DEUTSCHE
BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 26/2026

The consequences of optimism and pessimism for inflation dynamics and monetary policy

Stefan Hasenclever
(Deutsche Bundesbank and University of Cologne)

Mathias Hoffmann
(Deutsche Bundesbank)

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank, Press and Public Relations Division, at the above address or via email: www.bundesbank.de/contact

Internet http://www.bundesbank.de

DOI https://doi.org/10.71734/DP-2026-26
ISBN 978-3-98848-085-9
ISSN 2941-7503

The Consequences of Optimism and Pessimism for Inflation Dynamics and Monetary Policy*

Stefan Hasenclever[†] Mathias Hoffmann[‡]

September 2, 2026

Abstract

We study how imperfect information about the persistence of economic fundamentals affects inflation dynamics and monetary policy. We focus on total factor productivity (TFP) growth, which encompasses demand- and supply-side effects linked to the post-COVID inflation surge. Agents in our model conduct Bayesian learning to infer the persistence of observed TFP growth rates. With imperfect information, waves of optimism and pessimism emerge, leading to a temporary divergence of inflation dynamics from fundamentals in terms of over- and undershooting. This mechanism helps explain the gradual rise of high inflation rates and the cautious monetary policy response in the aftermath of the COVID-19 crisis. Regarding optimal monetary policy, we find that a more hawkish stance is superior at reducing welfare losses and uncertainty arising in the imperfect-information model.

JEL Codes: E31, E52, D83

Keywords: Monetary Policy, Inflation Dynamics, Imperfect Information, Bayesian Learning

*We are grateful to our referee Fabian Seyrich for the valuable comments that substantially improved the paper. We also thank Francesco Bianchi, Roger Farmer, Benedikt Kolb, Moritz Lenel, Johannes Pfeifer, Frank Schorfheide, and seminar participants at Deutsche Bundesbank for their helpful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect the views of Deutsche Bundesbank.

[†]Deutsche Bundesbank and University of Cologne; stefan.hasenclever@bundesbank.de

[‡]Deutsche Bundesbank; mathias.hoffmann@bundesbank.de

1 Introduction

Following the COVID-19 pandemic and the Russian invasion of Ukraine, global inflation rates have risen sharply, reaching historically high levels. In the policy debate, the surge of inflation has been attributed to demand- and supply-side factors, with disagreements about its persistence. Two opposing camps have emerged: 'Team Transitory' has been assuming a transitory rise in inflation and has been arguing for moderate interest rate hikes, while 'Team Persistent' has been calling for decisive monetary policy responses in the face of expected prolonged inflation rates.[1] On that account, policymakers have had to act under significant uncertainty regarding the duration and sustained level of inflation rates (Londono et al., 2023). In light of this uncertainty, the Federal Reserve (Fed) misjudged the rise in inflation rates as more likely to be temporary. As a result, the Fed delayed raising interest rates until March 2022, even though inflation rates had been above the target of 2% for three consecutive quarters, starting from the second quarter of 2021, as highlighted in Figure 1.[2] In what followed, the Fed raised the Federal Funds Rate over seven consecutive quarters starting in the first quarter of 2022 until the third quarter of 2023, reaching a range between 5.25% and 5.50%.

To account for the uncertainty about the duration of elevated inflation rates, we analyze the implications for inflation dynamics and monetary policy when the economy has imperfect information about the persistence of the fundamentals driving the economy. For this purpose, we depart from the standard macroeconomic model, in which agents have perfect information about the persistence of realized events. To this end, we extend the work by Bianchi and Melosi (2016) and incorporate their learning mechanism about the persistence of TFP growth into a New Keynesian Model. This allows us to assess how the perception of economic fundamentals affects the economy's inflation dynamics and monetary policy responses. In particular, we focus on the role of TFP growth in shaping economic fundamentals, as it includes demand- and supply-side effects, which have been linked to the surge in inflation in the aftermath of the COVID-19 pandemic (Shapiro, 2026; Reis, 2026).

In our model, TFP growth can be realized as either high or low in each period. These TFP growth realizations are either transitory or persistent, with these regimes following a Markov-switching (MS) process. Agents observe realized TFP growth but do not know whether it is transitory or persistent. They form beliefs using Bayesian learning, updating probabilities over regimes as new observations arrive. By assigning probabilities to regimes,

[1] For a non-exhaustive overview of the debate, see Summers (2021), Blanchard (2021), Powell (2021), Beckmann et al. (2021), Krugman (2021), Schneeweiss (2021), and Kose et al. (2021).

[2] Fed Chairman Jerome Powell claimed in a speech in August 2021 that "Inflation at these levels is, of course a cause for concern. But that concern [...] [is] likely to prove temporary" (see: https://www.federalreserve.gov/newsevents/speech/powell20210827a.htm).

Figure 1: US CPI inflation vs. Federal Funds Rate

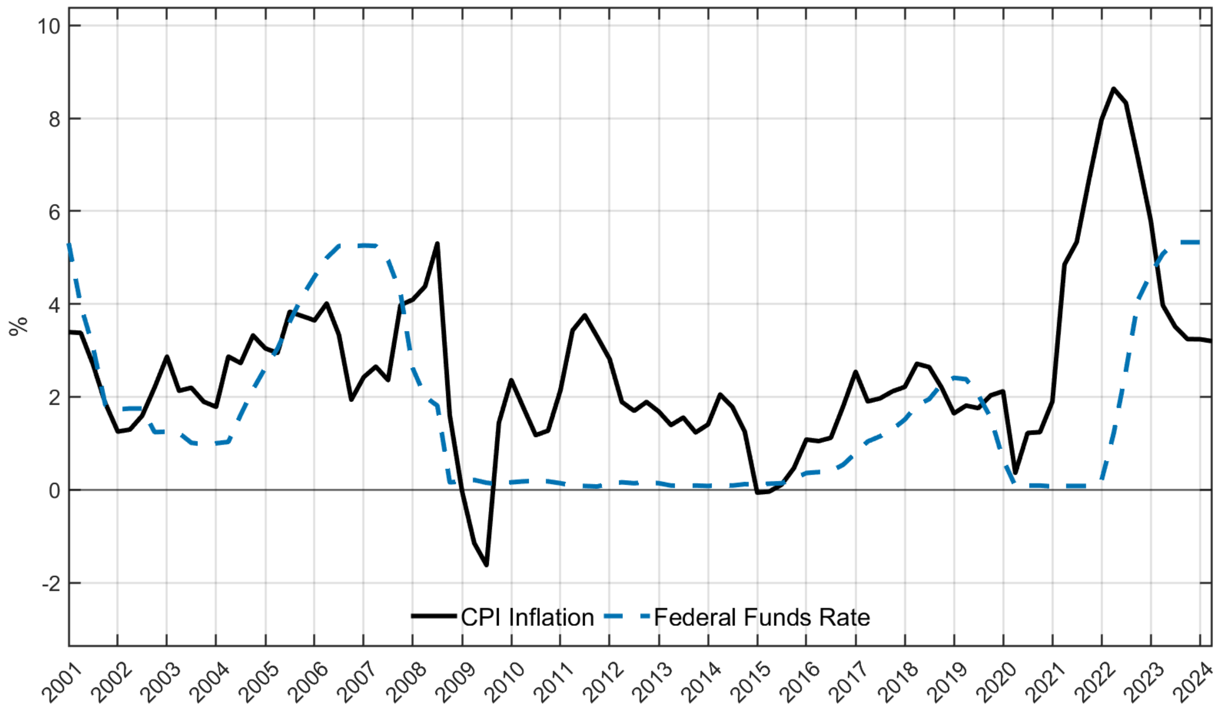

Notes: Comparison of quarterly CPI inflation rate (green line) and the Federal Funds Rate (blue line) from 2001Q1 until 2024Q2. Source: FRED St. Louis Fed.

agents form expectations about future growth, internalizing that these beliefs will change with future observations. In this sense, agents are *fully rational* as they know that they do not know. In the period when the economy shifts between the high- and low-growth states, agents apply prior beliefs about the regime for that period. The prior is designed to be dynamic, meaning that prior beliefs change over time depending on past beliefs and the history of events.

We compare the economic implications of our mechanism with the standard perfect information assumption, which is incorporated in our benchmark model. In the benchmark model, agents know the current regime and, thus, can form beliefs perfectly about future realizations according to the transition matrix. In this sense, the perfect-information model reflects economic dynamics fully aligned with the fundamentals. Throughout our analysis, we calibrate the transition matrix according to the persistence of TFP growth observed in the data from the 1970s onward. For this purpose, we use the data provided by Fernald (2012).

To compare both models, we generate a sequence of TFP growth regimes that replicate the business cycle history since the beginning of the 2000s. Here, we introduce two persistent high-growth phases reflecting the boom before the Global Financial Crisis (GFC) and the current post-COVID episode of persistently high inflation rates. In addition, we add two persistent low-growth phases, representing the GFC and the COVID-19 crisis, respectively. Between the GFC and the COVID-19 crisis, the economy fluctuates between transitory high- and low-growth regimes.

First, we find that waves of optimism and pessimism occur in the economy with imperfect information, leading to economic overheating and undershooting, respectively. In the persistent growth regimes, agents incorrectly assess the probability of being in the corresponding transitory regimes. This results in pessimism in the persistent high-growth regime and optimism in the persistent low-growth regime compared to the benchmark model. These divergences in beliefs have economic consequences as output and inflation undershoot during episodes of persistent high growth, for example, causing the policymaker to increase interest rates less aggressively than fundamentals would suggest. As agents become more certain that they are in the persistent growth regime, the longer the economy remains in a growth environment, output, inflation, and the interest rate gradually converge toward their levels implied by the fundamentals under perfect information.

This mechanism helps account for the hesitant post-COVID monetary policy response observed in the data. In the simulated post-COVID inflation phase, the interest rate rises for six consecutive quarters before converging to the level implied by fundamentals. These policy responses closely match the observed tightening cycle, as the Fed raised the policy rate for seven consecutive quarters in the aftermath of the COVID-19 crisis, as discussed above. In this regard, our learning mechanism brings the standard New Keynesian model closer to the sluggish adjustment of beliefs and macroeconomic variables observed in the data, thereby introducing an endogenous source of persistence. This is achieved without the need to include ad hoc adjustments such as price indexation or habit formation, in line with Milani (2007) and Fuhrer (2017).

Second, our model provides an explanation for why the Fed acted decisively at the onset of the COVID-19 crisis, yet responded cautiously at the beginning of the post-COVID inflation era. We show that after having been in the GFC for a long time and after some subsequent shifts between transitory high- and low-growth regimes, agents assign higher prior beliefs to being in the persistent low-growth regime when entering the COVID-19 crisis compared to the start of the GFC. When considering the Zero Lower Bound (ZLB) on interest rates, the more pessimistic outlook implies that the ZLB is reached sooner in the COVID-19 crisis compared to the GFC. Analogously, after having been confronted with two negative events in a short period, agents are more certain that the recovery following the COVID-19 crisis will be transitory compared to their beliefs at the beginning of the pre-GFC boom. This dynamic adjustment of prior beliefs illustrates how historical experience shapes expectations, providing an explanation for why the Team Transitory interpreted the surge in inflation as temporary and why the Fed adopted a cautious approach to tightening monetary policy after the COVID-19 crisis. These findings are in line with evidence that the state of the economy matters for agents' learning, e.g., Weber et al. (2025).

Third, we show that during transitory episodes, the learning mechanism results in an opposite evolution of beliefs compared to that observed in persistent regimes. For instance, in transitory high-growth phases, imperfectly informed agents are too optimistic as they incorrectly consider the probability of being in the persistent high-growth regime. This leads to an amplified economic expansion, with higher inflation rates compared to the perfect-information model. As a consequence, the central bank raises the interest rate more sharply than justified by fundamentals. In contrast, at the beginning of transitory low-growth episodes, agents are too pessimistic about future outcomes as they incorrectly consider the probability of being in the persistent low-growth regime, leading to an undershooting of inflation and a lower interest rate compared to the benchmark model.

Fourth, we study the implications of a looking-through policy in the aftermath of the COVID-19 crisis, following the evidence of Nakamura et al. (2025). Under this policy, the central bank delays lifting the interest rate from the ZLB until there is sufficient evidence that the recovery and associated inflation dynamics are persistent, as in Buss and Traficante (2025). As discussed above, it took three quarters for the Fed to lift the policy rate from the ZLB despite inflation exceeding the 2% target. According to transcripts, during this time, the Fed learned that the inflation-driving factors were persistent.[3] To replicate these three prolonged periods at the ZLB in our model framework, we assume that the central bank must be almost 90% certain that the recovery is persistent before raising the interest rate. We show that the macroeconomic variables' responses are amplified by the presence of the ZLB, leading to an excessive increase in inflation and output compared to the model without this constraint. This highlights the crucial importance of the timing of a central bank's exit from the ZLB for keeping inflation under control and underscores the risks of delaying policy tightening if inflationary pressures prove to be persistent rather than temporary.

Fifth, and finally, we explore the evolution of welfare and uncertainty arising from the imperfect-information model and examine how monetary policy should optimally respond to reduce the associated welfare losses, using the methods developed by Bianchi (2016). We show that uncertainty about output, inflation, and the interest rate, as well as welfare losses, increases with the duration of persistent growth regimes. This is because, as agents become more confident that the regime is persistent, potential adjustments become large when the economy shifts to the other growth realization.

[3]According to FOMC statements, it took the Federal Reserve around nine months to determine that inflation dynamics were persistent. In April 2021, the statement referred to inflation as being driven by "transitory factors"—a phrase that was dropped in December 2021, as noted by Ihrig and Waller (2024).

Given these patterns of welfare and uncertainty, we determine the optimal monetary policy response by minimizing welfare losses over a parameter space of Taylor rule coefficients. We find that the most hawkish monetary policy stance minimizes both welfare losses and uncertainty most effectively. This result is in line with the findings of Nakamura et al. (2025). Notably, the optimal policy does not require more aggressive adjustments of the nominal interest rate than those implied by the standard Taylor rule parameters. This is because the commitment to a more hawkish stance better anchors inflation expectations, allowing the real interest rate to adjust more effectively to smooth business cycles. Overall, the more hawkish monetary policy stance increases welfare nearly ninefold relative to the standard Taylor rule parameters.

Taking together our looking-through policy and welfare analysis suggests that policymakers should pursue a decisive monetary policy stance in periods when the persistence of fundamentals is uncertain, as advocated by Team Persistent.

