or SELU), an FNN is expressed as:

$$\mathbf{p}\left(\mathbf{x}\right) \approx \mathcal{F}_\nu\left(\mathbf{x}\right) = \mathcal{N}^L\left(\mathbf{x}\right) = A^L \circ \sigma_{L-1} \circ A^{L-1} \circ \cdots \circ \sigma_1 \circ A^1\left(\mathbf{x}\right). \tag{16}$$

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

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

[12] 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.

[13] 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 5.1).

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 $\mathbf{p}(\mathbf{x})$ begins by randomly initializing the parameters ν, representing an initial guess for the unknown approximate policy function. Next, a sequence of $N_{\text{path length}}$ states is 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, $\mathcal{F}_\nu(\mathbf{x})$, combined with the remaining model-implied dynamics.

If the (approximate) policy function satisfying the equilibrium conditions were known, equation (15) would hold along the simulated path. However, since the neural network is initialized with random coefficients, $\mathbf{G}(\mathbf{x}_t, \mathcal{F}_\nu(\mathbf{x}_t)) \neq 0$ 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:

$$\ell_\nu = \frac{1}{N_{\text{path length}}} \sum_{\mathbf{x}_t \text{ on sim. path}} \sum_{m=1}^{N_{\text{eq}}} \left(\mathbf{G}_m(\mathbf{x}_t, \mathcal{F}_\nu(\mathbf{x}_t))\right)^2, \tag{17}$$

where $\mathbf{G}_m(\mathbf{x}_t, \mathcal{F}_\nu(\mathbf{x}_t))$ represents each of the N_{eq} first-order equilibrium conditions of the model, i.e., $\mathbf{G}(\mathbf{x}_t, \mathcal{F}_\nu(\mathbf{x}_t)) = \sum_{m=1}^{N_{\text{eq}}} (\mathbf{G}_m(\mathbf{x}_t, \mathcal{F}_\nu(\mathbf{x}_t))$. Equation (17) is then used to update the network weights using any variant of (stochastic) gradient descent,[14] specifically,

$$\nu_k' = \nu_k - \alpha^{\text{learn}} \frac{\partial \ell(\nu)}{\partial \nu_k}, \tag{18}$$

where ν_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.

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

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 1×10^{-5}, initially in single precision, and then resumed in double precision.

4 Response to persistent supply shock under a monetary policy rule

In this Section, we focus on the case in which the central bank implements its monetary policy via a Taylor rule. We first introduce the efficient and flexible-price allocations as benchmarks, and then discuss how the presence of persistent supply shocks affect macro dynamics.

4.1 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:[15]

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

Combining this result with the aggregate budget constraint

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

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

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

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 $\hat{r} = \frac{1}{\beta} - 1$, which corresponds to 1 percent in annual terms. 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

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

rate of transformation, as

$$h_t^{*\omega} c_t^{*\gamma} = \frac{A_t}{(1+\tau_t)\,\mathcal{M}} = \frac{A_t}{\eta_t \xi_t \,(1-\bar{\tau})\,\mathcal{M}} = \frac{A_t}{\eta_t \xi_t},$$

where, in the last equality, we apply the fact that the labor subsidy neutralizes the average markup.[16] Here, $\eta_t \xi_t$ represents the cost-push shock, which has a mean of one in the normal-times regime, but a mean of $\mathcal{M}$ in the persistent supply shocks regime. Consumption in this case satisfies the equation

$$c_t^* + g_t = y_t^* = A_t h_t^* = A_t^{\frac{1}{\omega}+1} (\eta_t \xi_t)^{-\frac{1}{\omega}} c_t^{*-\frac{\gamma}{\omega}}. \tag{19}$$

This can be rewritten as

$$\left(\frac{c_t^* + g_t}{A_t}\right)^{\omega} - \frac{A_t}{(\eta_t \xi_t) c_t^{*\gamma}} = 0. \tag{20}$$

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 temporary and persistent shocks are zero, and the temporary shocks remain at their mean values. This is an adaptation of the standard concept of stochastic steady state (SSS) to the case of a Markov-switching model, which is instrumental in understanding model dynamics. Given that the persistent shock has two different values, the economy exhibits two SSSs, one in each regime.

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, this equation implies

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

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

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

$$1 + r^*_{1,ss} \approx \frac{1}{\beta} \frac{{c^*_{2,ss}}^{\gamma}}{\left(p_{12}{c^*_{1,ss}}^{\gamma} + (1 - p_{12})\, {c^*_{2,ss}}^{\gamma}\right)}, \tag{21}$$

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 (19):

$$c^*_{n,ss} + \bar{g} = (\eta_n \xi_{n,ss})^{-\frac{1}{\omega}} (c^*_{n,ss})^{-\frac{\gamma}{\omega}},$$

where $\eta_n \xi_{n,ss}$ equals one in regime 1, and $\mathcal{M}$ in regime 2. The values are $c^*_{1,ss} = 0.9377$ and $c^*_{2,ss} = 0.8877$. Since $c^*_{2,ss} < c^*_{1,ss}$, the denominator on the right-hand side of equation (21) 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 shown in the first column of Table 2.

The large 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.

4.2 The economy under a standard Taylor rule

Ergodic distribution and SSSs. In the baseline case with nominal rigidities and a Taylor rule of the form (11), the central bank controls nominal interest rates to steer the economy towards an inflation level $\bar{\pi}$. Figure 1 shows the ergodic distribution in this case. It 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 the persistent supply shock regime.

Several results emerge. First, the considered variables (inflation, output gap, nominal and real interest rates) exhibit *bimodality*: the distribution of realizations clusters around two distinct points. These two points correspond to the stochastic steady states (SSS) of each variable, as reported in the second column of Table 2.

	Flex. prices	Taylor rule	Mod. Taylor rule
Inflation			
normal times	0.0%	-0.9%	0.0%
bad times	0.0%	1.6%	0.0%
Output gap			
normal times	0.0%	-0.1%	0.0%
bad times	-5.5%	-5.4%	-5.5%
Real interest rates			
normal times	0.0%	0.1%	0.0%
bad times	2.7%	2.6%	2.7%
Nominal interest rates			
normal times	0.0%	-0.8%	0.0%
bad times	2.7%	4.2%	2.7%

Table 2: Stochastic steady state values.

Second, long-term inflation in this model is not zero. In the normal times regime it is negative (-0.9 percent) and in the bad times regime it is positive (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

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

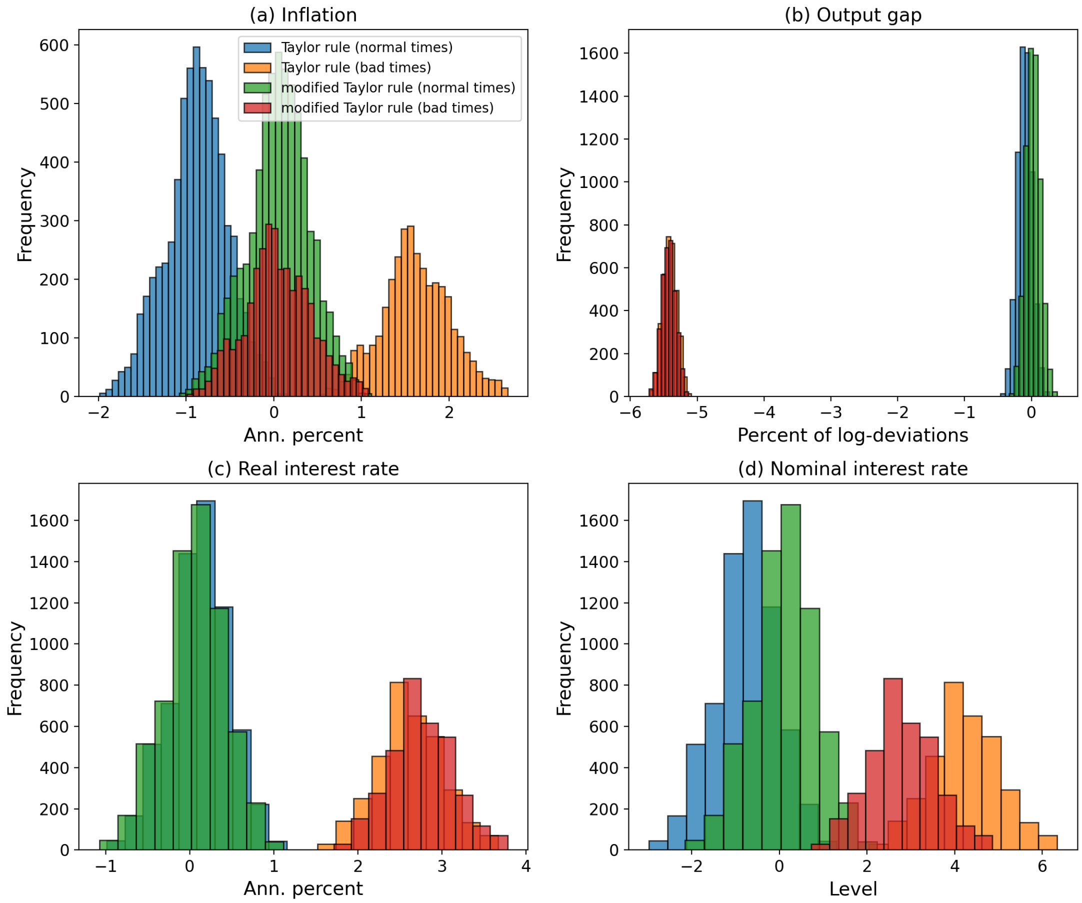

Figure 1: Ergodic distribution

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.

where $\bar{r}$ is the real rate target of the central bank, which, under the Taylor rule (11), 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

$$\pi_{n,sss} \simeq \bar{\pi} + \frac{r_{n,ss} - \bar{r}}{\psi - 1}. \tag{22}$$

Equation (22) 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 (22) on average: $0.9\% - 1\% = -0.1\%$.

Third, real rates are slightly higher than natural rates in normal times and lower in the persistent shocks regime. This small divergence between the long-term real rates and the natural rates is due to the different values of SSS consumption. Compared to the flex-price allocation, consumption is lower in normal times and higher in bad times: the SSS values are $c_{1,ss} = 0.9368$ and $c_{2,ss} = 0.8883$. Following the logic of equation (21) again, the jump in consumption between the two regimes becomes slightly narrower, and so does the gap in real interest rates.