Related Literature: Our research adds to the literature debating the causes and consequences of the recent spike in inflation. Bernanke and Blanchard (2025) argue that the rise in inflation was mainly driven by shocks to prices and sectoral supply shortages. Dao et al. (2024) also emphasize the role of energy price shocks across advanced and emerging economies, with the notable exception of the United States (US), where demand-side factors also played a key role. Moreover, Amiti et al. (2024) highlight supply chain disruptions, the shift of consumption from services towards goods, and labor market constraints. Relatedly, Merendino and Monacelli (2026) emphasize the role of transportation uncertainty in amplifying supply chain disruptions. Ball et al. (2022) also present evidence that labor market tightness contributed to core inflation, while Lorenzoni and Werning (2023) emphasize the relevance of price-wage spirals. Other contributions rely on non-linear Phillips curves, as in Harding et al. (2023) and Benigno and Eggertsson (2023), or on rational inattention, as in Pfauti (2023), to account for the rise in inflation. Bergholt et al. (2025), Bianchi and Melosi (2022) and Bianchi et al. (2023) instead focus on demand-side effects, particularly fiscal shocks.

In addition to these contributions, we propose a mechanism that explains the hesitant monetary policy response to the post-COVID surge in inflation. Closely related to our work, Erceg et al. (2024) combines a non-linear Phillips curve, learning about the persistence of a cost-push shock, endogenous inflation indexation, and a central bank's forecast-based targeting rule. Together, these features can account for the delayed monetary policy response and the resilient economic performance observed after the COVID-19 pandemic, despite the occurrence of supply shocks. To learn about the persistence of the cost-push shock, agents solve a signal extraction problem. This is similar to Buss and Traficante (2025), who also model a signal extraction problem to replicate the tentative post-COVID monetary policy response, assuming the central bank refrains from lifting rates from the ZLB if it assesses the cost-push shock to be temporary. In contrast to these studies, we consider TFP growth shocks, capturing both supply- and demand-side effects identified as key drivers of the post-COVID inflation. Furthermore, in our approach, prior beliefs do not have to be set arbitrarily but instead adjust dynamically based on past observations. In turn, our model presents *one explanation* of why central banks were less willing to consider the post-COVID inflation as persistent after having experienced the GFC and the COVID-19 crisis recently. Moreover, unlike learning by solving a signal extraction problem, our mechanism ensures rationality as agents take into account that their beliefs will change according to future states of the economy. This feature also distinguishes our work from Williams and Sargent (2026), who also study the Fed's slow post-COVID response. In their paper, the Fed solves an optimal control problem each period, updating its beliefs about inflation persistence and treating the updated estimate as fixed going forward, consistent with anticipated utility.

In the context of optimal monetary policy in the aftermath of the COVID-19 crisis, our paper is related to the work of Brandão-Marques et al. (2024), who demonstrate that during periods of uncertainty about inflation persistence, the central bank should respond aggressively to inflation spikes and maintain a tightening bias to ensure that inflation returns to target. In contrast to our approach, however, their setup assumes that only the central bank lacks information about the inflation persistence, while the private sector is fully informed. Moreover, their framework does not rely on a learning mechanism—rather, the central bank's beliefs are set ad hoc, as is the timing at which the truth about the underlying persistence is revealed. Furthermore, the associated policy changes are not anticipated by households and firms. This contrasts with our setting, in which agents and the central bank rationally take into account that the states of the economy will change, and accordingly, so will their beliefs. In addition to our explanation-based mechanism, Nakamura et al. (2025) identify three reasons why optimal monetary policy may follow a looking-through approach under well-anchored long-run inflation expectations: forward guidance, the interaction of supply and demand shocks, and transmission lags during transitory shocks.

Outline. The rest of the paper is structured as follows: Section 2 introduces the New Keynesian model, the learning mechanism, and the solution techniques. Section 3 describes the calibration of the model. Subsequently, Section 4 provides a discussion of the results, before Section 5 finally summarizes the findings of the paper.

2 New Keynesian Model with incomplete information

We incorporate the learning mechanism developed by Bianchi and Melosi (2016) into a standard New Keynesian model. The model features nominal rigidities, as we assume monopolistic competition in the production sector and that prices are set according to staggered nominal contracts à la Calvo (1983). Monetary policy follows either a Taylor rule or is constrained by the ZLB, which is modeled as a MS process.

2.1 Households

The economy is populated by a continuum of households. The typical household derives utility from consumption C_t and disutility from supplying labor N_t

$$U(C_t, N_t) = \ln(C_t) - \frac{N_t^{1+\phi}}{1+\phi}, \tag{1}$$

where ϕ denotes the inverse Frisch elasticity.

The typical household receives a nominal wage W_t for supplying labor and spends her income on consumption expenditures. The price for consumption goods is denoted by P_t. The household also trades non-contingent bonds B_t that pay a nominal gross interest rate R_t. Altogether, the household's budget constraint reads:

$$P_t C_t + B_t = B_{t-1} R_{t-1} + W_t N_t. \tag{2}$$

Maximizing the utility function (1) subject to the budget constraint (2) yields the labor-leisure equation and the Euler equation

$$\frac{W_t}{P_t} = N_t^{\phi} Y_t, \tag{3}$$

$$Y_t^{-1} = \beta \, \mathbb{E}_t \, Y_{t+1}^{-1} \frac{R_t}{\Pi_{t+1}}, \tag{4}$$

where $\Pi_{t+1} = \frac{P_{t+1}}{P_t}$ reflects the gross inflation rate, while β and ϕ denote the discount factor and the inverse Frisch elasticity, respectively. Since market clearing implies that output is entirely spent on consumption $Y_t = C_t$, we substitute output for consumption.

2.2 Firms

Final goods production. Final goods producers aggregate a continuum of intermediate goods $Y_t(i)$ to produce the final output good

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}.$$

Here, ϵ denotes the price elasticity of demand for intermediate goods.

We assume that final goods production is perfectly competitive and that final goods producers sell the final good at the aggregate price index P_t

$$P_t = \left(\int_0^1 P_t(i)^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}.$$

Intermediate goods producers. Intermediate goods are produced by a continuum of monopolistic firms, which set a markup $\mu = \frac{\epsilon}{\epsilon-1}$ over their nominal marginal costs. To produce the intermediate goods, firms utilize labor and TFP Z_t according to the linear production function

$$Y_t(i) = Z_t N_t(i). \tag{5}$$

The demand for intermediate goods depends positively on aggregate demand Y_t and negatively on its prices $P_t(i)$

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t. \tag{6}$$

Following Calvo (1983), we assume that each firm can reset its price with a probability of $1-\theta$ in any given period. If a firm has the opportunity to adjust its price, it sets the price optimally $P_t^*(i)$ to maximize its expected profits $E_t \sum_{k=0}^{\infty} \theta^k Q_{t,t+k} \{P_t^*(i) Y_{t+k}(i) - W_{t+k} N_{t+k}(i)\}$ subject to the production function (5) and the demand function (6).[4]

2.3 The total factor productivity process

We assume that total factor productivity evolves according to an exogenous process

$$\ln Z_t = \ln Z_{t-1} + \mu(\xi_t^Z), \tag{7}$$

where the drift μ depends on a discrete Markov process ξ_t^Z. We assume ξ_t^Z can take four regimes, $\xi_t^Z \in \{1,2,3,4\}$, which evolve according to the transition probability matrix

[4] $Q_{t,t+k}$ denotes the stochastic discount factor (for details see Appendix B).

$$P = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \\ p_{41} & p_{42} & p_{43} & p_{44} \end{bmatrix}, \tag{8}$$

where p_{ij} denotes the probability of switching from regime i to regime j.

The regimes map into the drift of TFP such that $\xi_t^Z \in \{1,2\}$ result in a high TFP drift, that is $\mu(\xi_t^Z) = \mu_h$, while $\xi_t^Z \in \{3,4\}$ result in a low TFP drift, that is $\mu(\xi_t^Z) = \mu_l$.

We define the high TFP realizations of Regime 1 and Regime 2 as Block 1, while Block 2 comprises the low TFP realizations from Regime 3 and Regime 4. Within each of the blocks, the growth realizations do not differ in magnitude, but in their persistence, such that $p_{22} > p_{11}$ and $p_{44} > p_{33}$. Therefore, we define Regime 1 and Regime 3 as transitory growth regimes and Regime 2 and Regime 4 as persistent growth regimes. In addition, we assume that regimes do not switch within the blocks, that is, $p_{12} = p_{21} = p_{34} = p_{43} = 0$. To sum up the notation, it holds that

$$\left.\begin{array}{l}\text{Regime 1: transitory high-growth regime} \\ \text{Regime 2: persistent high-growth regime}\end{array}\right\}\text{Block 1}$$

$$\left.\begin{array}{l}\text{Regime 3: transitory low-growth regime} \\ \text{Regime 4: persistent low-growth regime}\end{array}\right\}\text{Block 2}$$

We set the entries of P to match the model to the data. Therefore, we use the quarterly TFP growth rates from 1970Q1 until 2022Q4 provided by Fernald (2012).[5] In the first step, we calculate the mean TFP growth rate and define all growth rates above and below the mean as high-growth and low-growth realizations, respectively. Moreover, we define persistent regimes as those occurring for at least three consecutive periods of the same growth realization, while transitory regimes represent episodes of at most two consecutive realizations of the same growth rate. This procedure pins down the sequences of Regime 1, Regime 2, Regime 3, and Regime 4 in the data.

As described above, the entries of P are the probabilities of shifting from Regime i to Regime j, that is p_{ij} with $i \in \{1,2,3,4\}$ and $j \in \{1,2,3,4\}$. To calculate p_{ij} we divide the number of shifts from Regime i to Regime j, $\sum s_{ij}$, by the total amount of transitions from Regime i ($\sum_{j=1}^{4} s_{ij}$) observed in the data. For example, the probability to shift from Regime 1 to Regime 3 is given by $\frac{\sum s_{13}}{\sum_{j=1}^{4} s_{1j}}$.

[5]The choice of the sample is motivated by the results of Kiley (2024), who show that a Bayesian decision maker can better predict the possible persistence of the high inflation rates in 2021 when taking data from the second half of the 20th century into account. In contrast, 21st-century data alone contain little information on which the Bayesian decision maker might consider inflation to be persistent.

Following these two steps results in the transition matrix

$$P = \begin{bmatrix} 0.21 & 0 & 0.56 & 0.23 \\ 0 & 0.77 & 0.19 & 0.04 \\ 0.58 & 0.20 & 0.22 & 0 \\ 0.11 & 0.06 & 0 & 0.83 \end{bmatrix}. \tag{9}$$

The transition probabilities of the transitory regimes are quite similar, with a likelihood of around 20%. For both transitory regimes, the most likely transition is to the other transitory regime, at around 60%. The transition probabilities of the persistent regimes differ mainly in the probability of remaining in their respective regime. This probability is higher in the persistent low-growth regime than in the persistent high-growth regime.

Moreover, we calculate the high TFP growth rate μ_h and the low TFP growth rate μ_l as the average growth rate above and below the mean in the data ($\bar{\mu}^{Data} = 0.84\%$), yielding $\mu_h = 3.4\%$ and $\mu_l = -1.5\%$. With these values, we can compute the ergodic TFP growth rate, $\bar{\mu} = (\bar{p}_1 + \bar{p}_2)\mu_h + (\bar{p}_3 + \bar{p}_4)\mu_l$, where $\bar{p}_i$ denotes the ergodic probability of being in regime i. We obtain $\bar{\mu} = 0.80\%$, which is close to the average growth rate in the data.

Figure 2: TFP growth: Simulation vs. data

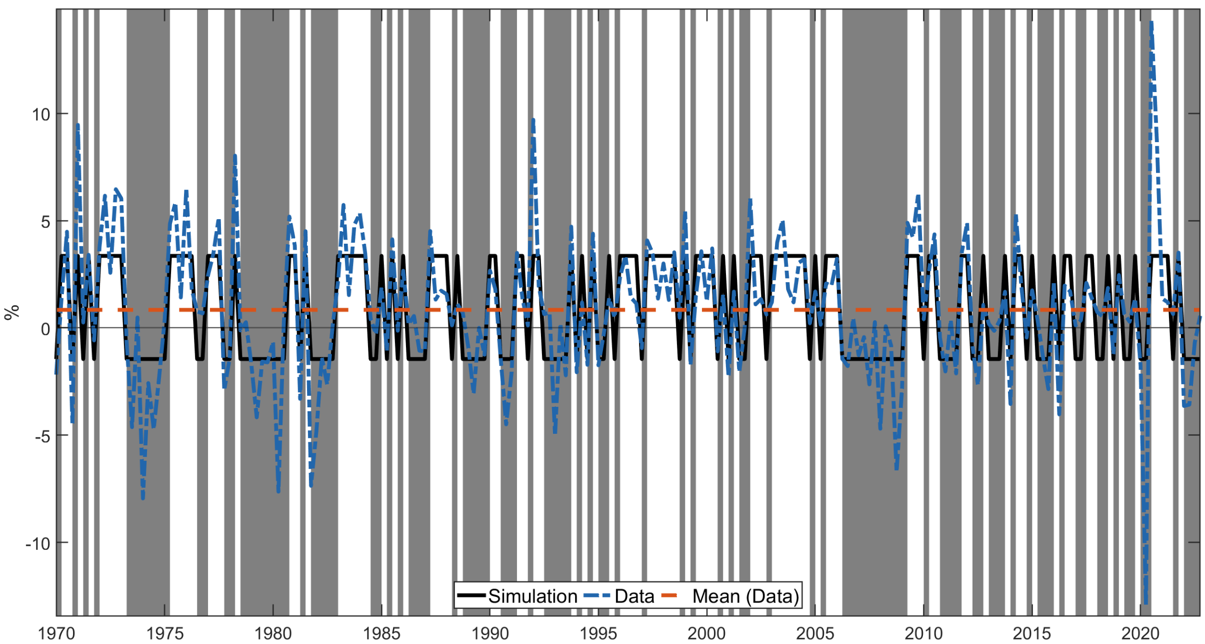

Notes: Comparison of TFP growth in the data (blue-dashed line) and simulated TFP growth (black-solid line) around the TFP growth mean in the data (orange-dashed line). Source: Fernald (2012).

Figure 2 highlights how our calibration matches the TFP growth in the data by plotting the realized TFP growth rate (blue-dashed line) against our simulation (black-solid line) based on our definition of long- and transitory regimes described above. The white and gray areas indicate high- and low-growth episodes, respectively, while the red-dashed line displays the mean of TFP growth in the data. Our two-state approximation of values below and above the mean captures the evolution of TFP growth well, except for large outliers such as the sharp drop and recovery during the COVID-19 crisis in 2020.