The difference in consumption concerning the flex-price allocation is a consequence of the long-term impact of non-zero inflation on markups. Combining equations (8) and (12) in steady state, we have

$$w_{n,ss} = A_t \left(\eta_n \xi_{n,ss}\right)^{-1} \left(\frac{1 - \theta \left(1 + \pi_{n,ss}\right)^{\epsilon - 1}}{(1 - \theta)} \right)^{\frac{1}{1-\epsilon}} \frac{1 - \theta\beta \left(1 + \pi_{n,ss}\right)^{\epsilon}}{1 - \theta\beta \left(1 + \pi_{n,ss}\right)^{\epsilon - 1}}.$$

In the flex-price allocation, $\theta = 0$ and this expression simplifies to $w_{n,ss} = A_t \left(\eta_n \xi_{n,ss}\right)^{-1}$, such that in normal times wages coincide with those in the efficient allocation, $w_{1,ss} = A_t$, whereas in the bad times regime they are distorted $w_{1,ss} = A_t \left(\bar{\eta}\right)^{-1} < A_t$, which leads to lower labor and consumption in equilibrium. In the economy with nominal rigidities, long-term wages are also affected by long-term inflation. This distortion operates *against* the effect of persistent cost-push shocks. In the normal times regime, there is no distortion due to cost-push shocks. Still, negative inflation introduces an additional distortion on wages, which leads to slightly lower output and consumption. In the bad-times regime, however, the distortion due to cost-push shocks is high, and positive inflation mitigates it to a limited extent, thus marginally increasing consumption. This can be confirmed by comparing the values of the output gap, -0.1 percent and -5.4 percent, with those under flexible prices, 0 percent and -5.5 percent. In both regimes, non-zero inflation increases price dispersion, but this effect is second order compared to that on the average markup.

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 deglobalization process.

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

$$i_t = \bar{r}_t + \psi \left(\pi_t - \bar{\pi} \right), \tag{23}$$

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 (22) is compatible with a zero inflation target $\bar{\pi} = 0$. The third column in Table 2 confirms this: inflation is zero in both SSS and real rates and output gaps coincide with those in the flex-price allocation. The red and green bars in Figure 1 display the ergodic distribution under this modified Taylor rule. They are again centered on the SSS values. In particular, 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 persistent supply shocks the central bank

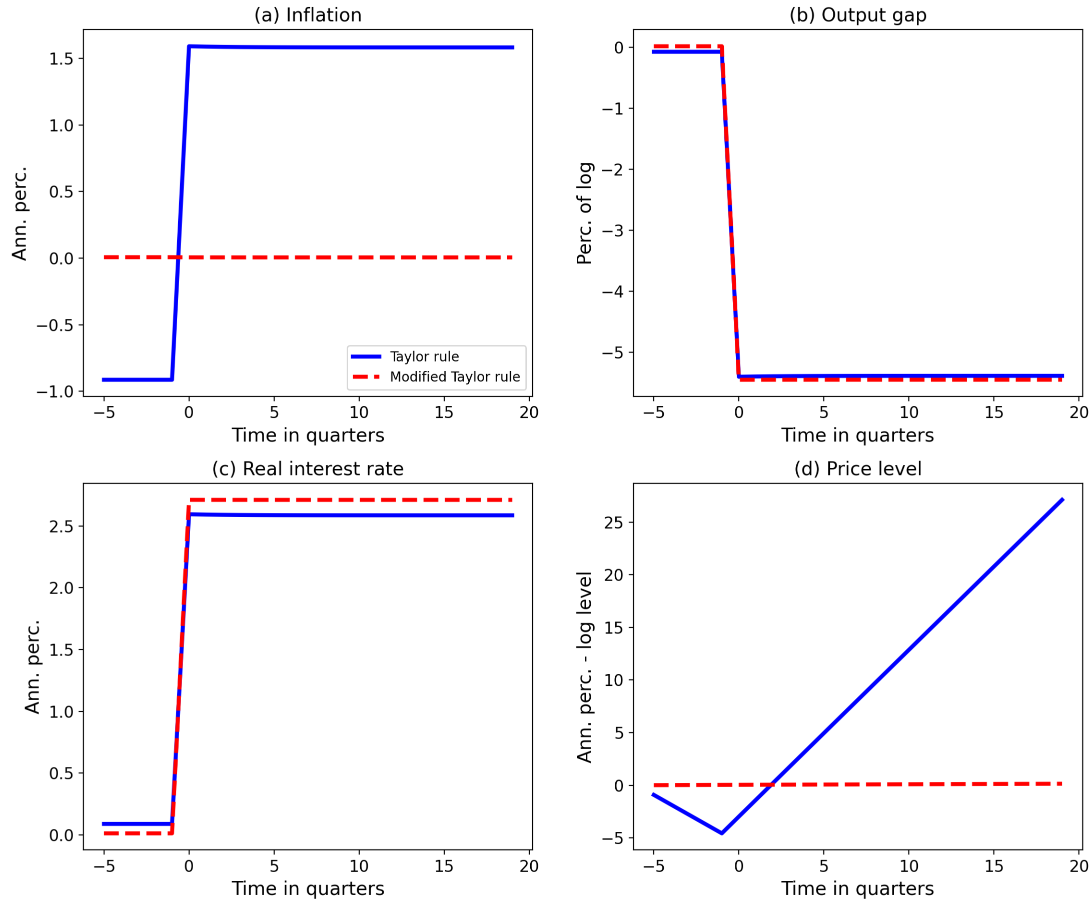

Figure 2: Response to a regime change (Taylor rule).
Note: The figure displays the transition from a normal times regime to a bad times regime in the case of a Taylor rule (solid blue line) and modified Taylor rule (dashed red line). We set the innovations to the temporary shocks at zero. The economy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

should endogenously adapt its interest rate target to track the natural rate, which becomes a regime-contingent object.

Figure 2 displays the transitional dynamics after a regime change. It assumes that the economy starts at the SSS of the normal-times regime and, at time zero, a transition happens to the persistent-supply-shock regime, leaving the realization of all shocks at zero.

Under the Taylor rule (solid blue line), inflation, the output gap, and the real interest rate (panels a-c) jump after the shock arrival and remain constant at their new SSS values. The price level (panel d) changes its normal-times deflationary trend to an inflationary trend. The modified Taylor rule (dashed red line) corrects the latter effect by keeping inflation anchored at zero.

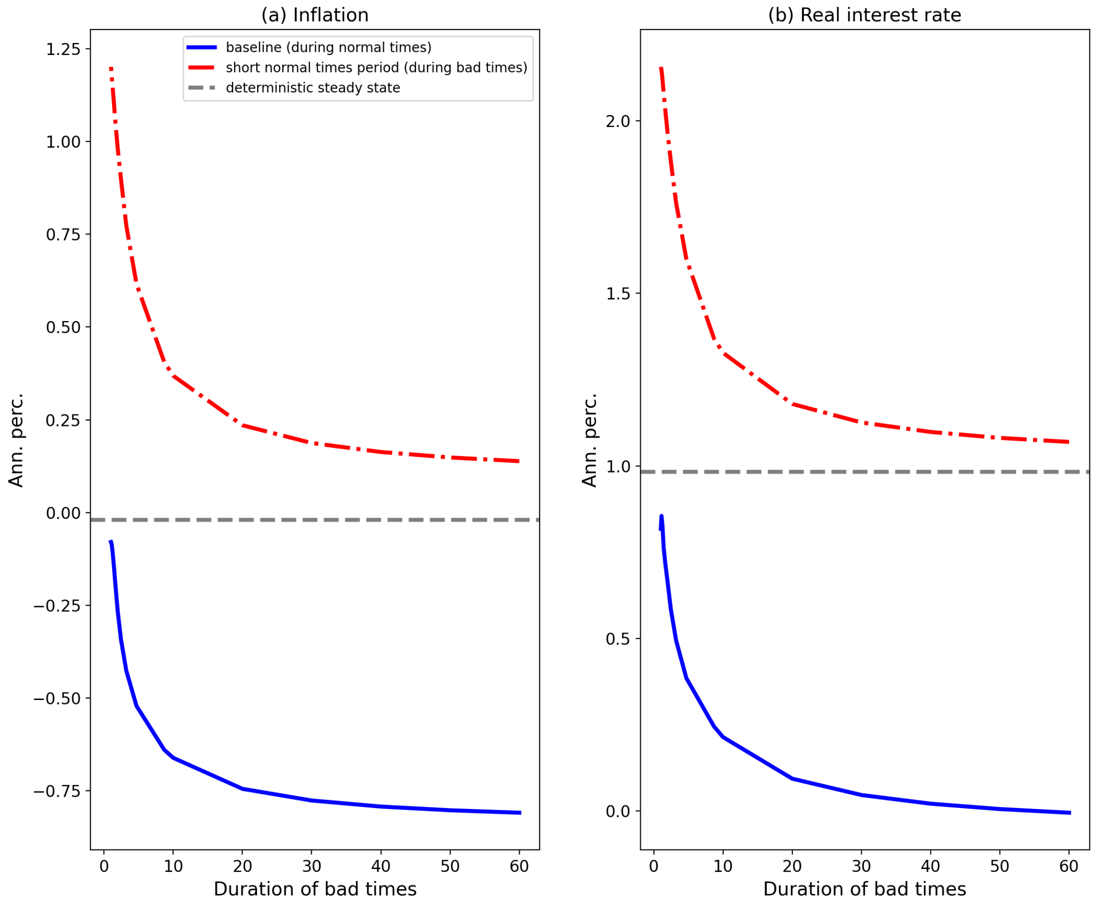

Figure 3: Sensitivity to regime length (Taylor rule).
Note: The figure displays the SSS values of inflation and the real interest rate for alternative values of the average duration for the bad times regime. The blue line displays the normal-times SSS values in the baseline model and the dashed red line 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.

Sensitivity to regime length. Figure 3 displays the sensitivity of SSS inflation and real rates to alternative average lengths of the persistent supply shock period under the Taylor rule (11). 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 persistent cost-push shock. The baseline calibration implies an average duration of 48 quarters. The SSS values of inflation and the real interest rate are very nonlinear with respect to changes in the average ratio of the persistent supply shock. 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 3 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.