2.4 Incomplete information about the persistence of TFP growth

We consider two model specifications with respect to agents' knowledge of TFP persistence. In our benchmark model, perfectly informed agents can observe the history of regimes and the transition matrix P, which allows them to assign a probability to the persistence of the realized TFP at any point in time. Under perfect information, the following model featuring the MS process described above can be solved using existing methods in the literature.[6]

In the alternative specification, we assume that agents have imperfect information. They know the transition matrix P and the history of the blocks, that is, whether TFP growth is high or low, but cannot observe the history of regimes. In other words, agents observe TFP growth but do not know its persistence, whether the current realization is transitory or persistent. This means imperfectly informed agents cannot predict the most likely future TFP realizations. To solve the model, we use the fact that agents can observe TFP growth and know the transition matrix P following the two-step procedure of Bianchi and Melosi (2016). In the first step, we calculate the evolution of agents' beliefs about the current regime conditional on the history of the blocks by applying Bayesian learning. Based on these beliefs, we calculate the transition probabilities across blocks in the second step in order to recast the model as a perfect information rational expectations model.

2.5 Learning about the persistence of TFP growth

As we show in detail in Subsection 2.5.1, in the incomplete information model, agents apply Bayes' Theorem to form beliefs about the regimes conditional on having been in a specific block for consecutive periods of time. When the block shifts, agents must rely on prior beliefs about the prevailing regime. The prior is designed to be dynamic as it ensures that past beliefs affect prior beliefs as discussed in Subsection (2.5.2).

[6] For example, Schorfheide (2005), Liu et al. (2011), Davig and Leeper (2007), Farmer et al. (2009), Cho (2016), and Foerster et al. (2016).

2.5.1 Beliefs within the blocks

As highlighted above, agents can observe the history of blocks. In this sense, they know the number of consecutive periods spent in the current block, denoted as τ_t^i, with $i \in \{1, 2\}$ indexing the two blocks. For example, if the economy is in Block 1 at time t, $\tau_t^1 > 0$ and $\tau_t^2 = 0$. In the next period, the economy can either remain in Block 1, implying that $\tau_{t+1}^1 = \tau_t^1 + 1$ and $\tau_{t+1}^2 = 0$, or it can switch to Block 2, in which case $\tau_{t+1}^1 = 0$ and $\tau_{t+1}^2 = 1$.

Given this information, they can apply Bayes' Theorem to compute the conditional probability of being in the transitory regime given the consecutive realizations of Block 1 ($\tau_t^1 > 1$). As shown by Bianchi and Melosi (2016), the conditional probability using Bayes' Theorem can be expressed as

$$prob(\xi_t^Z = 1|\tau_t^1) = \frac{prob(\xi_{t-1}^Z = 1|\tau_{t-1}^1)p_{11}}{prob(\xi_{t-1}^Z = 1|\tau_{t-1}^1)(p_{11} - p_{22}) + p_{22}}. \tag{10}$$

Consequently, the probability that the economy is in the first regime of Block 1 depends on the previous period's probability and the entries of P.

Similarly, the belief that the economy is in the transitory low-growth regime ($\xi_t^Z = 3$) given the consecutive realizations of Block 2 ($\tau_t^2 > 1$) can be calculated as

$$prob(\xi_t^Z = 3|\tau_t^2) = \frac{prob(\xi_{t-1}^Z = 3|\tau_{t-1}^2)p_{33}}{prob(\xi_{t-1}^Z = 3|\tau_{t-1}^2)(p_{33} - p_{44}) + p_{44}}. \tag{11}$$

Equations (10) and (11) determine agents' beliefs about the regimes within the blocks. As mentioned before, the aim is to recast the imperfect-information model as a perfect-information model by expanding the number of regimes indexed by agents' beliefs. For this purpose, the beliefs have to converge within the blocks to ensure that the expanded set of regimes is tractable. As shown in Appendix A.1, three assumptions about the entries of P ensure the convergence of beliefs within the blocks. This implies that after a sufficient number of consecutive periods within each block (τ_1^*, τ_2^*), agents become almost certain to be in the persistent regime, with the probability of the transitory regime becoming negligible. In other words, beliefs under incomplete information converge towards those under perfect information if the economy stays sufficiently long in a block. Consequently, as shown in Appendix A.1, $prob(\xi_t^Z = 1|\tau_t^1)$ and $prob(\xi_t^Z = 3|\tau_t^2)$ converge to zero at the convergence points of Block 1 and Block 2 denoted by $\tilde{\lambda}_{b1}$ and $\tilde{\lambda}_{b2}$.[7]

2.5.2 Prior beliefs across blocks

According to equations (10) and (11), beliefs about the regimes in period t depend on the beliefs in $t-1$, which, in turn, are conditional on the fact that the same block was realized

[7] In this sense, the imperfect-information model nests the perfect-information model in the limiting cases in which the probability of being in the transitory regime converges to zero.

in $t-1$. Hence, the initial beliefs at the moment the blocks switch cannot be pinned down by the recursions. To determine the initial beliefs, agents apply prior beliefs, which are based on the relative probabilities of being in the regimes of the new block. For instance, when the economy shifts from Block 2 to Block 1, agents apply the relative probability of being in Regime 1 as

$$p(\xi_t^Z = 1|\xi_{t-1}^Z = 3, \tau^1 = 1) = \frac{p_{31}}{p_{31} + p_{32}},$$

if the economy was in Regime 3 before the block switched, or as

$$p(\xi_t^Z = 1|\xi_{t-1}^Z = 4, \tau^1 = 1) = \frac{p_{41}}{p_{41} + p_{42}},$$

if the economy was in Regime 4 before the block switched. Analogously, one can pin down the relative probabilities of entering Regime 3 when the economy shifts from Regime 1 or Regime 2.

Using the calibrated transition matrix (9), the relative probabilities are

$$\frac{p_{31}}{p_{31} + p_{32}} = 0.74 > \frac{p_{41}}{p_{41} + p_{42}} = 0.64,$$

$$\frac{p_{13}}{p_{13} + p_{14}} = 0.71 < \frac{p_{23}}{p_{23} + p_{24}} = 0.85.$$

The calibration of the transition matrix implies that the relative probabilities of entering the regimes in the new block differ depending on the previous regime. The relative probability of entering the transitory regime in Block 1 is higher when the economy has been in Regime 3 compared to Regime 4, whereas the opposite holds for entering the transitory regime in Block 2. This indicates that beliefs about the previous regime matter when forming prior beliefs.

Accordingly, agents' prior beliefs are an average of the relative probabilities outlined above, weighted by the beliefs about the previous regime before the blocks switched. That is, the prior probability of being in the transitory regime when the economy shifts to Block 1 is defined as

$$p(\xi_t^Z = 1|I_t) = \frac{prob\{\xi_{t-1}^Z = 3|I_{t-1}\}p_{31} + (1 - prob\{\xi_{t-1}^Z = 3|I_{t-1}\})p_{41}}{prob\{\xi_{t-1}^Z = 3|I_{t-1}\}(p_{31} + p_{32}) + (1 - prob\{\xi_{t-1}^Z = 3|I_{t-1}\})(p_{41} + p_{42})}, \tag{12}$$

whereas the prior probability of being in the transitory regime at the moment the economy shifts to Block 2 is

$$p(\xi_t^Z = 3|I_t) = \frac{prob\{\xi_{t-1}^Z = 1|I_{t-1}\}p_{13} + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})p_{23}}{prob\{\xi_{t-1}^Z = 1|I_{t-1}\}(p_{13} + p_{14}) + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})(p_{23} + p_{24})}, \tag{13}$$

with the information set I_t containing the history of blocks and beliefs. In this sense, the evolution of beliefs in the imperfect-information model can be determined by the history of blocks and initial beliefs.

To solve the imperfect-information model, we follow Bianchi and Melosi (2016) and recast the model as a rational expectations perfect-information model that tracks agents' beliefs with a new set of regimes. More precisely, the recast perfect-information model is governed by a new transition matrix $\hat{P}$ which contains the probabilities to stay in or to shift between the blocks, that is, the new set of regimes ζ_t^Z. These probabilities are in turn calculated using the imperfect beliefs pinned down by (10-11) and (12-13). Consequently, the entries of $\hat{P}$ keep track of agents' beliefs. For instance, the probability to remain in Block 1 is given by $prob\{\xi_{t-1}^Z = 1|I_{t-1}\}p_{11} + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})p_{22}$, while the complementary probability represents the probability of switching to Block 2.

Prior beliefs change over time, as they depend on previous beliefs, and the paths of beliefs within a block in turn also vary, because they depend on the prior beliefs. As a result, there are potentially infinitely many paths of beliefs. To overcome the infinite-dimensionality problem, we follow Bianchi and Melosi (2016) and set up a finite grid to approximate agents' beliefs and to pin down the transition matrix $\hat{P}$. Appendix A.2 provides a detailed explanation of this procedure.

2.6 Model stationarization

Since the TFP process (7) exhibits a unit root, we make the model stationary by expressing the exogenous TFP process in growth rates $\mu_t = \ln\left(\frac{Z_t}{Z_{t-1}}\right)$ and by scaling output and wages by Z_t, that is, $\tilde{Y}_t = \frac{Y_t}{Z_t}$, $\tilde{W}_t = \frac{W_t}{Z_t}$. As in Schorfheide (2005), Liu et al. (2011), and Bianchi and Melosi (2016), we define the steady state as the stationary equilibrium with no regime switches and TFP growth at its ergodic level $\bar{\mu}$. In Appendix B, we provide details on how we make the model stationary and compute the steady state. In what follows, the steady states of output, inflation, and the interest rate are denoted by Y_{ss}, π_{ss}, and R_{ss}, respectively.

2.7 Monetary policy

We assume that monetary policy follows a two-regime MS process $\xi_t^{MP} \in \{1, 2\}$ that governs the nominal interest rate R_t

$$\left(\frac{R_t}{R_{ss}}\right) = (1 - D_{\xi_t^{MP}})\left[\left(\frac{\pi_t}{\pi_{ss}}\right)^{\phi_\pi}\left(\frac{\tilde{Y}_t}{Y_{ss}}\right)^{\phi_y}\right] + D_{\xi_t^{MP}}\frac{1}{R_{ss}}, \tag{14}$$

where ϕ_π and ϕ_y denote the inflation and output Taylor rule coefficients, while $D_{\xi_t^{MP}}$ is a dummy variable indicating the regime of monetary policy. If $D_{\xi_t^{MP}}$ takes the value 0, the ZLB does not bind, and the central bank follows a standard Taylor rule ($\xi_t^{MP} = 1$), referred to as 'normal monetary policy'. If $D_{\xi_t^{MP}}$ takes the value 1, the ZLB binds, and the central bank sets the interest rate to zero ($\xi_t^{MP} = 2$).

The monetary policy regimes evolve according to the transition matrix H, that is

$$H = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix}. \tag{15}$$

H_{11} and H_{22} represent the probability of remaining within the normal monetary policy regime and the ZLB regime, respectively, whereas H_{12} and H_{21} denote the probabilities of switching between the two regimes. We calibrate the entries of H based on observed interest rate movements from 1970Q1 to 2022Q4. Specifically, we count the number of times the regimes persisted in the data (s_{11} and s_{22}) and the number of times the regimes switched to the other regime (s_{12} and s_{21}). Using this information, we calibrate the entries of H as relative probabilities. For example, we set $H_{12} = \frac{\sum s_{12}}{\sum_{j=1}^{2} s_{1j}}$.

This procedure yields the following transition matrix:

$$H = \begin{bmatrix} 0.99 & 0.01 \\ 0.06 & 0.94 \end{bmatrix}.$$

It is worth emphasizing that, in both the perfect- and imperfect-information models, agents know the monetary policy regime in place and its transition probabilities, as given by the matrix H. Thus, the only difference between the models lies in the information about the persistence of TFP shocks.

2.8 Four equation New Keynesian Model with Bayesian learning

By log-linearizing the stationary model, we can represent our model as a four-equation model.[8] In this regard, $\hat{y}_t$, $\hat{\pi}_t$, and $\hat{r}_t$ denote the log-deviations of stationary output, inflation, and the interest rate from their steady-state levels, while $\hat{\mu}_t$ and $\hat{\mu}(\xi_t^Z)$ reflect the deviations of TFP growth and the TFP drift from their ergodic mean $\bar{\mu}$.

[8]For further details, see Appendix B.1-B.5.

$$\hat{y}_t = \mathbb{E}_t\, \hat{y}_{t+1} + \mathbb{E}_t\, \hat{\mu}_{t+1} - (\hat{r}_t - \mathbb{E}_t\, \hat{\pi}_{t+1}),$$

$$\hat{\pi}_t = \beta\, \mathbb{E}_t\, \hat{\pi}_{t+1} + \frac{(1-\beta\theta)(1-\theta)}{\theta}(1+\phi)\hat{y}_t,$$

$$\hat{r}_t = (1 - D_{\xi_t^{MP}})(\phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t) - D_{\xi_t^{MP}} \log R_{ss},$$

$$\hat{\mu}_t = \hat{\mu}(\xi_t^Z).$$

2.9 The core mechanism in a nutshell

In the model, TFP growth captures both supply- and demand-side effects associated with the surge in inflation in the aftermath of the COVID-19 pandemic. The intuition is that higher expected TFP growth reflects anticipated improvements in productivity due to the easing of COVID-19 restrictions and the resolution of supply chain disruptions. Higher expected future productivity, in turn, raises expected future resources and thereby induces forward-looking households to increase current demand, as implied by the IS curve.

As is evident from the four-equation model, only expectations about future TFP growth enter, so that its dynamics are entirely driven by these expectations. These expectations are formed based on agents' information sets, which differ between the complete- and incomplete-information models.[9] Perfectly informed agents know which regime they are in when forming expectations. In contrast, imperfectly informed agents assign probabilities to being in either the transitory or the persistent growth regime, which gives rise to waves of optimism and pessimism that depend on the underlying TFP growth rate. For instance, during a persistent high-growth regime, they attach positive probability to being in a transitory regime and are thus too pessimistic, while in a persistent low-growth regime, they are too optimistic. The opposite pattern arises in transitory regimes, implying that optimism and pessimism systematically vary with both the true regime and the level of TFP growth. This leads to systematic overshooting and undershooting in TFP growth expectations, which in turn drive macroeconomic dynamics and monetary policy that diverge from those under perfect information. Importantly, since the Taylor rule responds to endogenous variables driven by private sector expectations, the central bank effectively shares the same information set as private agents in both models.

[9] We refrain from introducing state variables because we want differences in dynamics across models to be driven entirely by the information asymmetry.