5 Optimal monetary policy response to persistent supply shocks

We turn next to analyze optimal monetary policy under both discretion and commitment. In Section 5.1 we introduce each of these problems, and in 5.2 we discuss the main results.

5.1 Problem Statement

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 Δ_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. This problem, therefore, includes one endogenous state variable, three exogenous variables, and a regime change. The value function satisfies the Bellman equation, that is,

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

subject to the equilibrium conditions

$$c_t^{-\gamma} = h_t^{\omega}/w_t, \tag{24}$$

$$1 = \theta\left(1+\pi_t\right)^{\epsilon-1} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}, \tag{25}$$

$$\Delta_t = \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}, \tag{26}$$

$$y_t = A_t h_t \left(\Delta_t\right)^{-1}, \tag{27}$$

$$y_t = c_t + g_t. \tag{28}$$

$$p_t^* = \mathcal{M}\frac{y_t w_t\left(1+\tau_t\right)\left(A_t\right)^{-1} + \mathbb{E}_t\left[\theta\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right]}{y_t + \mathbb{E}_t\left[\theta\Lambda_{t,t+1}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right]}. \tag{29}$$

Under discretion, the central bank cannot commit to future policy paths. It, however, understands how its policies may affect price dispersion, which in turn affect 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.

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

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

subject to the equilibrium conditions (24)-(28) and the constraints

$$\begin{aligned}
p_t^*\Xi_t^D &= \mathcal{M}\Xi_t^N, \\
\Xi_t^N &= \left(c_t+g_t\right)^{1+\omega}\left(\frac{\Delta_t}{A_t}\right)^{\omega} c_t^{\gamma}\left(1+\tau_t\right)\left(A_t\right)^{-1} + \mathbb{E}_t\left[\beta\theta c_t^{\gamma}c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon}\Xi_{t+1}^N\right], && (30)\\
\Xi_t^D &= \left(c_t+g_t\right) + \mathbb{E}_t\left[\beta\theta c_t^{\gamma}c_{t+1}^{-\gamma}\left(1+\pi_{t+1}\right)^{\epsilon-1}\Xi_{t+1}^D\right]. && (31)
\end{aligned}$$

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 (30)-(31).

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 (30) and (31), evolve around their SSS values.[17]

5.2 Optimal response to persistent supply shock

Ergodic distribution and SSSs. Similar to our approach in Section 4, we first analyze the ergodic distribution of the main macroeconomic variables under the optimal policy. Figure 4 displays the ergodic distributions in the case of discretion. The results show that inflation, the output gap, and real interest rates still exhibit bimodal distributions, as in the standard Taylor rule discussed earlier, but the distribution differs markedly from the previous case.

First, inflation clusters around two points: zero inflation and 2.6 percent inflation. This is confirmed by the SSS analysis, shown in the first column of Table 3. Compared to the Taylor rule, inflation in normal times now clusters around zero, as the central bank correctly identifies the long-run natural rate in this regime and thus achieves its inflation target. This contrasts with the bad times regime, where inflation centers around 2.6 percent; a value significantly higher than under a Taylor rule. This difference is not due to a systematic error by the central bank in assessing the natural rate, as was the case with the Taylor rule, but rather 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 persistent supply shock, as it offsets the optimal labor subsidy leading to a distorted average markup.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 incorporates it into their expectations, leading to higher prices. The outcome is that the economy ends up with a systematically higher level of inflation due to the central bank's lack of commitment.

The situation is quite different in the case of optimal monetary policy under commitment (cf.

[17] 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.

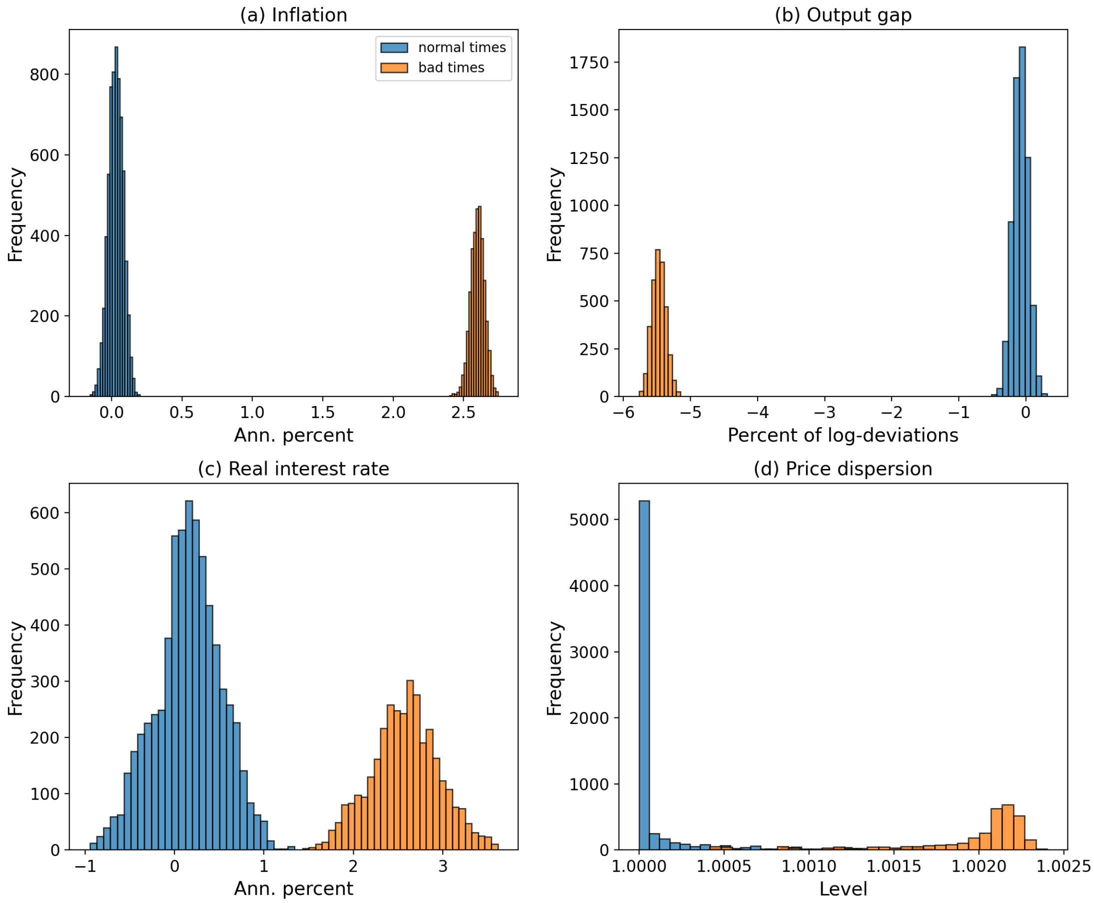

Figure 4: Ergodic distribution: discretion.

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.

Figure 5). Inflation is now tightly centered around zero in both regimes, as the central bank can credibly commit to maintaining long-run inflation at zero. The SSS values of the real interest rate differ from the natural rates, as shown in the second column of Table 3. This occurs because consumption is higher in the bad times regime compared to the flexible-price equilibrium, and lower in normal times. The reason is that, as we will see next, monetary policy provides some cushion during the dynamic transition between the two regimes, thereby reducing precautionary saving motives.[18]

[18]It is for this reason that all cases with nominal rigidities, whether under Taylor rules or optimal policy, show lower real interest rates between the two regimes compared to the natural rates. However, the effect is most pronounced in the case of the optimal policy under commitment.

	Discretion	Commitment
Inflation		
normal times	0.0%	0.0%
bad times	2.6%	0.0%
Output gap		
normal times	0.0%	0.0%
bad times	-5.5%	-5.4%
Real interest rates		
normal times	0.1%	0.4%
bad times	2.6%	2.2%
Nominal interest rates		
normal times	0.1%	0.4%
bad times	5.2%	2.2%

Table 3: Stochastic steady state values.

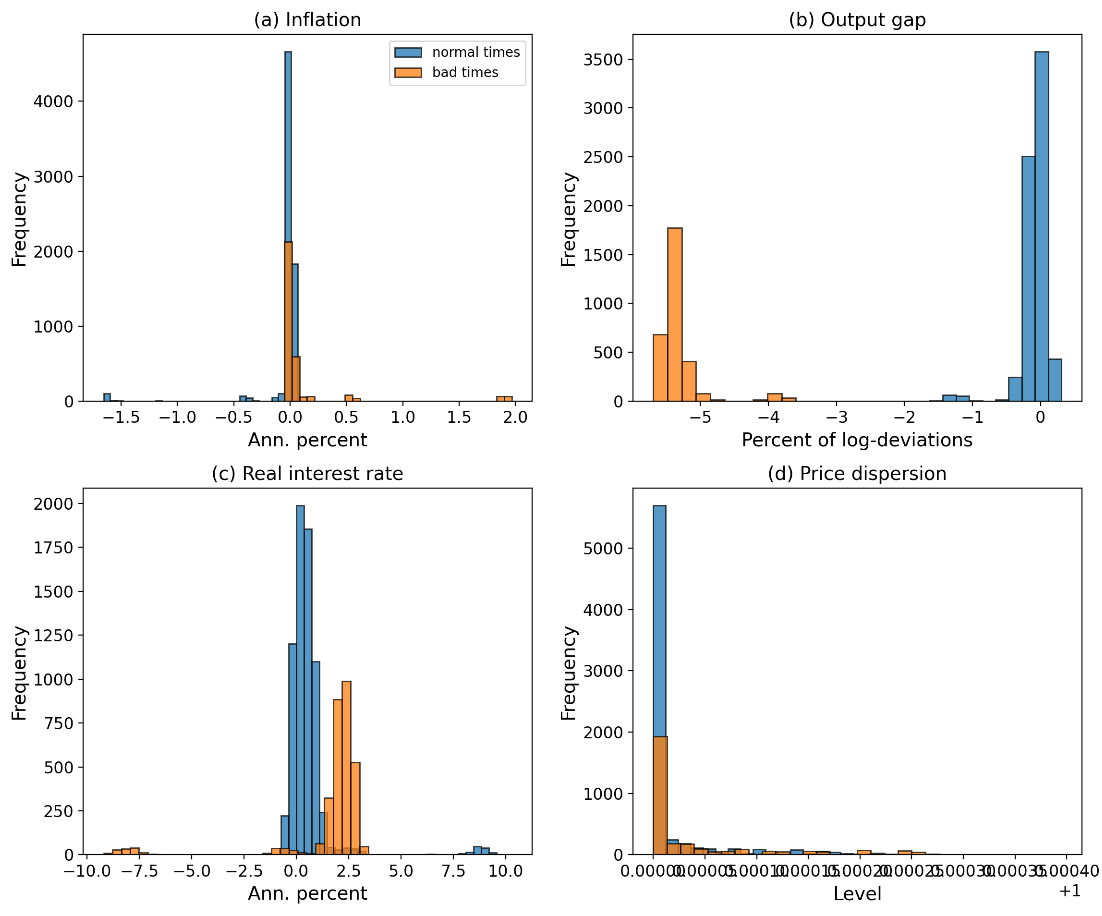

Figure 5: Ergodic distribution: commitment.

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 to bad times. The figure is produced by simulating the model for a large number of periods.

Dynamics. Figures 4 and 5 also illustrate that the variance of inflation is quite low in both models compared to the cases with a Taylor rule, as the central bank now optimally responds to shocks. This is especially pronounced under commitment, as the "divine coincidence" (Blanchard and Galí, 2007) holds, and the central bank responds to TFP and government spending shocks by maintaining inflation at zero.[19] This contrasts with the case of the modified Taylor rule (23) above, where the central bank responds following the same rule, regardless of the origin of the shock and the variance of macroeconomic variables is larger.

In the case of temporary (autoregressive) cost-push shocks, the prescription in the New Keynesian model is to "lean against the wind" (Galí, 2008), tolerating a rise in inflation to partially cushion the fall in the output gap. This differs from the dynamics in the flexible-price equilibrium, where inflation is always zero. By tolerating the rise in inflation, the central bank aims to minimize the distortions associated with the temporary increase in markups, albeit at the cost of higher price dispersion. When the central bank can commit, it ensures that the initial rise in inflation is later offset by a deflationary period, bringing the price level back to its pre-shock position. The solid blue line in Figure 6 shows the optimal response to a temporary AR(1) shock when the economy is in the bad times regime. The optimal policy exhibits this price-level targeting feature, as the price level returns to its initial value. This contrasts with the modified Taylor rule (dashed red line), where "bygones are bygones" and the price level increases permanently. Compared to the modified Taylor rule, the strategy under commitment results in lower inflation and reduced output gap volatility.

The optimal prescription regarding cost-push shocks changes when these shocks become persistent. Figure 7 shows the transition from normal times to the persistent-supply-shock regime, similar to Figure 2, but this time comparing discretion and commitment.

Two conclusions emerge. First, in the case of *discretion*, the transition to the high-inflation SSS occurs almost immediately after the regime shift, without any overshooting of inflation. The only exception is price dispersion, which progressively evolves towards its higher SSS value. This result complements the findings in Afrouzi et al. (2023), who analyze the transitional dynamics

[19]The divine coincidence only holds in normal times, as the SSS in the persistent supply shock regime is distorted, and government purchases are present (see Benigno and Woodford, 2005).

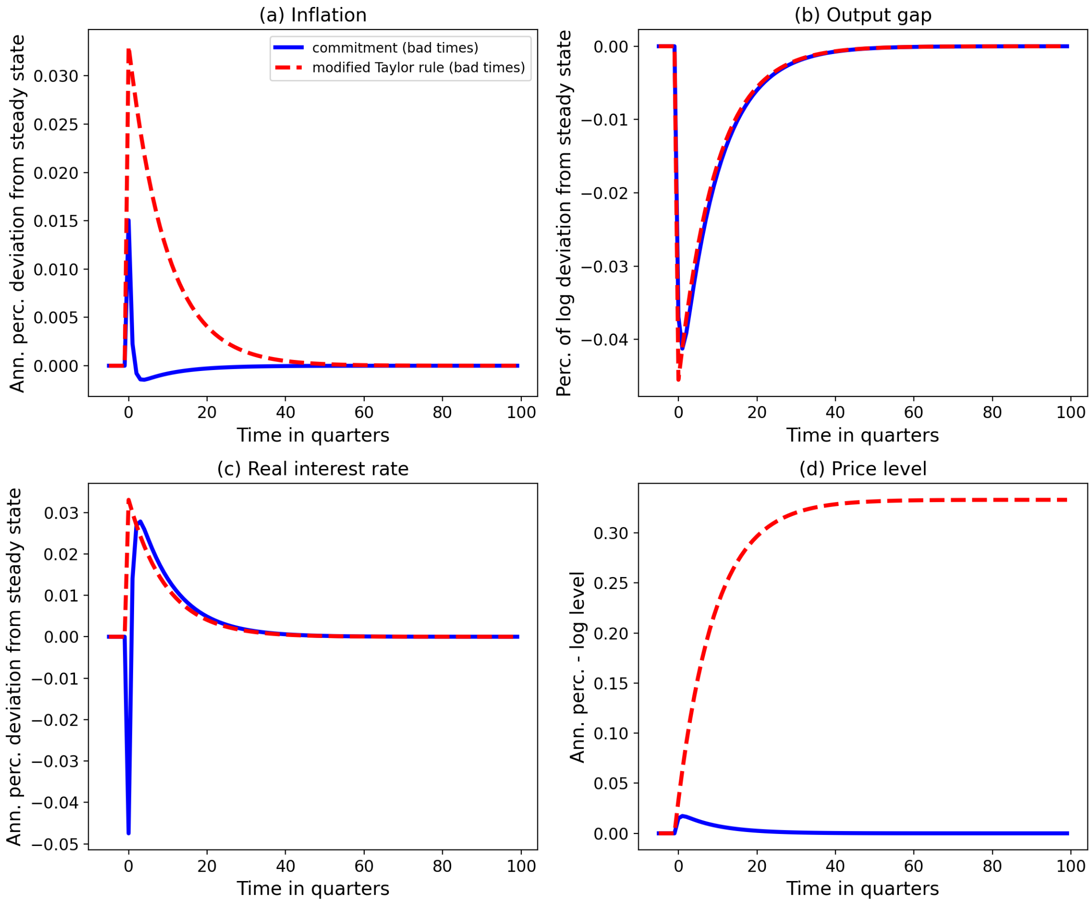

Figure 6: Impulse response to a transitory cost-push shock (commitment versus Taylor rule).
Note: The figure displays the impulse response to a temporary AR(1) cost-push shock when the economy is in the bad times regime. The figure displays the response in the case of an optimal policy under commitment (solid blue line) and with a modified Taylor rule (dashed red line).

under discretion from a low to a high non-stochastic steady state in a similar model. The main difference between the two studies is that while risk is anticipated here, Afrouzi et al. (2023) examine transitional dynamics following a one-off, unanticipated shock under perfect foresight. They find that inflation overshoots before converging to the high-inflation steady state. In our model, there is only one non-stochastic steady state, and we do not consider transitional dynamics.

Second, the optimal policy under *commitment* tolerates a temporary increase in inflation in response to a permanent cost-push shock. This is similar to the case with temporary shocks, but the key difference is that inflation now reverts back to zero without a subsequent disinflation period. In other words, bygones are bygones, and the central bank does not attempt to reverse the increase in the price level.

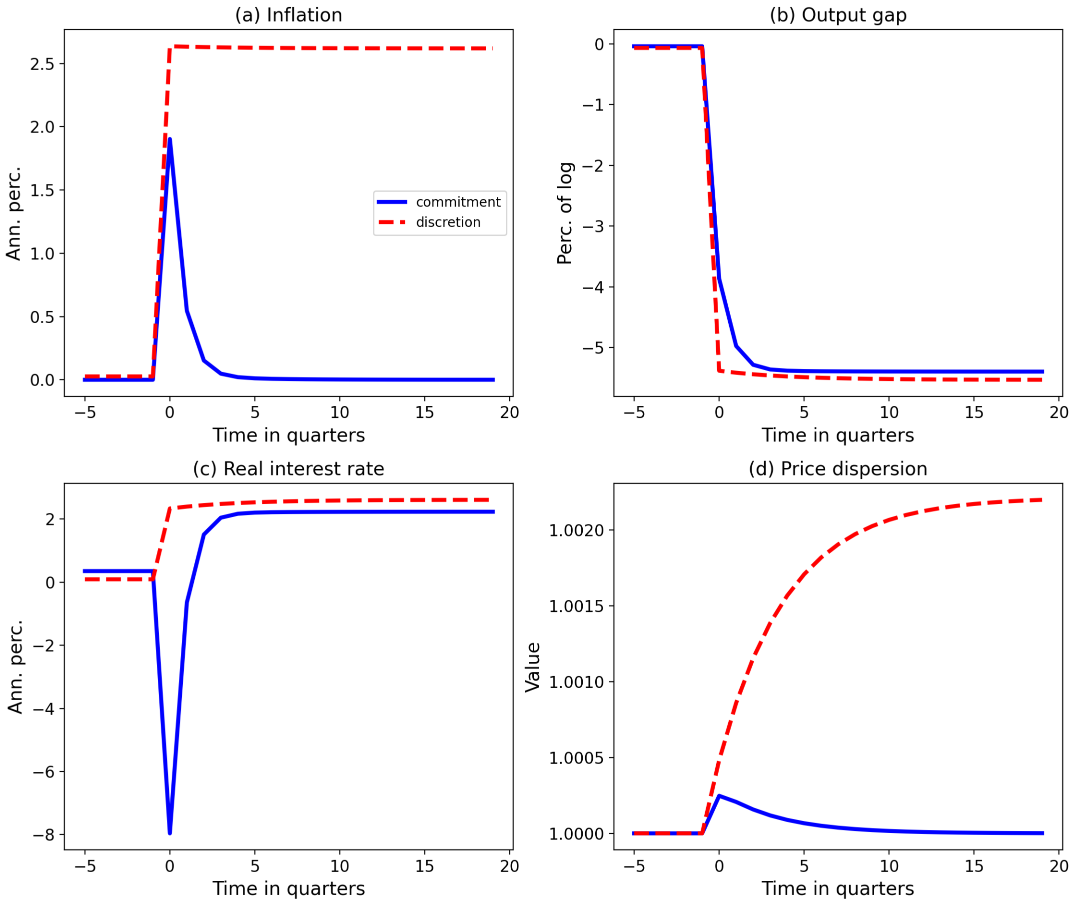

Figure 7: Response to a regime change: commitment versus discretion.
Note: The figure displays the transition from a normal times regime to a bad times regime in the case of the optimal policy under commitment (solid blue line) and discretion (dashed red line). We set the innovations to the temporary shocks at zero. The economy starts at the SSS of the normal times regime.