2.10 Solution techniques

The perfect-information model featuring the MS process with transition matrix P can be expressed as

$$\Gamma_0(\xi_t^Z)S_t = \Gamma_c(\xi_t^Z) + \Gamma_1(\xi_t^Z)S_{t-1} + \Upsilon\eta_t, \tag{16}$$

where S_t denotes the vector of the model variables, η_t contains endogenous expectation errors, and Γ_0, Γ_c, Γ_1 contain the model's parameters that are controlled by the hidden variable ξ_t^Z.

As discussed above, we recast the imperfect-information model as a perfect-information model in terms of ζ_t^Z governed by $\hat{P}$, as follows

$$\Gamma_0(\zeta_t^Z)S_t = \Gamma_c(\zeta_t^Z) + \Gamma_1(\zeta_t^Z)S_{t-1} + \Upsilon\eta_t. \tag{17}$$

Solution under normal monetary policy

In Subsections 4.1 and 4.3, we consider the model under a standard Taylor rule, with $\xi_t^{MP} = 1$ and $D_{\xi_t^{MP}} = 0$ throughout.[10] Here, the regime changes enter additively and only affect the constants $\Gamma(.)$ in (16) and (17). Therefore, we follow Schorfheide (2005), Liu et al. (2011), Bianchi and Melosi (2016), and Bianchi et al. (2018) and expand the state space with dummy variables to solve the perfect- and imperfect-information model using *gensys* (Sims, 2002).

In either case, the solution can be expressed as an MS vector autoregression. For instance, solving model (17) results in

$$S_t = c(\zeta_t^Z, \hat{P}) + T(\zeta_t^Z, \hat{P})S_{t-1}.$$

Since our model lacks a state variable and stochastic shocks, it follows that $T = 0$. Hence, the model's dynamics are purely driven by expectations, which depend on the regime in place and the associated probabilities of shifting across the regimes. As a result, beliefs about the persistence become crucial for the model's dynamics.

Solution under the ZLB

In Subsection 4.2, the ZLB on interest rates binds, that is, $D_{\xi_t^{MP}} = 1$ and $\xi_t^{MP} = 2$. To incorporate the dynamics of the TFP process and the ZLB on the interest rate, we expand the regimes by combining the Markov chains, yielding, for instance, $\xi_t \equiv [\zeta_t^Z, \xi_t^{MP}]$ in the

[10] In these subsections, the ZLB regime and the associated monetary policy transition dynamics are not considered.

imperfect-information model. As a result, the number of regimes expands, meaning that the learning mechanism takes place under normal monetary policy and under the ZLB. The corresponding transition matrix of $\hat{P}H$ is the Kronecker product of $\hat{P}$ and H, that is $\hat{P}H \equiv \hat{P} \otimes H$. Consequently, agents take into account the joint probabilities of the system shifting between the TFP growth regimes and the monetary policy regimes. The imperfect-information model capturing the ZLB can be expressed as

$$\Gamma_0(\xi_t)S_t = \Gamma_c(\xi_t) + \Gamma_1(\xi_t)S_{t-1} + \Upsilon\eta_t.$$

In the perfect and imperfect information models incorporating the ZLB, regime changes enter multiplicatively due to structural breaks in the Taylor rule coefficients. Hence, we solve these models using the solution technique developed by Farmer et al. (2009).

3 Calibration

The model's parameters are calibrated using standard values defined in the literature. We set the discount factor β to 0.995, consistent with an annualized steady-state interest rate of roughly 3% given the ergodic TFP growth rate of $\bar{\mu} = 0.8\%$ as described in Section 2.3. We set the inverse Frisch elasticity ϕ equal to 2 following Chetty (2012). In line with Christiano et al. (2011), we fix the price contract duration parameter θ at 0.85, while the price elasticity of demand for intermediate goods ϵ equals 11, implying a markup of 1.1. These values translate into a Phillips curve slope of 0.08. The Taylor rule coefficients on inflation and output, ϕ_π and ϕ_y, are set to 1.5 and 0.2, respectively.

4 Results

In what follows, Subsection 4.1 examines the implications of incomplete information about TFP growth persistence for inflation and monetary policy, Subsection 4.2 analyzes the looking-through policy at the ZLB, and Subsection 4.3 studies welfare, uncertainty, and optimal monetary policy.

4.1 Incomplete information and monetary policy

To assess the implications of our learning mechanism, we construct a sequence of TFP growth realizations to replicate the business cycle fluctuations since the early 2000s. These fluctuations are matched to the cycles in US real GDP, obtained by filtering the data using the Hodrick-Prescott filter with smoothing parameter $\lambda = 1600$. The gray-dotted line in Figure 3 shows the filtered cycles in real GDP, while the blue-dotted line displays the detrended TFP. The gray bars indicate NBER recessions around the collapse of the dot-com bubble (2001Q2-2001Q4), the GFC (2008Q1-2009Q2) and the COVID-19 recession (2020Q1-2020Q2). To reproduce the cycles, we group them into the following segments, separated by the horizontal dashed line: the decline in economic activity around the dot-com bubble burst ('Dot-Com Recession'); the boom between the Dot-Com Recession and the GFC ('Pre-GFC Boom'); the GFC; the COVID-19 recession; and the post-COVID recovery with high inflation rates ('Post-COVID Inflation'). Since neither GDP nor TFP exhibits a clear trend between the GFC and the COVID-19 recession, we characterize this period as a sequence of shifts between transitory high- and low-growth regimes ('Transitory Regimes').

Figure 3: Detrended real GDP and TFP

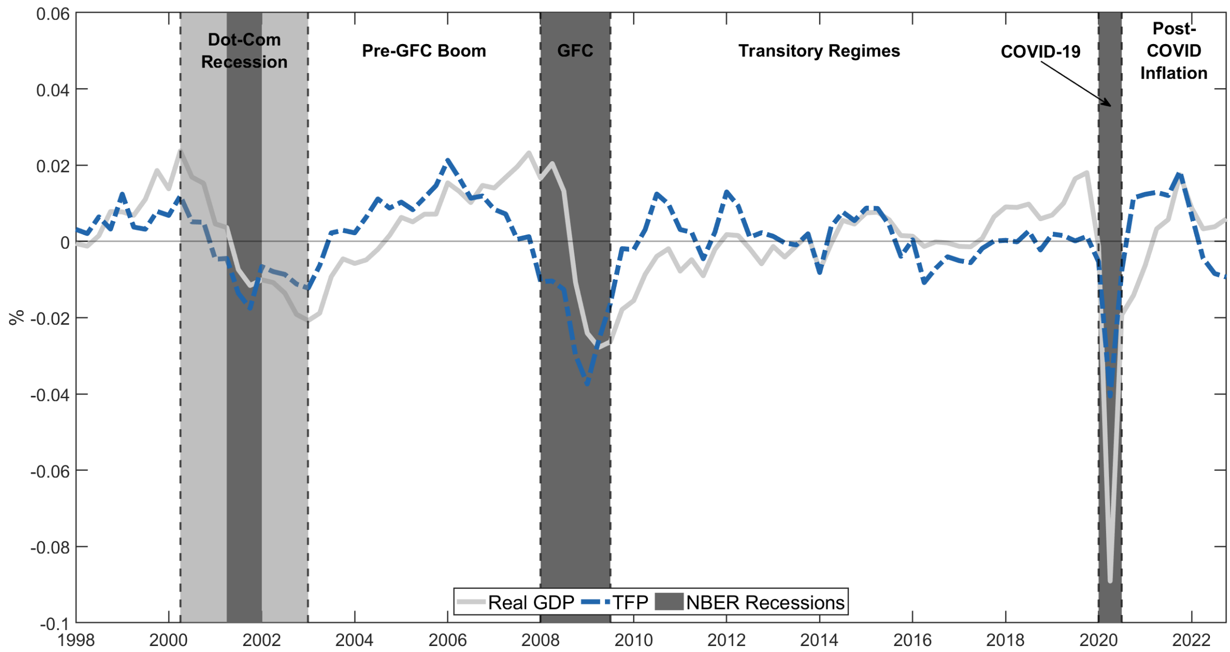

Notes: Detrended logarithmic real GDP and TFP filtered with the Hodrick–Prescott filter. Sources: Fernald (2012), FRED St. Louis, and NBER.

Motivated by the close relationship between detrended GDP and detrended TFP shown in Figure 3, we assume in our model that business cycles are solely driven by TFP growth. We design the end of the Dot-Com Recession, the GFC, and the COVID-19 recession as persistent low-growth regimes, while the Pre-GFC Boom and the Post-COVID Inflation are modeled as persistent high-growth regimes. The COVID-19 recession lasts for three periods and is modeled to be shorter than the GFC. For the sake of clarity, we consider a sequence of three transitory regimes between the GFC and the COVID-19 recession, which represent the fluctuations between the GFC and the downturn associated with the COVID-19 crisis. Each transitory regime lasts two periods. In Figure 4, the black- and white-shaded areas indicate the sequence of low- and high-growth regimes. The TFP

Figure 4: Sequence of simulated TFP growth

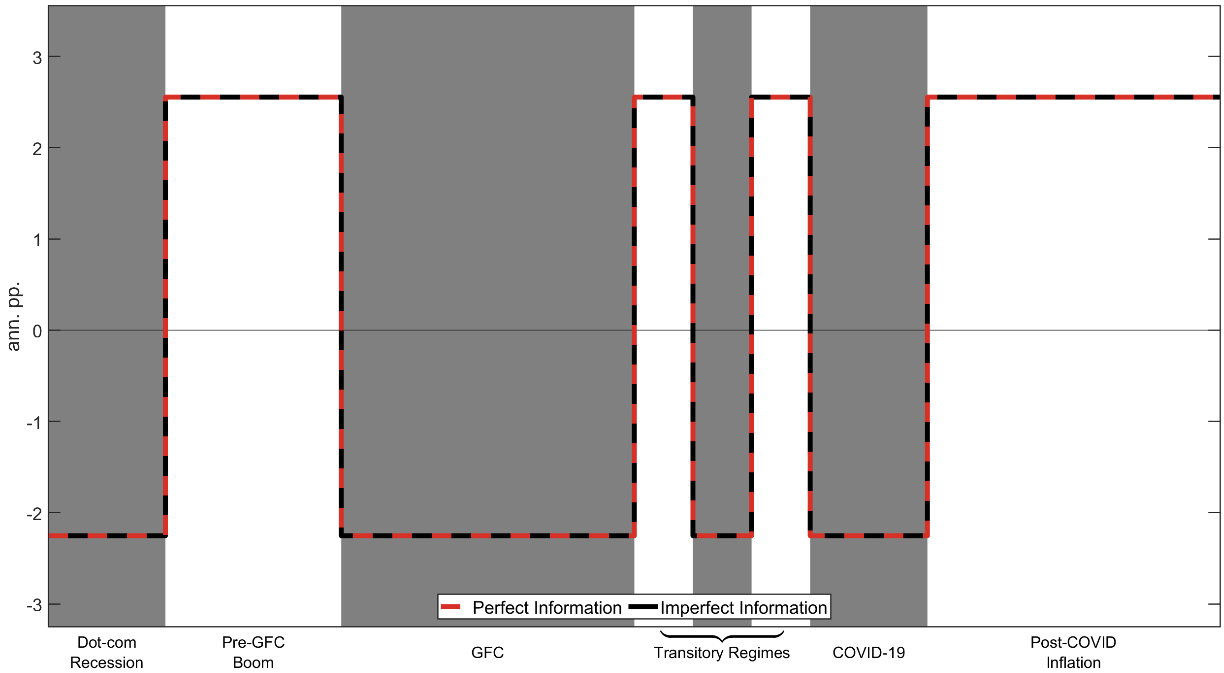

Notes: The figure shows the annualized percentage-point (ann. pp.) deviations of TFP growth from its ergodic mean for the imperfect-information model (black-solid line) and the perfect-information model (red-dashed line). The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas indicate episodes of high TFP growth.

growth realizations are identical for the perfect information and imperfect information models, which are represented by the red-dashed and black-solid lines, respectively.

Given this sequence of TFP growth, Figure 5 highlights the evolution of agents' beliefs in our model. Panel 1 visualizes the beliefs of being in the persistent low-growth regime, while Panel 3 shows the beliefs of being in the persistent high-growth regime. Here, the black-solid lines represent beliefs under imperfect information, and the red-dashed lines visualize beliefs under perfect information. Since perfectly informed agents are aware of the regime in place, their beliefs about being in the persistent regimes jump between zero and one.

We initialize imperfectly informed agents' beliefs such that they are relatively sure of being in a persistent low-growth regime at the beginning of the simulation. This indicates the end of the economic downturn associated with the dot-com bubble burst. Following detrended output and TFP growth in Figure 3, the Dot-Com Recession lasts for three more periods in our simulation, during which agents become almost completely sure that the economy is in the persistent low-growth regime, or, to put it differently, their beliefs converge towards those under perfect information as highlighted by Panel 1.[11]

[11]Observing TFP growth in period t, agents assign zero probability to the persistent high-growth

Figure 5: Regime changes and belief formation

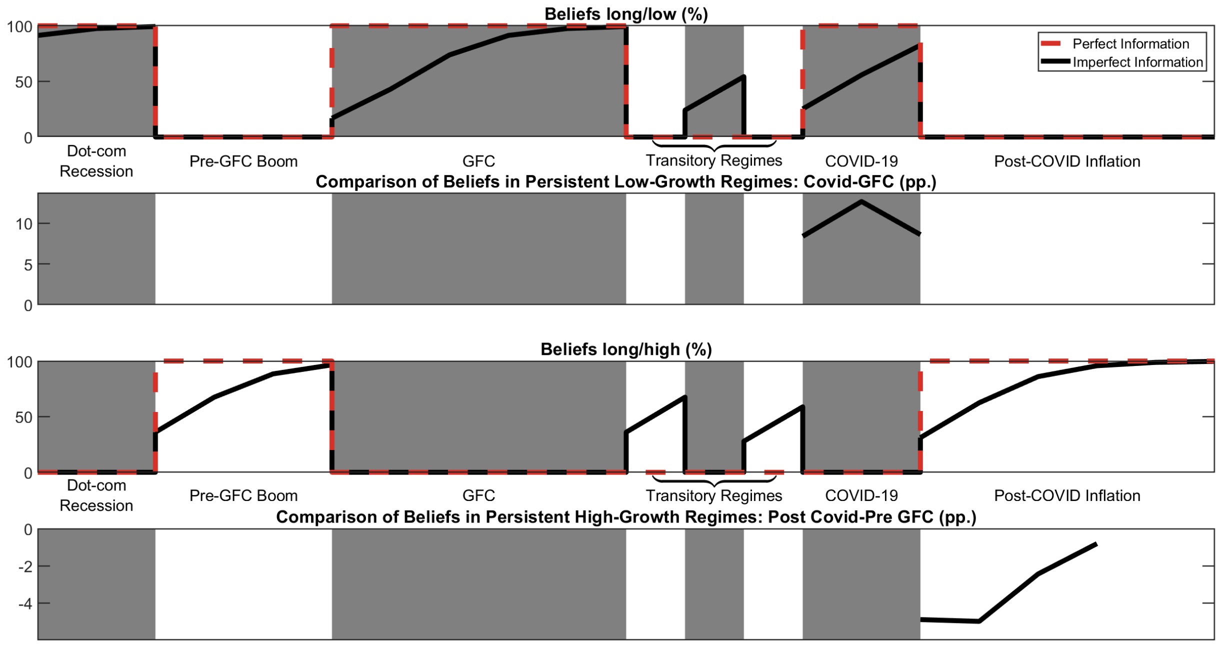

Notes: The evolution of beliefs in the imperfect-information model and the perfect-information model. The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas reflect episodes of high TFP growth. The first and third panels display the beliefs of being in the persistent low-growth regime (long/high) and the beliefs of being in the persistent high-growth regime (long/low). Panel 2 illustrates the differences in beliefs about being in the persistent low-growth regime between the three periods of the COVID-19 recession and the first three periods of the GFC. Panel 4 shows the differences in beliefs about being in the persistent high-growth regime between the first four periods of the post-COVID inflation era and the pre-GFC boom.