To provide some intuition for this result, Figure 8 plots the optimal response under commitment to two autoregressive cost-push shocks: one calibrated as in Table 1, and the other with a higher persistence, featuring an autoregressive coefficient of 0.99. The figure shows that, as 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:

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

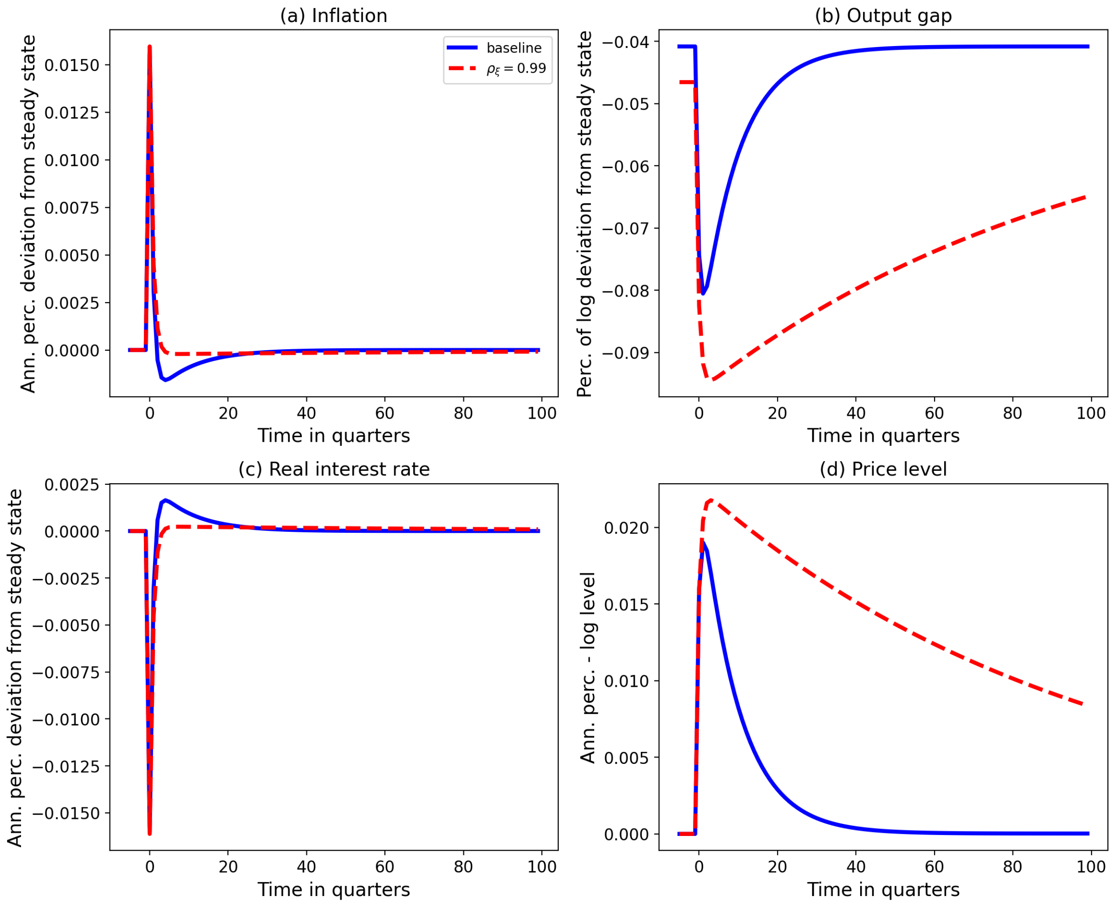

Figure 8: Impulse response to a transitory cost-push shock with different levels of persistence. *Note*: The figure displays the impulse responses to a temporary cost-push shock under the optimal policy with commitment. We consider two shocks of different autoregressive coefficients in a model without persistent cost-push shocks. The solid blue line is the baseline with $\rho^\tau = 0.90$ and the dashed red line is more persistent $\rho^\tau = 0.99$.

where x is the output gap. 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 and then remains constant, which implies that inflation increases on impact and then remains constant. 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 7.

It could be argued, however, that even if the central bank tolerates persistent deviations in the price level, the positive inflationary periods during the transition to the persistent-supply-shocks regime should offset the deflationary periods when the transition occurs in the opposite direction, thus keeping the average price level constant in the ergodic distribution. This is not the case, as shown in Figure 12 in Appendix B. The increase in inflation is greater than the decrease, resulting in an overall positive trend in prices.[20]

Robustness to shock size. Figure 14 in Appendix B 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 the baseline. The temporary increase in inflation is also more pronounced in this case, leading to a larger increase in price dispersion. Nonetheless, the key qualitative results remain unchanged. In particular, the optimal response still implies an increase in inflation that reverts back to zero after a few quarters, with no deflationary period.

6 Optimal policies with the zero lower bound

We now extend the model to include an occasionally-binding zero lower bound (ZLB) constraint, introducing an additional source of nonlinearity at the macroeconomic level. First, we analyze the results under Taylor rules in Section 6.1, and then we revisit the case of optimal monetary policy in Section 6.2.

6.1 Taylor rules

Standard Taylor rule. We consider the case in which the Taylor rule (11) includes a ZLB constraint:

$$i_t = \max\left\{\frac{(1+\bar{\pi})}{\beta} - 1 + \psi\left(\pi_t - \bar{\pi}\right), 0\right\}.$$

The ergodic distribution is displayed in Figure 9 in blue and orange, and the SSS values are shown in the first column of Table 4. In the absence of the ZLB, the nominal interest rate in the

[20] In the standard New Keynesian model, the optimal response to a large temporary cost-push shock deviates slightly from pure price-level targeting. This deviation is nevertheless tiny compared to the one introduced by the optimal response to a persistent cost-push shock, as displayed in Figure 13 in Appendix B.

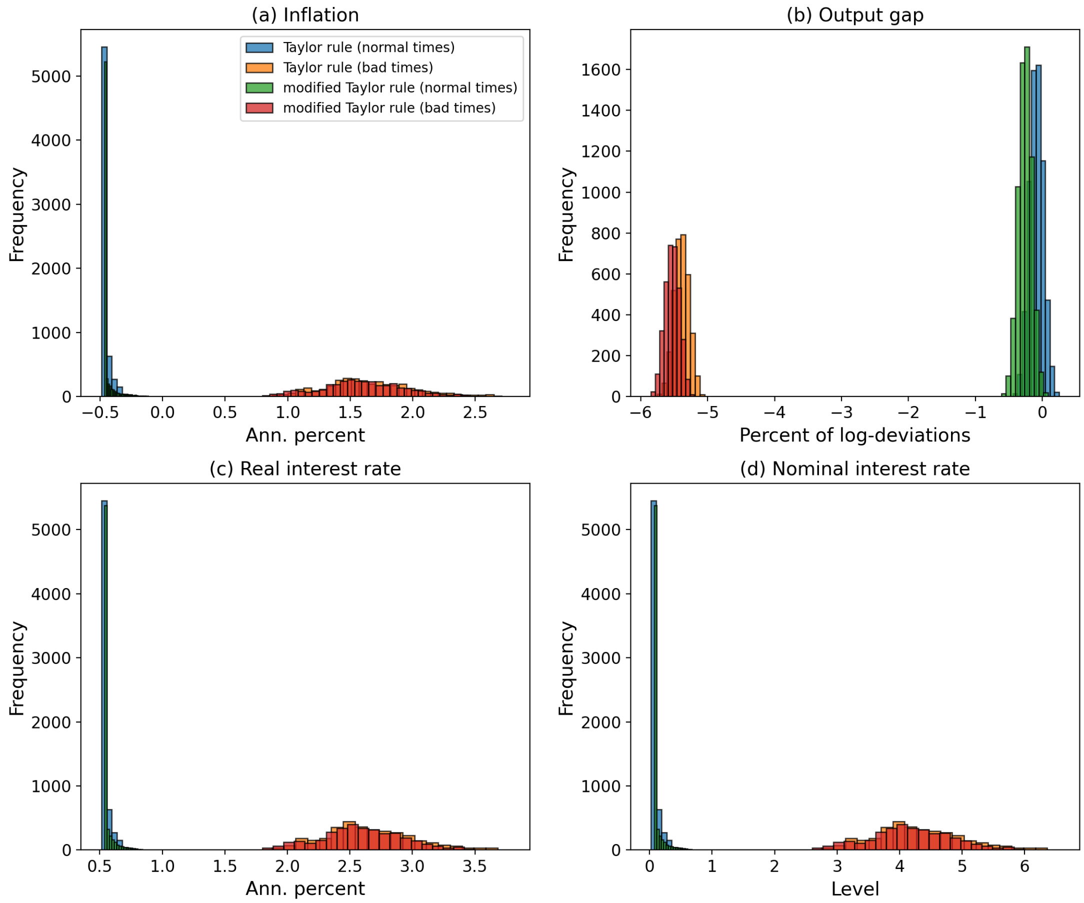

Figure 9: Ergodic distribution: Taylor rules with a ZLB

Note: The figure displays the ergodic distribution in the model under a standard Taylor rule and a modified one with an occasionally binding ZLB. Colors distinguish the two regimes: blue denotes the samples corresponding to the standard Taylor rule in normal times, and orange in bad time. 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.

SSS of the normal-times regime is negative (-0.8%, see Table 2), leading to a significant mass in the negative territory within the ergodic distribution of nominal rates. However, when the ZLB constraint binds, as shown in panel (d) of Figure 9, instead of negative nominal rates, we observe a mass point at zero interest rates. The nominal interest rate in the SSS of the normal-times regime is slightly above zero (0.1%), resulting from both a higher long-term rate (0.6% versus 0.1% without the ZLB) and higher inflation (-0.4% compared to -0.9%).

These results may seem counterintuitive in light of existing research. Several papers have found that the presence of the ZLB introduces a "deflationary bias" in both RANK and HANK models,

	Taylor rule	Mod. Taylor rule	Discretion	Commitment
Inflation				
Normal times	-0.4%	0.0%	0.0%	0.0%
persistent supply shock	1.6%	0.0%	2.6%	0.0%
Output gap				
Normal times	-0.2%	-0.1%	0.0%	-0.1%
persistent supply shock	-5.4%	-5.5%	-5.5%	-5.4%
Real interest rates				
Normal times	0.6%	0.1%	0.1%	0.3%
persistent supply shock	2.6%	2.7%	2.6%	2.3%
Nominal interest rates				
Normal times	0.1%	0.1%	0.1%	0.3%
persistent supply shock	4.2%	2.8%	5.2%	2.3%

Table 4: Stochastic steady-state values, model with ZLB.

decreasing rather than increasing the long-term real interest rate.[21] The intuition is that the central bank cannot provide the necessary degree of monetary policy accommodation during ZLB spells, which are both deflationary and contractionary. Households anticipate this, increasing their demand for safe assets even when the economy is outside the ZLB, depressing the natural rate.

The reasoning changes in the face of persistent supply shocks. During normal times, the nominal rate would be negative in the SSS absent the ZLB. As the nominal rate cannot be negative if the ZLB is active, the central bank should guarantee that SSS inflation is large enough. According to equation (22), SSS inflation is a function of the policy gap between the long-term real rate and the central bank's target of 1%, thus the real rate needs to increase in order to raise inflation. A long-term real rate of 0.6% and an inflation rate of -0.4% are the smallest pair of values that satisfy both the Fisher equation and equation (22).

Modified Taylor rule. We also extend the modified Taylor rule (23) to include a ZLB constraint. Results are displayed in Figure 9 in green and red, and the SSSs in the second column of Table 4. In this case, the presence of the ZLB does not change the SSS significantly. In terms of the ergodic distribution, the presence of the ZLB skews the distribution of inflation, real and nominal rates to the right.

[21] See, for instance, Nakata and Schmidt (2019), Bianchi et al. (2021), or Fernández-Villaverde et al. (2024).

6.2 Optimal monetary policy

We turn next to the design of optimal monetary policy. The ZLB introduces an additional occasionally binding constraint, $i_t \geq 0$, into the central bank's optimization problem. The complete set of first-order conditions can be found in Appendix A.4.

Ergodic distribution and SSS. Figure 15 in Appendix B and the third and fourth columns in Table 4 display the ergodic distribution as well as the SSS values in the case of optimal policy under discretion and commitment. In the case of discretion, results are very similar to those without the ZLB.[22] The main difference is the right skewness in the ergodic distribution of inflation and interest rates.

In the case of commitment, SSS optimal inflation remains at zero in both regimes, despite the presence of the ZLB. Real interest rates, however, are lower in normal times and higher in bad times compared to the case without the ZLB. The reason is, again, precautionary savings. The ZLB affects more during normal times, as the central bank cannot provide the necessary monetary policy stimulus when nominal rates hit the ZLB. Agents now save comparatively more in good time, and less in bad times, which explains the comparative change in rates.

Impulse responses. Finally, we analyze the optimal policy under commitment in response to a regime change when the economy is in the normal-times regime and nominal rates are close to the ZLB constraint. Figure 10 compares the baseline without the ZLB (solid blue line) with the case featuring an occasionally binding ZLB.

Once the regime change occurs, the central bank tolerates a temporary rise in inflation in both cases, as described above. However, the path of nominal rates differs significantly. In the baseline model, nominal rates abruptly fall on impact, as the central bank uses monetary policy to partially cushion the distortions in markups. When the ZLB is present, this abrupt reduction in nominal rates is not feasible as they hit the ZLB. The central bank tightens them at a slower pace compared to the baseline to compensate for the lack of "policy space". Consequently, in the ZLB

[22] The long-run real rate is slightly lower in normal times, 0.07% instead of 0.09%, consistent with the preemptive easing identified by Adam and Billi (2007) and Nakov et al. (2008).

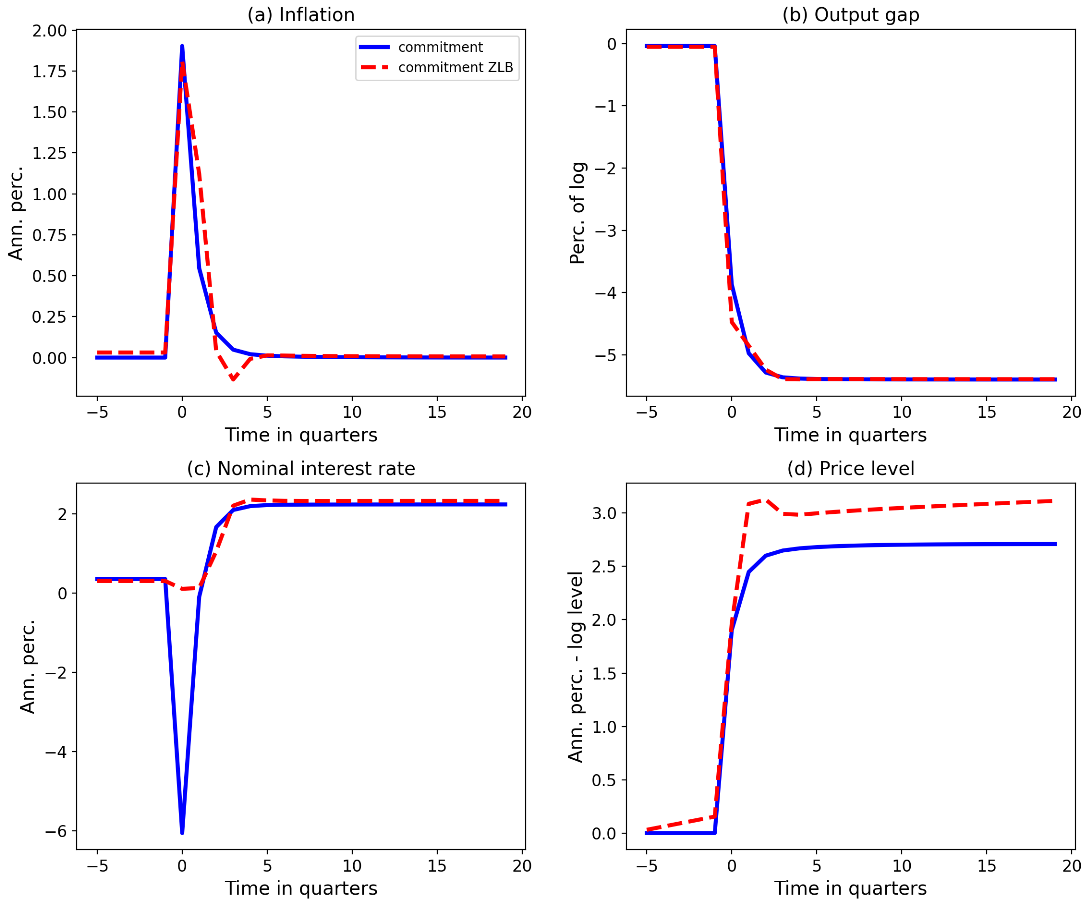

Figure 10: Response to a regime change: commitment with ZLB.
Note: The figure displays the transition from a normal times regime to a bad times regime in the case of the optimal policy under commitment without the ZLB (solid blue line) and with the ZLB (dashed red line). We set the innovations to the temporary shocks at zero. The economy starts at the SSS of the normal times regime in each case.

case, the increase in inflation is larger, leading to a greater deviation in the price level from its pre-shock value. The "bygones are bygones"effect is amplified at the ZLB. Notice that this result shares some of the logic of the "low for longer"optimal policy response to demand shocks in the presence of the ZLB (Eggertsson et al., 2003). The inability of the central bank to reduce nominal rates as much as needed due to the ZLB forces it to tighten nominal rates at a slower pace.

7 Conclusions

This paper advances the understanding of optimal monetary policy in the presence of persistent supply shocks by extending the standard New Keynesian framework. We introduce a novel, regime-switching cost-push shock that captures the economy's transition between "normal times" and "bad times", characterized by sustained increases in production costs.

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, we demonstrate that traditional Taylor rules fail to stabilize inflation across different regimes due to endogenous shifts in the natural interest rate. Second, an inflationary bias emerges during bad times when analyzing optimal policy under discretion because the central bank cannot commit to future policies. Third, under commitment, the central bank cushions the persistent 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 persistent shocks. This result is reinforced if monetary policy is constrained by the zero lower bound.

These findings have significant implications for both academic research and practical policy-making. They suggest that central banks need to consider the potential for persistent supply disruptions and the associated shifts in the natural rate when designing monetary policy frameworks. Traditional rules that do not account for regime changes may lead to systematic deviations from inflation targets and suboptimal economic outcomes.

Methodologically, our paper is, to the best of our knowledge, the first to solve such a New Keynesian model globally with a persistent, 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

Adam, K. and R. M. Billi (2006): "Optimal monetary policy under commitment with a zero bound on nominal interest rates," *Journal of Money, credit and Banking*, 1877–1905.

——— (2007): "Discretionary monetary policy and the zero lower bound on nominal interest rates," *Journal of Monetary Economics*, 54, 728–752.

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.

Andrade, P., J. Gali, H. L. Bihan, and J. Matheron (2019): "The Optimal Inflation Target and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 50, 173–255.

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.

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.

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.

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.

Daudignon, S. and O. Tristani (2023): "Monetary Policy and the Drifting Natural Rate of Interest," Working Paper Series 2788, European Central Bank.

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

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.

Eggertsson, G. B. et al. (2003): "Zero bound on interest rates and optimal monetary policy," *Brookings papers on economic activity*, 2003, 139–233.

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. Lauri'ere, 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.

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.

Nakata, T. and S. Schmidt (2019): "Conservatism and liquidity traps," *Journal of Monetary Economics*, 104, 37–47.

Nakov, A. et al. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4, 73–127.

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.

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 5.