In the fourth period, the economy shifts towards the high-growth block for four periods, indicating the boom before the GFC. In the first period, agents apply a prior probability of roughly 0.36 to be in the persistent high-growth regime. Subsequently, agents become more certain of being in the persistent regime with each consecutive high-growth realization, before the GFC hits the economy. In this process, beliefs evolve convexly, as learning gains are highest at the beginning of the block and decrease over time. The discrepancy between agents' beliefs under perfect and imperfect information visualizes that imperfectly informed agents incorrectly take the probability of being in the transitory regime into account. As a result, agents are too *pessimistic* about future growth rates.

In period eight, the beginning of the GFC, agents apply a prior belief of roughly 0.17 of being in the persistent low-growth regime before they become more certain with each consecutive period experiencing negative TFP growth realizations. By that means, imperfectly informed agents assign positive probability to being in the transitory low-growth regime and are thus too *optimistic* about future growth.

regime during low-growth episodes, and vice versa.

The opposite pattern emerges in periods of transitory high and low growth (periods 14–19). Perfectly informed agents assign zero probability to being in the persistent regime. Imperfectly informed agents, on the other hand, incorrectly consider the probability of being in the persistent regime during transitory regimes. As a result, agents are too optimistic in cases of transitory high-growth regimes and too pessimistic in cases of transitory low-growth regimes.

During the COVID-19 recession, which lasts in our simulation from period 20 to 22, imperfectly informed agents, once again, are too optimistic as they assign positive probability to the transitory low-growth regime. Conversely, in the subsequent recovery and post-COVID inflation era, they are too pessimistic, assigning positive probability to the transitory high-growth regime. After six consecutive periods in the high-growth environment, agents assign a probability of 0.99 to be in the persistent regime. In this sense, beliefs have almost fully converged towards those under perfect information.[12]

Under the dynamic prior, the history of TFP growth realizations and beliefs becomes relevant for forming prior beliefs at the moment the blocks shift. The different priors, in turn, affect the trajectory of beliefs within the blocks. Panels 2 and 4 show the relevance of the dynamic prior assumptions for the evolution of beliefs. The panels emphasize the differences in beliefs regarding the probability of being in the persistent regime between the second and first persistent periods for both blocks. In this regard, Panel 2 shows the difference in beliefs between the three periods of the COVID-19 recession (periods 20–22) and the first three periods of the GFC (periods 8–10). The difference is positive, indicating that agents are more certain of being in the persistent low-growth regime when entering the COVID-19 recession than at the beginning of the GFC, reflecting a comparatively more pessimistic economic outlook. This heightened pessimism stems from the historical experience of prolonged low growth during the GFC, coupled with subsequent fluctuations between transitory low and high-growth rates. The difference in beliefs peaks at approximately 13 percentage points in the second period before experiencing a slight decline in the third period. This finding adds another perspective on why the Fed swiftly moved toward a more accommodative monetary policy at the onset of the COVID-19 crisis, compared to the beginning of the GFC.

In addition, Panel 4 displays the differences in beliefs about being in the persistent high-growth regime between the first four periods of the post-COVID inflation era (periods 23–26) and the expansion prior to the GFC (periods 4–7). The negative difference indicates that agents are less certain about being in the persistent regime during the post-COVID inflation era than during the pre-GFC boom, pointing to a relatively more pessimistic

[12]Technically, beliefs fully converge to the perfect-information benchmark after 10 quarters, implying that agents assign probability one to the persistent regime.

outlook on the economy. In the second period of the post-COVID inflation era, the probability assigned to the persistent regime is almost five percentage points lower. This more pessimistic outlook can be explained by the recent history of two persistent low-growth phases and fluctuations between transitory regimes. The model thus helps explain why households, firms, and the central bank were initially more inclined to view inflation as transitory at the beginning of the post-COVID inflation period.

Figure 6 illustrates the economic consequences of these waves of optimism and pessimism in agents' beliefs. As discussed above, beliefs about the duration of TFP growth affect the economy through TFP growth expectations, which are shown in the first panel. Panel 2 displays the evolution of detrended output deviations from the steady state, which are intended to correspond to the output deviations observed in the data, as shown in Figure 3. The remaining panels report the dynamics of inflation and the interest rate. For all variables, results are presented under imperfect information (black-solid lines) and perfect information (red-dashed lines).

Figure 6: The effects of imperfect information on the economy

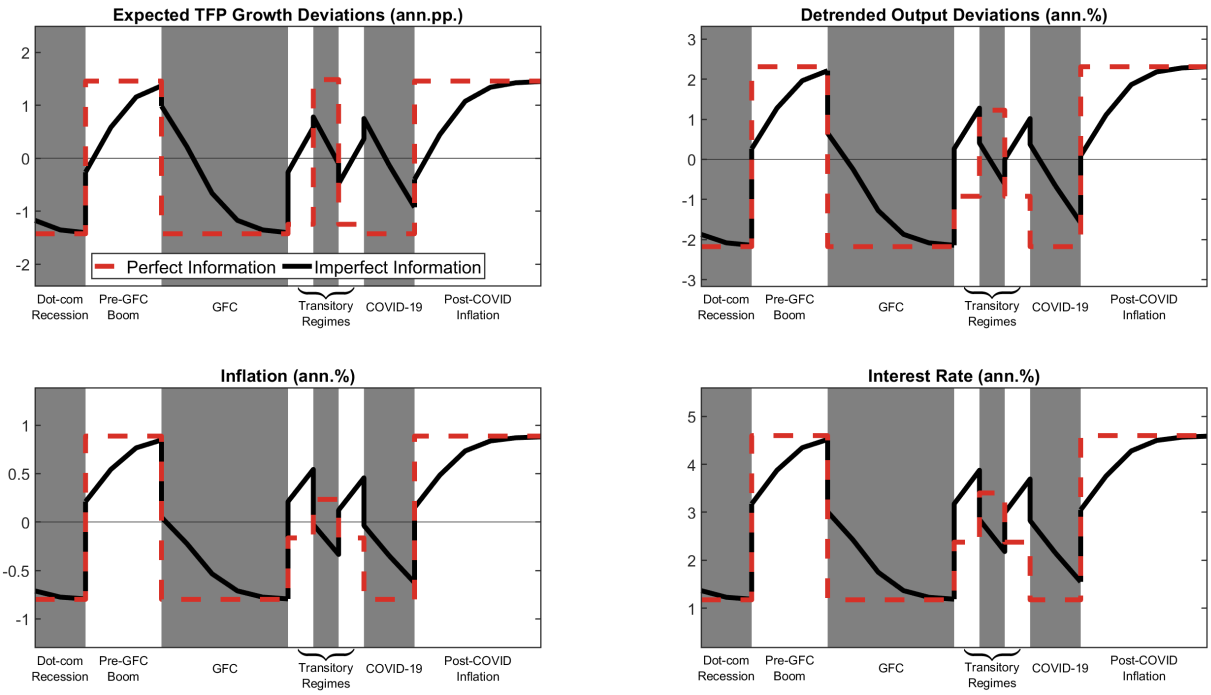

Notes: The effects of imperfect information (black-solid line) and perfect information (red-dashed line) on the economy. The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas denote episodes of high TFP growth. The first panel reflects the percentage-point (pp.) deviations of TFP growth expectations from its ergodic mean, while the second panel displays the percent deviations of detrended output from its steady state. The third and fourth panels display inflation and the interest rate.

At the onset of persistent regimes, TFP growth expectations are too low in high-growth regimes and too high in low-growth regimes under imperfect information relative to fundamentals, reflecting pessimistic and optimistic beliefs, respectively. However, TFP growth expectations rise in high-growth regimes and fall in low-growth regimes as agents become more confident that the economy is in a persistent regime, eventually converging to the perfect-information benchmark. In the perfect-information model, by contrast, TFP growth expectations jump instantly when the economy switches between blocks and remain constant until the next transition, reflecting the discrete adjustments in beliefs shown in Figure 5. Consequently, output, inflation, and the interest rate also adjust instantly after the blocks switch, remaining at their levels until the blocks change again.

Consistent with TFP growth expectations, in the imperfect-information model, the initial responses of output, inflation, and the interest rate are too small in persistent high-growth regimes and too large in persistent low-growth regimes relative to fundamentals. As the economy remains in the respective block and agents learn that the regime is persistent, the variables adjust gradually, converging to their perfect-information counterparts as certainty increases.[13] In this respect, the incomplete information model introduces an endogenous source of persistence without the need to add state variables, autocorrelated shocks, or ad hoc adjustments, in line with Milani (2007) and Fuhrer (2017). As a result, the imperfect-information model replicates the gradual rise in inflation and the gradual monetary policy tightening following the COVID-19 crisis (see Figure 1). In the model, it takes six quarters until the interest rate converges to the level justified by the fundamentals, which is 4.7%. This aligns with the data, considering that the Fed raised the Federal Funds Rate for seven consecutive quarters, starting in the first quarter of 2022, as highlighted by Figure 1.[14]

Optimism and pessimism in beliefs also matter in analyzing the implications of the transitory regimes. In the first transitory high-growth regime after the GFC (periods 14-15), the economy recovers in both models. However, TFP growth expectations, and thus output and inflation, rise more strongly in the imperfect-information model than in the perfect-information model. This stems from the fact that imperfectly informed agents are too optimistic, as they assign positive probability to the persistent regime. In subsequent transitory regimes, macroeconomic variables move differently in the perfect- and imperfect-information models. In the perfect-information model, agents are sure to be in the transitory regime, and, given the transition matrix P, the economy is most likely to shift from either of the transitory regimes to the other block with a probability

[13] Beliefs about being in the persistent regime reach 0.99 at the end of the GFC and the post-COVID inflation period. This is sufficient to ensure that, at these points, the macroeconomic variables in the incomplete information model equal their counterparts in the perfect-information model.

[14] One could argue that the Fed should have lifted the interest rate from the ZLB already three quarters earlier in the second quarter of 2021 when inflation exceeded the Fed's 2% inflation target. Consequently, it would have taken the Fed ten quarters to bring the interest rates to the level aligned with the fundamentals, which would still be close to the adjustment time suggested by our model. We analyze the implications of the central bank remaining at the ZLB for an additional three quarters in Subsection 4.2.

of roughly 0.8. Accordingly, TFP growth expectations along with output and inflation rise in the transitory low-growth regime and fall in the subsequent transitory high-growth regime. In contrast, imperfectly informed agents are too pessimistic during transitory low-growth regimes and too optimistic during transitory high-growth regimes, as they assign positive probability to the persistent regime. In each case, they assign a higher probability of remaining in the respective block than perfectly informed agents, so that output and inflation fall during transitory low-growth regimes and rise during transitory high-growth regimes.[15]

Since output and inflation move in opposite directions across models during transitory regimes, so does the interest rate. In the perfect-information model, the central bank responds procyclically to contemporaneous TFP growth realizations, raising the policy rate during transitory low-growth phases and lowering it during transitory high-growth phases. By contrast, in the imperfect-information model, the central bank's role is to moderate the economy during episodes of transitory high-growth and to stimulate it during periods of transitory low-growth. In other words, the monetary policy response to temporary TFP growth shocks is countercyclical. It is worth emphasizing that during persistent regimes, in contrast, the central bank's responses in both the imperfect and perfect information models align, as both pursue a countercyclical monetary policy.

4.2 Looking-through policy at the Zero Lower Bound

Since the central bank in our model follows a Taylor rule, it raises the policy rate when inflation increases. To better capture the early post-COVID episode, in which inflation was elevated while the policy rate remained at the ZLB, we examine the implications of the model under a looking-through policy at the ZLB. In what follows, we impose the ZLB exogenously, consistent with the MS structure of the policy regime, once the model-implied policy rate reaches zero. As demonstrated in Subsection 4.1, the TFP growth rates derived from historical realizations are insufficient to bring the model-implied nominal interest rate to the ZLB. To induce a binding constraint in this counterfactual analysis, we therefore specify a more pronounced negative realization of TFP growth, μ_l = -4.2%. This value closely approximates the average TFP growth rate observed during the GFC and the COVID-19 recession.

[15] In the incomplete information model, there is a discrepancy between TFP growth expectations and the macroeconomic variables at the beginning of the transitory periods. For example, in the first period of a transitory low-growth regime, agents tend to believe they are currently in a transitory low-growth regime. According to the transition matrix, however, it is most likely that the economy will switch to the transitory high-growth regime in the next period, which causes TFP growth expectations to increase. In the transitory high-growth regime, it is in turn most likely to move back to the low-growth block, leading to a decline in (expected) output and (expected) inflation. The same logic applies in reverse for the first period of a transitory high-growth regime.

For the evolution of beliefs displayed in Figure 5, Figure 7 presents the corresponding model dynamics when the ZLB binds during the GFC and the COVID-19 recession under a looking-through policy in the early post-COVID period.[16] During the GFC, it takes three quarters for the model-implied interest rate to reach the ZLB, whereas during the COVID-19 crisis, it takes only two. This mirrors the trajectory of interest rates in the US, as the policy rate hit the ZLB in a much shorter time span during the COVID-19 crisis than during the GFC (see Figure 1).[17] In the model, this difference can be traced back to agents' beliefs: as discussed in the previous section, agents are more pessimistic about being in the persistent low-growth regime at the onset of the COVID-19 crisis than at the beginning of the GFC, leading to a faster decline in TFP growth expectations and, thus, in output, inflation, and the interest rate.

Figure 7: Dynamics at the Zero Lower Bound under a looking-through policy

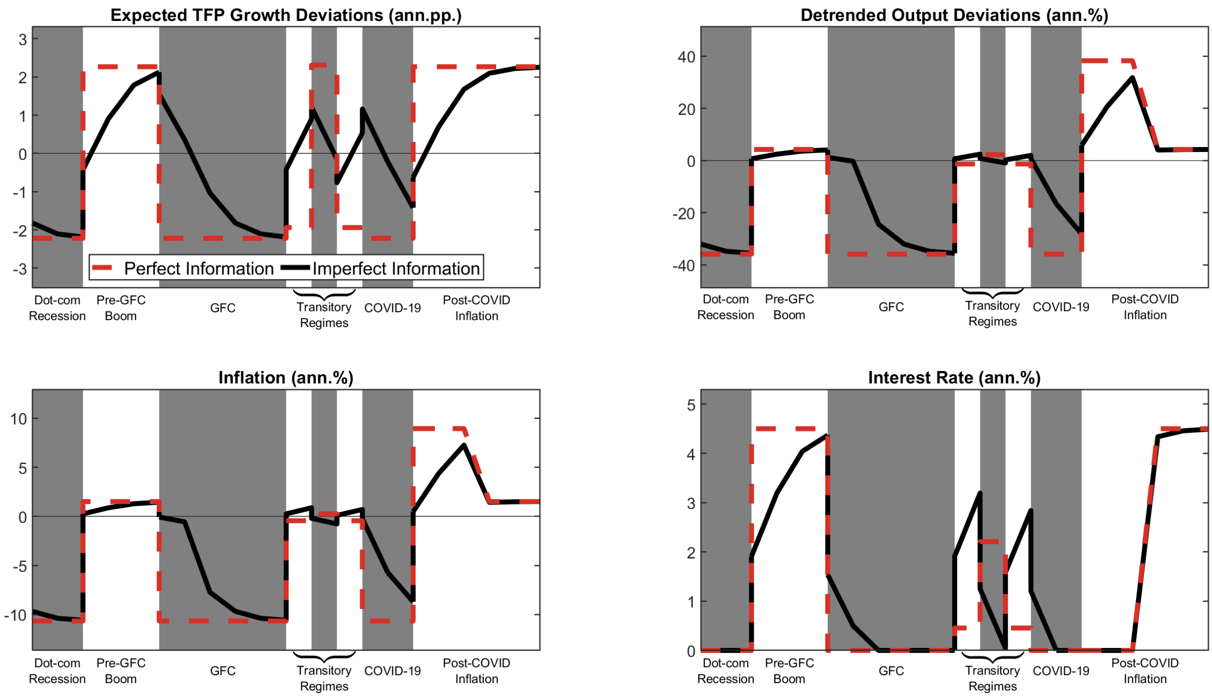

Notes: Dynamics at the Zero Lower Bound under a looking-through policy with imperfect information (black-solid line) and perfect information (red-dashed line). The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas denote episodes of high TFP growth. The first panel reflects the annualized percentage-point (ann. pp.) deviations of TFP growth expectations from its ergodic mean, while the second panel displays the annualized percentage deviations of detrended output from its steady state. The third and fourth panels display the annualized inflation and the interest rate.

[16]As the model contains only a single negative TFP growth realization, the ZLB binds also during the Dot-Com Recession, contrary to the historical evidence.

[17]During the GFC, the policy rate declined gradually to the ZLB from mid-2007 to the end of 2008. By contrast, at the onset of the COVID-19 crisis, the policy rate fell to the ZLB within a single quarter—indeed, within a single month at the beginning of 2020—although it had already been declining since the third quarter of 2019 (see Figure 1).

The ZLB reinforces the recession and the decline in inflation during the persistent low-growth regimes in the model, with the inflation rate dropping by up to eight percent and output dropping by up to 20 percent.[18] The stronger immediate adjustment under perfect information underscores the role of incomplete information in generating endogenous persistence, as agents gradually learn about the persistence of TFP growth. In the aftermath of the COVID-19 crisis, we explore the effects of a looking-through monetary policy. In this case, the central bank deliberately delays lifting the interest rate from the ZLB, pursuing a lower-for-longer strategy until it is confident that the economic recovery and the associated inflation dynamics are long-lived. This approach aims to ensure that the recovery is fully underway and not curtailed too early by a premature tightening of monetary policy. Like Buss and Traficante (2025), we assume that the central bank is willing to raise the interest rate once the probability that fundamentals driving inflation are long-lasting exceeds a certain threshold. As mentioned above in Section 1, the Fed raised the policy rate three quarters too late, considering that inflation had already exceeded the 2% target in the second quarter of 2021. Moreover, the FOMC statements suggest that the Fed required nine months to ascertain that inflation dynamics were persistent. In April 2021, the Fed added to its FOMC statement that inflation was driven by 'transitory factors'. This phrase was removed from its statement in December 2021, as noted by Ihrig and Waller (2024). To be consistent with the three quarters required to ensure certainty about the persistence of high TFP growth, we assume the central bank lifts off from the ZLB as soon as the belief about being in the persistent high-growth regime exceeds 0.87. That is, $prob(\xi_t^Z = 2|\tau_t^1) > 0.87$. According to Panel 2 in Figure 5, the beliefs about being in the persistent high-growth regime exceed the threshold of 0.87 in period 26 during the periods reflecting the post-COVID inflation era. In turn, in Figure 7, the interest rate remains at the ZLB from period 23 to 25, even though the economy has entered a high-growth block.

The additional periods at the ZLB during the early post-COVID recovery result in a sharp increase in output and inflation relative to the unconstrained case. For example, in the perfect-information model, inflation rises by more than a factor of five. Inflation and output increase gradually under imperfect information, albeit at higher levels. These results indicate that for central banks, early recognition of whether an expansion is temporary or persistent is crucial for the timing of exiting the ZLB and, ultimately, for their task of keeping inflation in check. Nevertheless, after three quarters, the central bank lifts the interest rate from the ZLB by increasing it in the 26th period, leading to a drop in output

[18] Compared to the declines in US GDP and inflation during the Great Recession observed in the data, our simulated contractions are rather high. This refers to the problem highlighted by Fernández-Villaverde et al. (2015) that one cannot create moderate drops in economic activity while letting the ZLB bind for an extended period of time.

and inflation. By this time, agents have updated their beliefs for four consecutive periods and are highly confident that TFP growth is persistent, assigning a probability of 0.96 to the persistent regime. This indicates that beliefs have nearly converged, leaving only minor differences between the perfect- and imperfect-information models once the central bank lifts the policy rate from the ZLB.

4.3 Welfare, uncertainty, and optimal monetary policy

How do welfare and uncertainty evolve in our model, and how should monetary policy respond optimally to reduce the welfare losses stemming from pessimism and optimism in agents' beliefs? Should the central bank adjust the interest rate more or less aggressively to dampen volatility in the variables? We answer these questions by drawing on the methods developed by Bianchi (2016) to measure welfare losses and uncertainty in MS-DSGE models. Welfare losses are measured by expected quadratic deviations of inflation and output from their respective steady states and uncertainty by the standard deviation of the variables at time $t+h$, both conditional on agents' information set I_t at time t. These measures take into account the uncertainty about the prevailing regime, the possibility of regime changes, and the evolution of agents' beliefs.

We derive the period welfare loss by applying a log-quadratic approximation of the household's utility function based on Rotemberg and Woodford (1999), Woodford (2003), and Galí (2008). Welfare losses are measured as expected quadratic deviations of inflation and detrended output from their steady-state levels. TFP growth does not enter the welfare approximation directly because it is exogenous. The detrended output deviations are weighted by the ratio of the Phillips curve slope, κ, and the elasticity of substitution between differentiated goods, ϵ:

$$\mathcal{L}_t \approx \sum_{s=0}^{\infty} \beta^s \left[\mathbb{E}_t \left(\hat{\pi}_{t+s}^2 \right) + \left(\frac{\kappa}{\epsilon} \right) \mathbb{E}_t \left(\hat{y}_{t+s}^2 \right) \right].$$

Figure 8 plots the evolution of welfare losses for the incomplete information model without the ZLB constraint, for the regime sequence discussed above. Interestingly, welfare losses decrease upon entering new regimes but increase the longer the economy remains within a given block. The same pattern emerges for uncertainty about inflation, output, and the interest rate, which is illustrated in Figure 9 at horizon h of four quarters. The rationale behind this is that the longer the economy remains in a block, the greater the adjustments once it shifts towards the other block, as visualized in Figure 6. During transitory regimes, these adjustments are smaller, which is reflected in lower uncertainty and welfare losses. This implies that the convergence of the imperfect-information model towards the perfect-information model in a given block is associated with a decline in welfare.

Figure 8: Welfare losses resulting from incomplete information

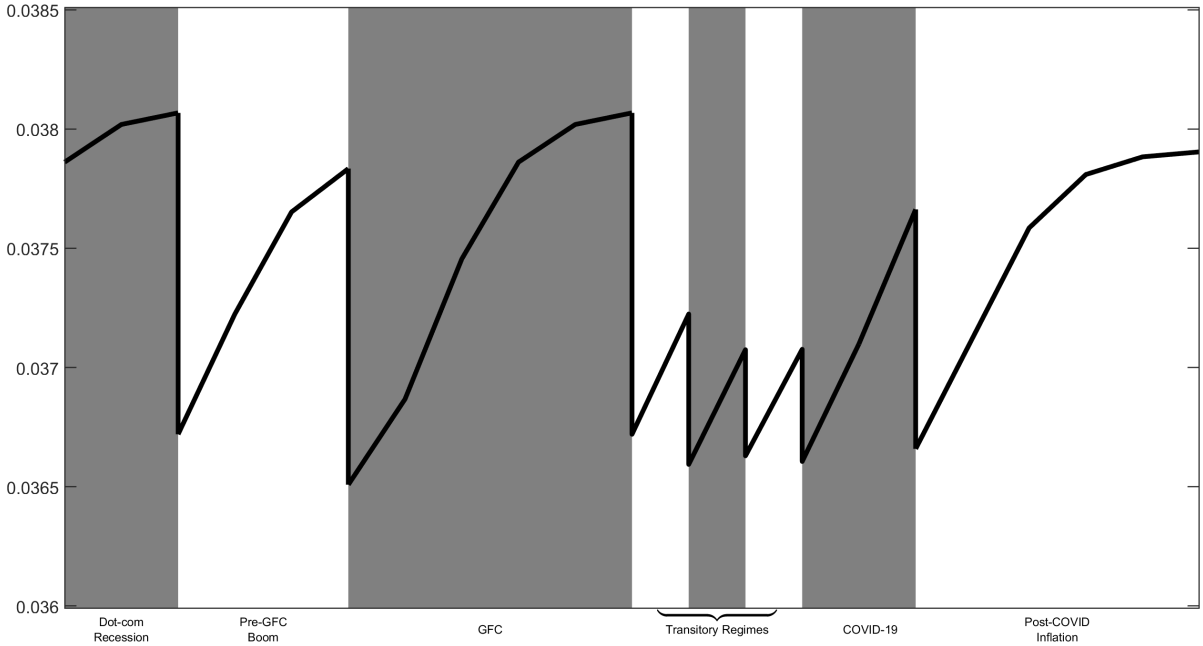

Notes: Welfare losses resulting from incomplete information. Welfare losses are measured as expected quadratic deviations of inflation and output from their steady-state levels. The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas denote episodes of high TFP growth.

In contrast, uncertainty about TFP growth declines the longer the economy remains in the same block, explaining the hump-shaped pattern of the uncertainty about output, inflation, and the interest rate. The reason is that future shifts in TFP growth become less likely once agents are sufficiently assured of being in the persistent regime, which means they assign a high probability of staying in the block, as shown in panel 5. However, entering a new block increases the uncertainty about future TFP growth as shifts become more likely, as implied by the transition matrix P. In this regard, TFP growth uncertainty peaks whenever the probability of remaining in a block in the next quarter is close to 0.5. These points are highlighted in panel 5 as blue (red) dots for the high (low) growth block, respectively.

Given that welfare losses and uncertainty are highest the longer the economy remains in a given block, we now examine how monetary policy should optimally respond to this pattern. To find the combination of ϕ_π and ϕ_y that minimizes welfare losses, we compute welfare losses over a grid of 30 equally spaced values, starting at 1.01 and ending at 3 for ϕ_π, and starting at 0 and ending at 1 for ϕ_y. Apart from the Taylor rule coefficients on inflation and output, the model and calibration remain the same, as described above. Figure 10 displays the average welfare losses over the parameter space for the simulated

Figure 9: The evolution of four-quarter-ahead uncertainty under standard and optimal monetary policy

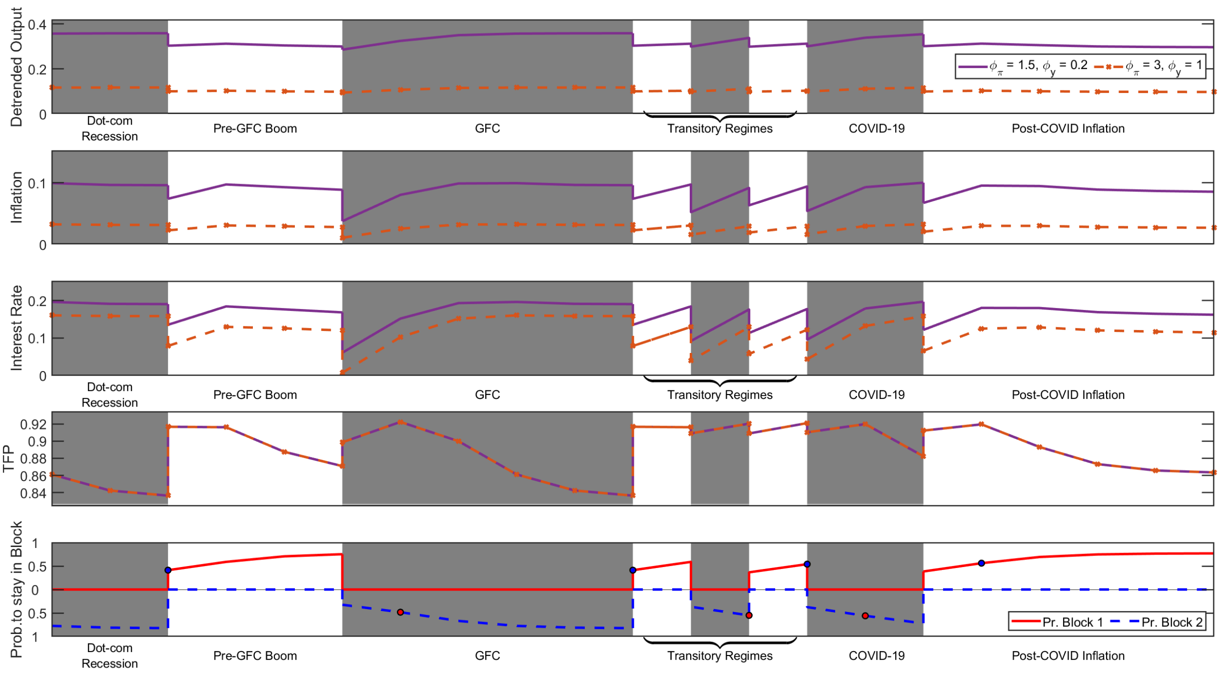

Notes: The evolution of four quarter ahead uncertainty about output, inflation, the interest rate, and TFP growth under the standard Taylor rule coefficients, $\phi_\pi = 1.5$, $\phi_y = 0.2$ (purple-solid line), and under optimal monetary policy (OMP) Taylor rule coefficients, $\phi_\pi = 3$, $\phi_y = 1$ (orange-dashed line). The last panel displays the probability of staying in Block 1 (red-solid line) and in Block 2 (blue-dashed line), respectively. The blue and red dots highlight the point at which the uncertainty about shifting blocks is at its peak. The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas reflect episodes of high TFP growth.

sequences. Notably, the most hawkish monetary policy response is optimal for minimizing welfare losses, i.e., $\phi_\pi = 3$ and $\phi_y = 1$, which is in line with the findings of Nakamura et al. (2025) and Merendino and Monacelli (2026). In our model, the hawkish monetary policy stance reduces the average welfare loss by a factor of 9.1 compared to the standard Taylor rule coefficients used for generating Figure 8, that is, $\phi_\pi = 1.5$ and $\phi_y = 0.2$.