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:

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

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

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

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:

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

The first-order condition with respect to consumption reads:

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

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

$$0=\rho_t\theta\left(\epsilon-1\right)\left(1+\pi_t\right)^{\epsilon-2}+\zeta_t\theta\epsilon\left(1+\pi_t\right)^{\epsilon-1}\Delta_{t-1}. \tag{32}$$

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

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

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

$$\begin{aligned} 0 = & -\frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} + \beta \mathbb{E}_t \left[\frac{\partial V}{\partial \Delta_t}\right] - \zeta_t + \\ & \mu_t \left[p_t^* c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial F}{\partial \Delta_t}\right] - \mathcal{M}\left((c_t+g_t)^{1+\omega} c_t^{\gamma} \frac{\omega}{\Delta_t}\left(\frac{\Delta_t}{A_t}\right)^{\omega} (1+\tau_t)\left(A_t\right)^{-1} + c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_t}\right]\right)\right]. \end{aligned} \tag{33}$$

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 (33):

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

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\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],\\
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},\\
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},\\
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},\\
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(\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-1}}\right]\right)\right],\\
c_t^{-\gamma} =& \lambda_t,\\
h_t^{\omega} =& w_t\lambda_t,\\
\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} + (1-\theta)\left(p_t^*\right)^{1-\epsilon},\\
\Delta_t =& \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}.
\end{aligned}
$$

A.3 Optimal policy under commitment

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

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

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 Ξ_t^N is given by:

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

Finally, the first order condition with respect to Ξ_t^D reads 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}
$$

A.4 ZLB

A.4.1 Discretion

We now consider the case in which the ZLB is a binding constraint. The Lagrangian is

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

where we define

$$
H\left(\Delta_{t-1}, A_t, \tau_t\right) \equiv \beta \left(\frac{1}{1+\pi_t}\right) c_t^{-\gamma}.
$$

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

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

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

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

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

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

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

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

The first order condition with respect to nominal rates

$$0 = \varphi_t E_t \left[H\left(\Delta_t, A_{t+1}, \tau_{t+1}\right)\right] + \chi_t,$$

where χ_t and φ_t are zero when the ZLB is not binding $\left(i_t \geq 0\right)$.

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

$$\begin{aligned} 0 =& -\frac{\left(\frac{(c_t+g_t)\Delta_t}{A_t}\right)^{1+\omega}}{\Delta_t} + \beta \mathbb{E}_t \left[\frac{\partial V}{\partial \Delta_{t-1}}\right] - \zeta_t \\ &+\mu_t \left[p_t^* c_t^\gamma \mathbb{E}_t \left[\frac{\partial F}{\partial \Delta_{t-1}}\right] - \mathcal{M}\left(\left(c_t + g_t\right)^{1+\omega} c_t^\gamma \frac{\omega}{\Delta_t}\left(\frac{\Delta_t}{A_t}\right)^\omega \left(1+\tau_t\right)\left(A_t\right)^{-1} + c_t^\gamma \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_{t-1}}\right]\right)\right] \\ &+\left(1+i_t\right) E_t \left[\frac{\partial H}{\partial \Delta_{t-1}}\right]. \end{aligned} \quad (34)$$

According to the envelope theorem, we have 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 out the following in (34):

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

The full set of equations is 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. \\
& \mathcal{M} \left(\left((1+\omega) c_t + \gamma \left(c_t + g_t\right)\right) \left(c_t + g_t\right)^{\omega} c_t^{\gamma-1} \left(\frac{\Delta_t}{A_t}\right)^{\omega} \left(1+\tau_t\right) \left(A_t\right)^{-1} \right. \\
& \left.\left. + \gamma c_t^{\gamma-1} \mathbb{E}_t \left[G\left(\Delta_t, A_{t+1}, \tau_{t+1}\right)\right]\right)\right], \\
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}, \\
0 =& \mu_t \Xi_t^D + \nu_t (1-\theta) \left(1-\epsilon\right) \left(p_t^*\right)^{-\epsilon} - \zeta_t (1-\theta) \epsilon \left(p_t^*\right)^{-\epsilon-1}, \\
0 =& \varphi_t E_t \left[H\left(\Delta_t, A_{t+1}, \tau_{t+1}\right)\right] + \chi_t, \\
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} \left(1+\tau_t\right) \left(A_t\right)^{-1} + c_t^{\gamma} \mathbb{E}_t \left[\frac{\partial G}{\partial \Delta_{t-1}}\right]\right)\right] \\
& + \left(1+i_t\right) E_t \left[\frac{\partial H}{\partial \Delta_{t-1}}\right] + \varphi_t \gamma c_t^{-\gamma-1}, \\
c_t^{-\gamma} =& \lambda_t, \\
h_t^{\omega} =& w_t \lambda_t, \\
\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} + (1-\theta) \left(p_t^*\right)^{1-\epsilon}, \\
\Delta_t =& \theta \left(1+\pi_t\right)^{\epsilon} \Delta_{t-1} + (1-\theta) \left(p_t^*\right)^{-\epsilon}.
\end{aligned}
$$

A.4.2 Ramsey

The Lagrangian is given by:

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

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

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

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

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

The first-order condition with respect to inflation is:

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

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

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

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

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

The first-order condition with respect to Ξ_t^D is:

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

Finally, the first-order conditions with respect to nominal rates:

$$
0 \quad = \quad \varphi_t \beta E_t \left[\frac{1}{1+\pi_{t+1}} c_{t+1}^{-\gamma}\right] + \chi_t,
$$

where χ_t and φ_t are zero when the ZLB is not binding $(i_t \geq 0)$.

The full set of equations is 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\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} \\
& + \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]\Bigg] \\
& + \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 + \left[\varphi_t - \beta^{-1}\varphi_{t-1}\left(\frac{1+i_{t-1}}{1+\pi_t}\right)\right] c_t^{-\gamma-1}\gamma c_t^{-\gamma-1}, \\
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(c_t+g_t)^{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}, \\
0 =& -\mu_t\Xi_t^D + \nu_t(1-\theta)\left(1-\epsilon\right)\left(p_t^*\right)^{-\epsilon} + \zeta_t(1-\theta)\left(-\epsilon\right)\left(p_t^*\right)^{-\epsilon-1}, \\
0 =& \mu_t\mathcal{M} - \vartheta_t + \beta^{-1}\vartheta_{t-1}\beta\theta c_{t-1}^{\gamma}c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon}, \\
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 - \beta^{t-1}\varphi_{t-1}\frac{1+i_{t-1}}{\left(1+\pi_t\right)^2}c_t^{-\gamma}, \\
0 =& -\mu_t p_t^* - \varrho_t + \beta^{-1}\varrho_{t-1}\beta\theta c_{t-1}^{\gamma}c_t^{-\gamma}\left(1+\pi_t\right)^{\epsilon-1}, \\
0 =& \varphi_t\beta E_t\left[\frac{1}{1+\pi_{t+1}}c_{t+1}^{-\gamma}\right] + \chi_t, \\
0 =& \chi_t i_t, \\
c_t^{-\gamma} =& \lambda_t, \\
h_t^{\omega} =& w_t\lambda_t, \\
\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} + (1-\theta)\left(p_t^*\right)^{1-\epsilon}, \\
\Delta_t =& \theta\left(1+\pi_t\right)^{\epsilon}\Delta_{t-1} + (1-\theta)\left(p_t^*\right)^{-\epsilon}.
\end{aligned}
$$

B Additional figures

In this Appendix, we provide additional illustrations of the paper's main results. Figure 11 depicts a stylized neural network.

Figure 12 reports the dynamics of the price level p_t based on simulating $10,000$ quarters of the model economy when monetary policy is optimal (under commitment). It addresses the concerns expressed in Section 5.2 regarding whether optimal policy with persistent shocks delivers a constant price level in the long run, as is the case in the standard New Keynesian model with AR(1) shocks. The figure shows that this is not the case, and there is a positive trend in prices.

Figure 13 compares a persistent shock with a temporary one calibrated to produce the same inflation on impact. It shows how, for large shocks, the optimal response to cost-push shocks deviates slightly from pure price-level targeting. Still, this deviation is small compared to that introduced by the optimal response to persistent shocks.

Figure 14 shows the differences in the response to a larger persistent shock. Results are qualitatively similar, but the response is quantitatively larger. This confirms that our qualitative results are independent of shock size.

Finally, Figure 15 displays the ergodic distribution of the model with the ZLB when the central bank implements optimal policy under commitment or discretion. It clearly shows how the ZLB creates a significant right skewness in inflation and interest rates.

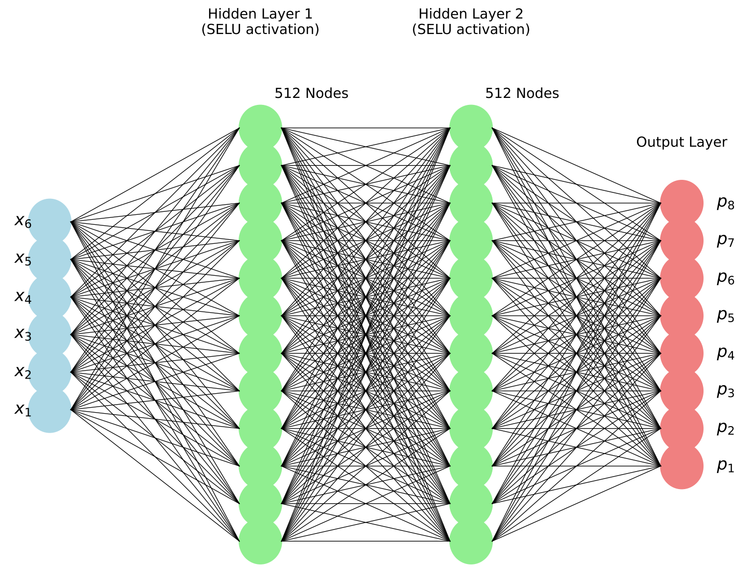

Figure 11: The figure above depicts a stylized neural network with an input $\mathbf{x} \in \mathbb{R}^6$. It consists of two hidden layers, each containing 512 neurons, using SELU activation functions, and produces an output $\mathbf{p}(\mathbf{x}) \in \mathbb{R}^8$.