It is worth emphasizing that the finding that higher Taylor rule coefficients reduce both uncertainty and welfare losses is not driven by the information set or the learning mechanism. The same result is obtained under perfect information, where OMP reduces welfare losses by a factor of 8.7 relative to the benchmark policy specification, albeit slightly less than under imperfect information. The similarity in welfare implications between the perfect and imperfect information models arises because the divine coincidence holds in both cases. An interesting extension would be to consider shocks that generate policy trade-offs and to compare the resulting welfare implications across the two models.[19]

[19]The divine coincidence implies that the optimal monetary policy involves Taylor rule coefficients that tend to infinity. For illustrative purposes, we impose an upper bound on these coefficients.

Figure 10: Average welfare losses over the parameter space

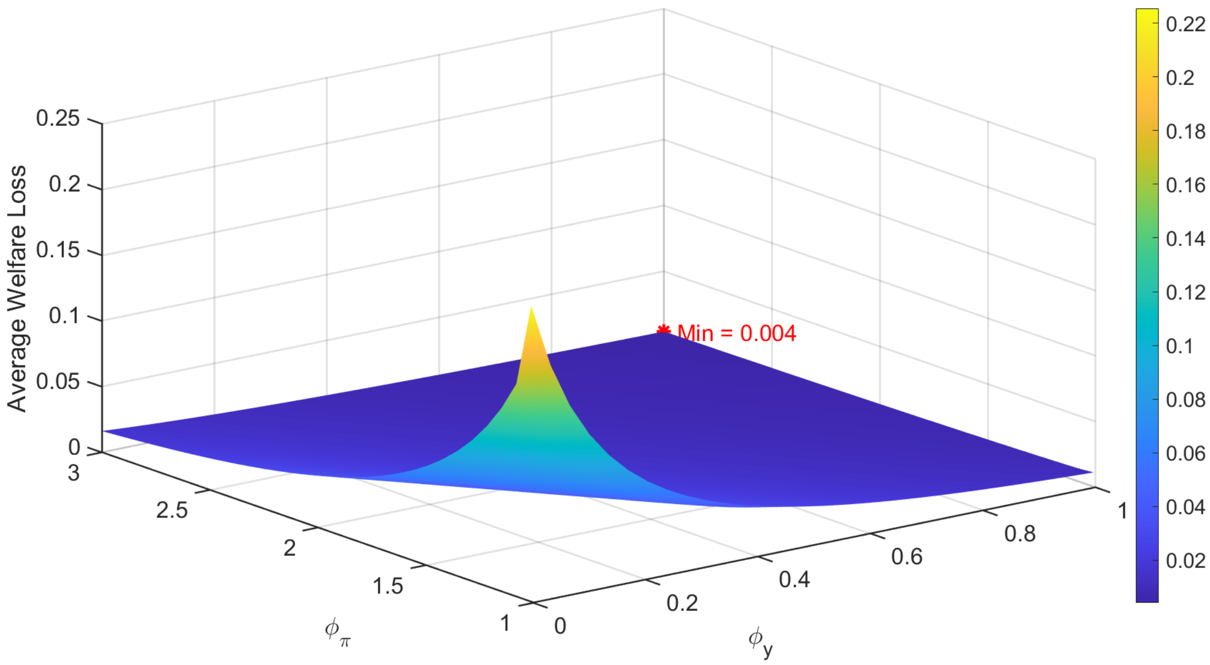

Notes: Average welfare losses over the parameter space for the simulated sequences. Welfare losses are measured as expected quadratic deviations of inflation and output from their steady-state levels. 'Min' refers to the minimum welfare loss in the parameter space.

Consistent with this, the optimal monetary policy (OMP) response leads to a reduction in the four-quarter-ahead uncertainty about output, inflation, and the interest rate, as illustrated by the difference between the solid-purple and the orange-dashed line in the first three panels of Figure 9. Nevertheless, uncertainty related to the interest rate declines somewhat less than that of output and inflation. This can be attributed to the fact that higher Taylor rule coefficients tend to induce stronger interest rate adjustments.

The same result holds when examining the actual variations in the variables over the simulated sequence for the optimal, more hawkish monetary policy (orange-dashed line) and the standard, more dovish Taylor parameters (black-solid line), see Table 1 and Figure 11. That being said, the optimal monetary policy regime is more effective at dampening fluctuations in output and inflation, and the policy rate across low- and high-growth regimes.

In line with the result for uncertainty, the difference in variation is smallest for the nominal interest rate when comparing the two monetary policy regimes. Yet, it is worth emphasizing that the higher Taylor rule coefficients do not lead to higher nominal interest rates in high-growth regimes or lower rates in low-growth regimes; instead, they result in slightly lower nominal rates. This reflects the reduced volatility of inflation expectations under the more hawkish policy stance compared to the more dovish one. Since agents anticipate that the central bank will respond more aggressively to deviations, inflation expectations remain better anchored around the steady state of the inflation rate. The stabilized inflation expectations, in turn, strengthen real interest rate adjustment without requiring aggressive increases in nominal rates, thereby reducing economic fluctuations, as reflected in Table 1.

Table 1: Standard deviations of model variables under standard and optimal monetary policy

	$\hat{y}_t$	$\hat{\pi}_t$	$\mathbb{E}_t\,\hat{\pi}_{t+1}$	$\hat{i}_t$	$\hat{i}_t - \mathbb{E}_t\,\hat{\pi}_{t+1}$
$\phi_\pi = 1.5,\ \phi_y = 0.2$	0.38	0.60	0.48	1.20	0.60
$\phi_\pi = 3,\ \ \phi_y = 1$	0.12	0.20	0.16	1.07	0.87

Notes: The table shows the standard deviations of output, inflation, expected inflation, the nominal interest rate, and the real interest rate for the simulated regime sequences for standard Taylor rule coefficients ($\phi_\pi = 1.5$, $\phi_y = 0.2$) and optimal monetary policy ($\phi_\pi = 3$, $\phi_y = 1$).

Figure 11: Model dynamics under standard and optimal monetary policy

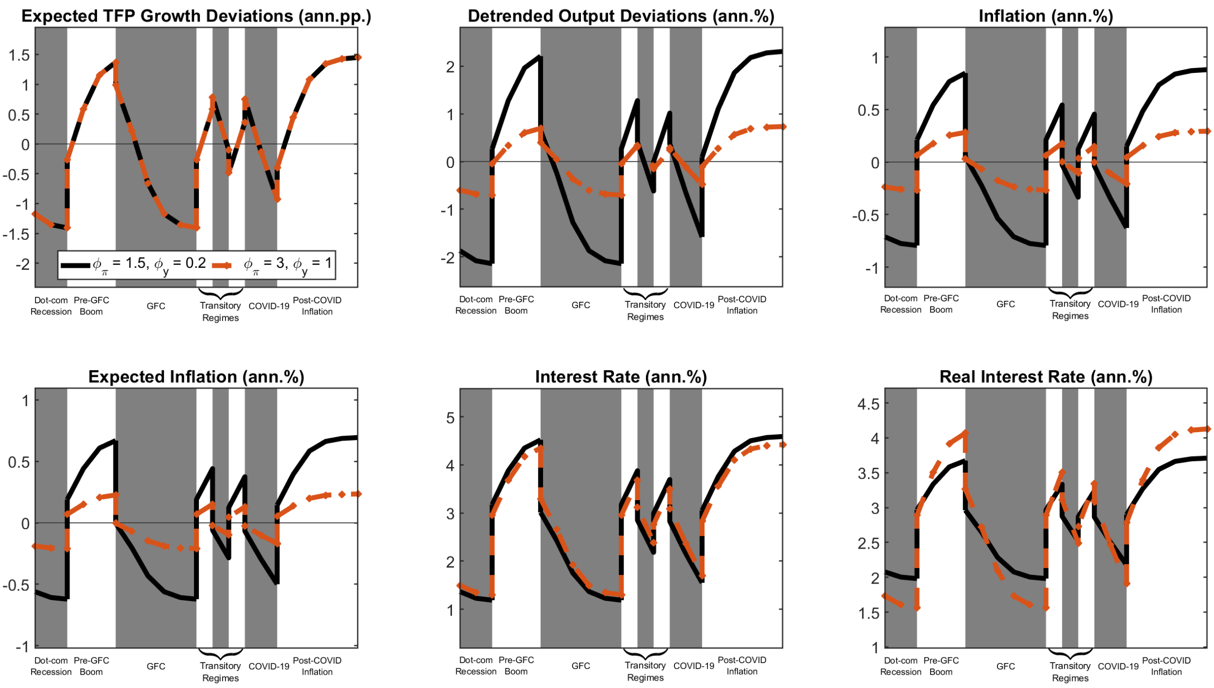

Notes: Dynamics of different central bank regimes under incomplete information. The black-solid line presents the results from monetary policy with standard Taylor rule coefficients ($\phi_\pi = 1.5$, $\phi_y = 0.2$), while the orange-dashed line represents the results under optimal monetary policy ($\phi_\pi = 3$, $\phi_y = 1$). The first panel reflects the annualized percentage-point (ann. pp.) deviations of TFP growth expectations from its ergodic mean, while the second panel displays the annualized percentage deviations of detrended output from its steady state. The third and fourth panels display the annualized inflation and expected inflation rates, while the fifth and sixth panels show the annualized nominal and real interest rates, respectively. The sequence of regime changes represents the business cycles since the beginning of the 21st century. Black areas depict episodes of low TFP growth, while white areas denote episodes of high TFP growth.

5 Conclusion

This paper contributes to the policy debate on inflation dynamics and monetary policy by introducing a New Keynesian model in which agents learn about the persistence of underlying TFP growth realizations through Bayesian learning. Departing from the standard assumption of perfect information, we show that uncertainty about the duration of TFP growth gives rise to waves of optimism and pessimism that shape inflation, output, and the interest rate dynamics in ways consistent with recent macroeconomic developments, including the Fed's delayed response following the COVID-19 crisis. In this sense, our model captures the gradual adjustment of macroeconomic variables and provides an expectation-based explanation for sluggish policy responses without resorting to ad hoc features like price indexation or habit formation. In addition, the dynamic adjustment of agents' beliefs based on past observations helps explain why inflation was initially perceived as transitory during the early post-COVID recovery.

Moreover, we show that a looking-through policy—under which the central bank delays lifting the interest rate from the ZLB until persistence becomes more certain—can result in an excessive increase in inflation and output. This underscores the critical role of timing in monetary tightening and the risks associated with waiting too long to raise rates when inflation pressures may turn out to be persistent rather than transitory.

We further demonstrate that a more hawkish monetary policy stance not only reduces welfare losses and uncertainty but also anchors expectations more effectively, reducing the need for aggressive interest rate adjustments. This underscores the stabilizing role of a stronger policy response when the persistence of TFP growth shocks is uncertain consistent with the position of Team Persistent.

Our findings emphasize the importance of accounting for uncertainty about the persistence of shocks in policy design and point to the need for further research aimed at identifying indicators to help policymakers assess the duration of macroeconomic disturbances. In this context, incorporating signals that inform agents about the persistence of the underlying regime would represent a valuable extension of our framework. Additionally, introducing belief discrepancies between policymakers and the private sector would be an interesting extension of our work, an avenue recently explored in the literature, for example, in studies by Iovino et al. (2022), Bachmann et al. (2022), Brandão-Marques et al. (2024), and Bullard et al. (2026). This would allow for a more precise analysis of imperfect information on the central bank side.

References

Amiti, M., Heise, S., Karahan, F., and Şahin, A. (2024). Inflation Strikes Back: The Role of Import Competition and the Labor Market. *NBER Macroeconomics Annual*, 38(1):71–131.

Bachmann, R., Gödl-Hanisch, I., and Sims, E. R. (2022). Identifying Monetary Policy Shocks Using the Central Bank's Information Set. *Journal of Economic Dynamics and Control*, 145(C).

Ball, L., Leigh, D., and Mishra, P. (2022). Understanding US Inflation during the COVID-19 Era. *Brookings Papers on Economic Activity*, 53(2 (Fall)):1–80.

Beckmann, J., Gern, K.-J., Hauber, P., Jannsen, N., and Stolzenburg, U. (2021). Rising Inflation: Transitory or Cause for Concern? Monetary Dialogue Papers PE 695.453, European Parliament, Policy Department for Economic, Scientific and Quality of Life Policies, Brussels.

Benigno, P. and Eggertsson, G. B. (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. Working Paper 31197, National Bureau of Economic Research.

Bergholt, D., Canova, F., Furlanetto, F., Maffei-Faccioli, N. F., and Ulvedal, P. (2025). What Drives the Recent Surge in Inflation? The Historical Decomposition Roller Coaster. *American Economic Journal: Macroeconomics (forthcoming).*

Bernanke, B. and Blanchard, O. (2025). What Caused the US Pandemic-Era Inflation? *American Economic Journal: Macroeconomics*, 17(3):1–35.