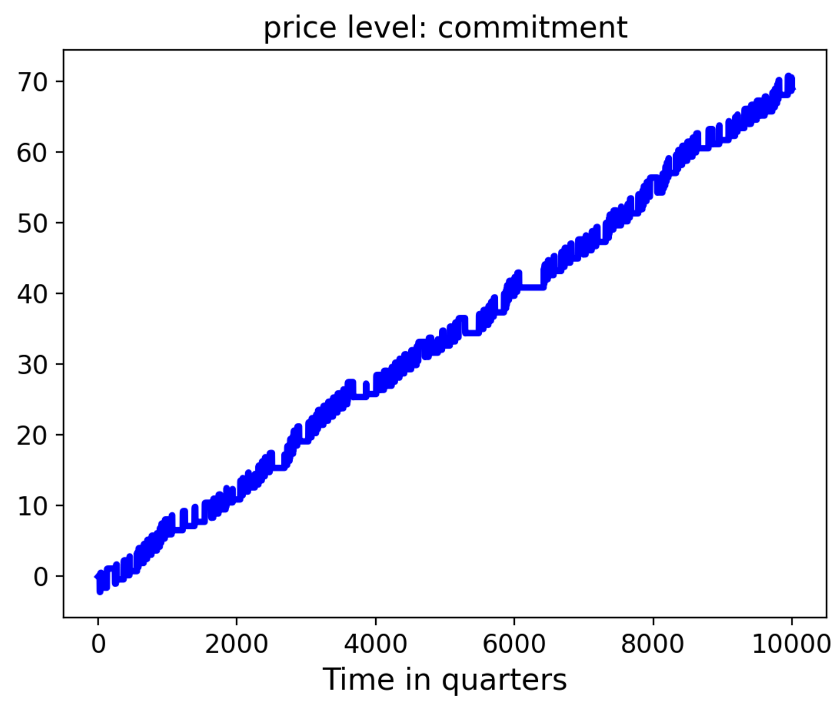

Figure 12: Price level p_t dynamics under commitment.

Note: The figure shows the result of simulating $10,000$ quarters of the model economy when monetary policy is optimal (under commitment).

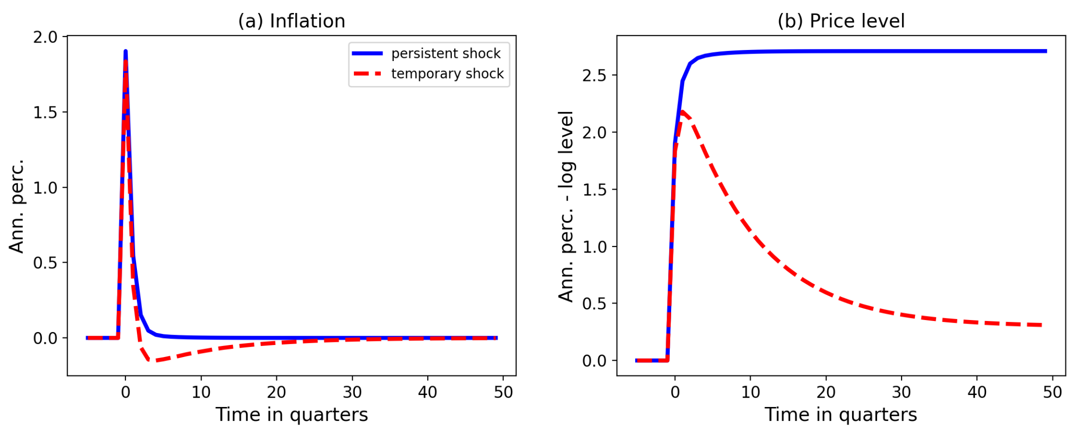

Figure 13: Comparison between persistent and temporary shocks.

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.

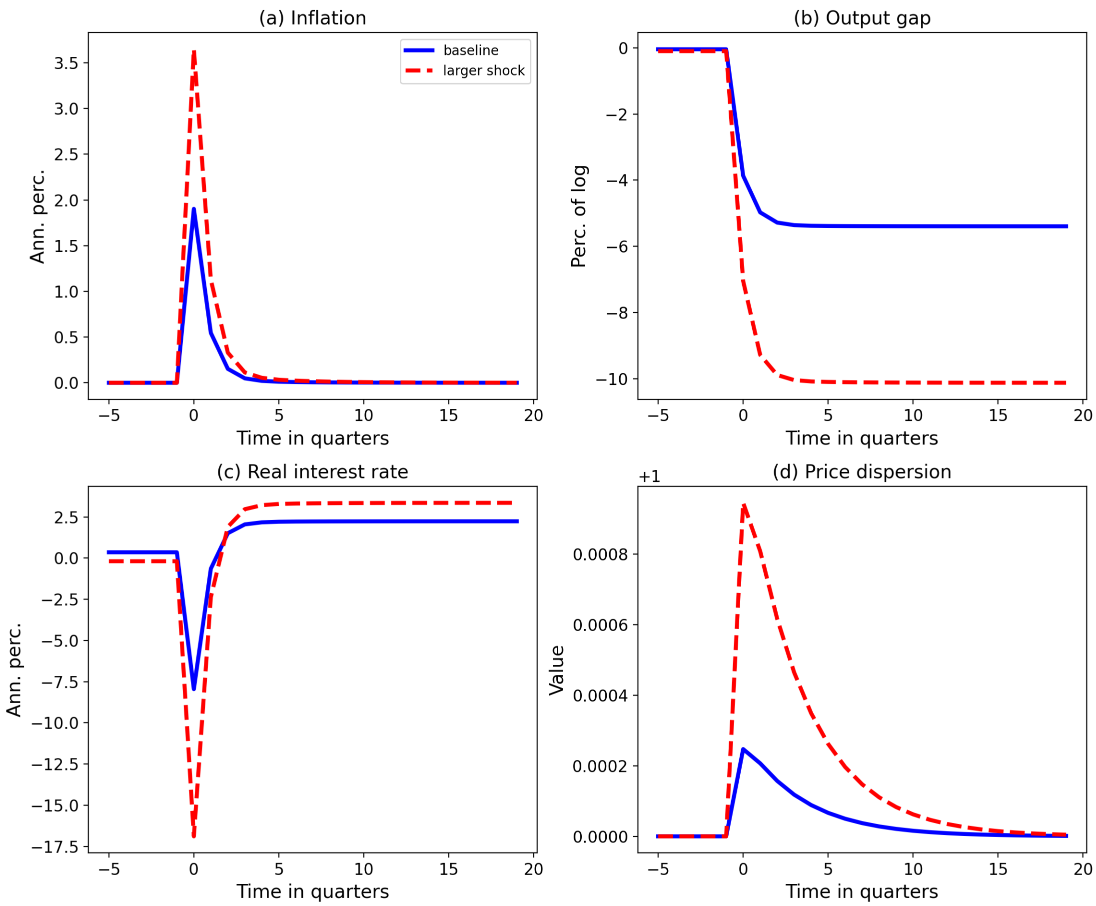

Figure 14: Response to a regime change: shock size.

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.

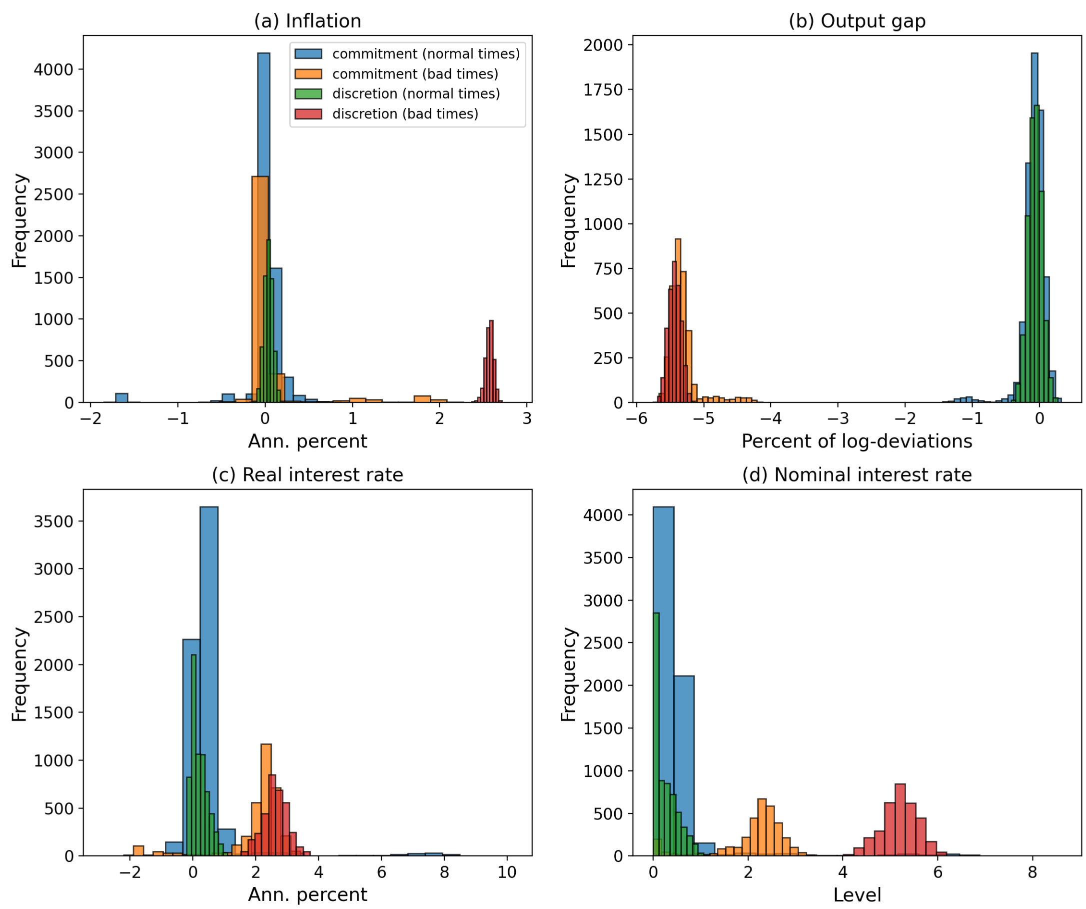

Figure 15: Ergodic distribution: optimal policy at the ZLB.

Note: The figure displays the ergodic distribution in the model under optimal policy with an occasionally binding ZLB. Colors distinguish the two regimes: blue denotes the samples corresponding to commitment in normal times, and orange to bad times. Green is the optimal discretionary policy in normal times and red in bad times. The figure is produced by simulating the model for a large number of periods.