Bianchi, F. (2016). Methods for Measuring Expectations and Uncertainty in Markov-Switching Models. *Journal of Econometrics*, 190(1):79–99.

Bianchi, F., Faccini, R., and Melosi, L. (2023). A Fiscal Theory of Persistent Inflation*. *The Quarterly Journal of Economics*, 138(4):2127–2179.

Bianchi, F., Ilut, C. L., and Schneider, M. (2018). Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle. *Review of Economic Studies*, 85(2):810–854.

Bianchi, F. and Melosi, L. (2016). Modeling The Evolution of Expectations And Uncertainty In General Equilibrium. *International Economic Review*, 57(2):717–756.

Bianchi, F. and Melosi, L. (2022). Inflation as a Fiscal Limit. Working Paper Series WP 2022-37, Federal Reserve Bank of Chicago.

Blanchard, O. (2021). In Defense of Concerns over the $1.9 Trillion Relief Plan. Real Time Economics. Peterson Institute for International Economics. Available at: `https://www.piie.com/blogs/realtime-economics/defense-concerns-over-19-trillion-relief-plan`.

Brandão-Marques, M. L., Meeks, M. R., and Nguyen, V. (2024). Monetary Policy with Uncertain Inflation Persistence. IMF Working Papers 2024/047, International Monetary Fund.

Bullard, J., Grimaud, A., Salle, I., and Vermandel, G. (2026). Soft Landing and Inflation Scares. *Journal of Monetary Economics*, 157:103871.

Buss, G. and Traficante, G. (2025). The Return of Inflation: Look-Through Policy Under Incomplete Information. Working Papers 2025/02, Latvijas Banka.

Calvo, G. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics*, 12(3):383–398.

Chetty, R. (2012). Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, 80(3):969–1018.

Cho, S. (2016). Sufficient Conditions for Determinacy in a Class of Markov-Switching Rational Expectations Models. *Review of Economic Dynamics*, 21:182–200.

Christiano, L., Eichenbaum, M., and Rebelo, S. (2011). When Is the Government Spending Multiplier Large? *Journal of Political Economy*, 119(1):78–121.

Dao, M. C., Gourinchas, P.-O., Leigh, D., and Mishra, P. (2024). Understanding the International Rise and Fall of Inflation since 2020. *Journal of Monetary Economics*, 148:103658. Inflation in the COVID Era and Beyond.

Davig, T. and Leeper, E. M. (2007). Generalizing the Taylor Principle. *American Economic Review*, 97(3):607–635.

Erceg, C. J., Lindé, J., and Trabandt, M. (2024). Monetary Policy and Inflation Scares. IMF Working Papers 2024/260, International Monetary Fund.

Farmer, R. E., Waggoner, D. F., and Zha, T. (2009). Understanding Markov-Switching Rational Expectations Models. *Journal of Economic Theory*, 144(5):1849–1867.

Fernald, J. G. (2012). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Federal Reserve Bank of San Francisco Working Paper Series, 2012-19, Updated April 2014.

Fernández-Villaverde, J., Gordon, G., Guerrón-Quintana, P., and Rubio-Ramírez, J. F. (2015). Nonlinear Adventures at the Zero Lower Bound. *Journal of Economic Dynamics and Control*, 57(C):182–204.

Foerster, A., Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2016). Perturbation Methods for Markov-Switching Dynamic Stochastic General Equilibrium Models. *Quantitative Economics*, 7(2):637–669.

Fuhrer, J. (2017). Expectations as a Source of Macroeconomic Persistence: Evidence from Survey Expectations in a Dynamic Macro Model. *Journal of Monetary Economics*, 86(C):22–35.

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

Harding, M., Lindé, J., and Trabandt, M. (2023). Understanding post-COVID Inflation Dynamics. *Journal of Monetary Economics*, 140(S):101–118.

Ihrig, J. E. and Waller, C. J. (2024). The Federal Reserve's Responses to the Post-Covid Period of High Inflation. FEDS Notes 2024-02-14, Board of Governors of the Federal Reserve System (U.S.).

Iovino, L., La'O, J., and Mascarenhas, R. (2022). Optimal Monetary Policy and Disclosure with an Informationally-Constrained Central Banker. *Journal of Monetary Economics*, 125(C):151–172.

Kiley, M. T. (2024). Anchored or Not: How Much Information Does 21st Century Data Contain on Inflation Dynamics? *International Journal of Central Banking*, 20(1).

Kose, M. A., Ohnsorge, F., and Ha, J. (2021). Inflationary Pressures: Likely Temporary but Challenging for Policy Design. CEPR VoxEU Column, July 14, `https://cepr.org/voxeu/columns/inflationary-pressures-likely-temporary-challenging-policy-design`.

Krugman, P. (2021). Wonking Out: I'm Still on Team Transitory. The New York Times. `https://www.nytimes.com/2021/09/10/opinion/transitory-inflation-covid-consumer-prices.html`.

Liu, Z., Waggoner, D. F., and Zha, T. (2011). Sources of Macroeconomic Fluctuations: A Regime-Switching DSGE Approach. *Quantitative Economics*, 2(2):251–301.

Londono, J. M., Ma, S., and Wilson, B. A. (2023). Global Inflation Uncertainty and its Economic Effects. FEDS Notes 2023-09-25, Board of Governors of the Federal Reserve System (U.S.).

Lorenzoni, G. and Werning, I. (2023). Inflation is Conflict. Working Paper 31099, National Bureau of Economic Research.

Merendino, A. and Monacelli, T. (2026). Supply Chain Uncertainty, Energy Prices, and Inflation. Unpublished manuscript, February 2026.

Milani, F. (2007). Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54(7):2065–2082.

Nakamura, E., Riblier, V., and Steinsson, J. (2025). Beyond the Taylor Rule. NBER Working Papers 34200, National Bureau of Economic Research, Inc.

Pfauti, O. (2023). The Inflation Attention Threshold and Inflation Surges. Papers 2308.09480, arXiv.org.

Powell, J. H. (2021). Monetary Policy in the Time of COVID. Speech at the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, August 27, `https://www.federalreserve.gov/newsevents/speech/powell20210827a.htm`.

Reis, R. (2026). Why Did Inflation Rise and Fall in 2021–24? Channels and Evidence from Expectations. CEPR Discussion Paper 21277, CEPR, Paris and London.

Rotemberg, J. and Woodford, M. (1999). The Cyclical Behavior of Prices and Costs. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, Part B, chapter 16, pages 1051–1135. Elsevier, 1 edition.

Schneeweiss, Z. (2021). ECB's Lagarde Says Inflation Spike Is 'Largely Transitory'. Bloomberg, October 16, `https://www.bloomberg.com/news/articles/2021-10-16/ecb-policy-will-continue-supporting-the-economy-lagarde-says`.

Schorfheide, F. (2005). Learning and Monetary Policy Shifts. *Review of Economic Dynamics*, 8(2):392–419. Monetary Policy and Learning.

Shapiro, A. H. (2026). Decomposing Supply- and Demand-Driven Inflation. *Journal of Money, Credit and Banking*, 58(2):365–388.

Sims, C. (2002). Solving Linear Rational Expectations Models. *Computational Economics*, 20(1-2):1–20.

Summers, L. H. (2021). The Inflation Risk Is Real. Opinion, The Washington Post, May 24, https://www.washingtonpost.com/opinions/2021/05/24/inflation-risk-is-real/.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., Kenny, G., and Ponce, J. (2025). Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings. *Econometrica*, 93(1):229–264.

Williams, N. and Sargent, T. J. (2026). Rationalizing Fed Interest Rate Decisions in the 2020s. Unpublished manuscript, January 2026.

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

A Appendix

A.1 Convergence of beliefs

Bianchi and Melosi (2016) prove that three assumptions on the transition matrix P ensure that the beliefs within the blocks converge. That is,

- Nontriviality assumption: $p_{11} + p_{12} + p_{21} + p_{22} \neq 0$ & $p_{33} + p_{34} + p_{43} + p_{44} \neq 0$, which ensure that regimes within blocks are relevant for agent
- Nontrivial learning assumption: Either $p_{11} \neq p_{22}$ or $p_{12} \neq p_{21}$ and either $p_{33} \neq p_{44}$ or $p_{34} \neq p_{43}$, which ensures that agents need to learn about the regimes
- Nonjumping assumption: $p_{11} + p_{22} > 0$ and $p_{33} + p_{44} > 0$, which ensures that in each block at least one regime can last more than one period

By imposing these assumptions, Bianchi and Melosi (2016) show that for any $\epsilon > 0$ there exists a τ_1^* and τ_2^* such that

$$prob(\xi_t^Z = 1|\tau_1^*) - prob(\xi_t^Z = 1|\tau_1^* + 1) < \epsilon,$$

$$prob(\xi_t^Z = 3|\tau_2^*) - prob(\xi_t^Z = 3|\tau_2^* + 1) < \epsilon.$$

This means that when beliefs converge at τ_1^* and τ_2^*, the additional gain from learning in the next period becomes negligibly small.

Moreover, Bianchi and Melosi (2016) show that the convergence points for $prob(\xi_t^Z = 1|\tau_1^*)$ and $prob(\xi_t^Z = 3|\tau_2^*)$, denoted as $\tilde{\lambda}_{b1}$ and $\tilde{\lambda}_{b2}$, are the stable characteristic roots of equation (10) and equation (11), respectively, which are

$$\tilde{\lambda}_{b1} = \frac{p_{11} - p_{22} + \sqrt{(p_{11} - p_{22})^2}}{2(p_{11} - p_{22})}, \tag{A.1}$$

$$\tilde{\lambda}_{b2} = \frac{p_{33} - p_{44} + \sqrt{(p_{33} - p_{44})^2}}{2(p_{33} - p_{44})}. \tag{A.2}$$

As mentioned above, the definition of the regimes implies that the persistence probability of the persistent regimes is higher than the persistence probability of the transitory regimes, that is, $p_{22} > p_{11}$ and $p_{44} > p_{33}$. With these assumptions, it follows that the probability of being in the transitory regime decreases the longer the economy stays within a block according to equations (10) and (11) and eventually converges towards 0 as implied by equations (A.1) and (A.2). In other words, after a certain amount of consecutive periods in a block, agents become certain to be in the persistent regime, that is, $prob(\xi_t^Z = 2|\tau_1^*) \equiv 1 - prob(\xi_t^Z = 1|\tau_1^*) = 1$ and $prob(\xi_t^Z = 4|\tau_2^* + 1) \equiv 1 - prob(\xi_t^Z = 3|\tau_2^* + 1) = 1$. This implies that beliefs under incomplete information converge towards those under perfect information if the economy stays sufficiently long in a block.

A.2 Approximating agents' beliefs and pinning down $\hat{P}$

To overcome the infinite dimensionality problem, we follow Bianchi and Melosi (2016) and set up a finite grid to approximate agents' beliefs. The grid G consists of two subgrids, one for each block. That is, $G_{b1} = \{G_1, ..., G_{g_1}\}$ contains the beliefs for $prob\{\xi_t^Z = 1|I_t\}$ and $G_{b2} = \{G_{g_1+1}, ..., G_{g_1+g_2}\}$ contains the beliefs for $prob\{\xi_t^Z = 3|I_t\}$, where $0 \leq G_i \leq 1$ for all $1 \leq i \leq g = g_1 + g_2$. We define a new set of regimes $\zeta_t^Z \in \{1, ...g_1 + g_2\}$ as the knots in the grid that best approximate agents' beliefs at each point in time. These new regimes evolve according to the transition matrix $\hat{P}$. We use ζ_t^Z to recast the model under imperfect information (see Section 2.10). To pin down $\hat{P}$, we apply the algorithm developed by Bianchi and Melosi (2016), which is briefly described as follows.

The algorithm consists of two main steps, which are identical for both blocks. In the following, we only describe the procedure for Block 1. In the first step, the transition matrix for the new regimes ζ_t^Z is set to $\hat{P} = 0_{g\times g}$. For each grid point G_i, the probability of staying in the block is calculated. That is, for any grid point $G_i \in G_{b1}$, $1 \leq i \leq g_1$, the probability to stay in Block 1 is given by

$$\hat{P}(i,j) = prob\{\xi_{t-1}^Z = 1|I_{t-1}\}p_{11} + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})p_{22}, \tag{A.3}$$

where $prob\{\xi_{t-1}^Z = 1|I_{t-1}\} = G_i$ and $j \leq g_1$ is the regime that approximates agents' beliefs best if the system stays in Block 1. That is, j is chosen such as to $\min |prob\{\xi_t^Z = 1|I_t\} - G_j|$. Therefore, $prob\{\xi_t^Z = 1|I_t\}$ is calculated using the recursive equation (10) by approximating $prob\{\xi_{t-1}^Z = 1|I_{t-1}\} = G_i$. In other words, j is the regime that applies if the economy remains in Block 1 in the next period. The algorithm assigns the value to element (i, j) in $\hat{P}$ given the expression in equation (A.3).

In the second step, for each grid point G_i, the probability of shifting toward the other block is calculated. For example, for Block 1 this given by $\hat{P}(i,l) = 1 - \hat{P}(i,j)$. Here, $l > g_1$ denotes the grid point that best approximates agents' beliefs if the system switches to Block 2 in the next period. That is the regime the economy switches to if the block changes. In this regard, l satisfies

$$\min \left| \frac{prob\{\xi_{t-1}^Z = 1|I_{t-1}\}p_{13} + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})p_{23}}{prob\{\xi_{t-1}^Z = 1|I_{t-1}\}(p_{13} + p_{14}) + (1 - prob\{\xi_{t-1}^Z = 1|I_{t-1}\})(p_{23} + p_{24})} - G_l \right|,$$

with $G_i = prob\{\xi^Z_{t-1} = 1|I_{t-1}\}$.

In what follows, we set up a grid with 100 equally spaced knots for each block to approximate the beliefs $prob\{\xi^Z_t = 1|I_t\}$ and $prob\{\xi^Z_t = 3|I_t\}$. Moreover, to improve the speed and accuracy of the approximation, we make the grid finer around the convergence points $\tilde{\lambda}_{b1}$ and $\tilde{\lambda}_{b2}$ by adding 185 knots to the grid. After removing unnecessary regimes, we obtain 52 grid points for Block 1 and 64 grid points for Block 2.[20]

B Model derivation and steady-state analysis

B.1 Households

Without population growth, labor does not grow along a balanced growth path. Consequently, output grows at the same rate as technology

$$(\ln Y_t - \ln Y_{t-1}) = (\ln Z_t - \ln Z_{t-1}),$$

which implies that

$$\frac{Y_t}{Z_t} = \frac{Y_{t-1}}{Z_{t-1}}.$$

Thus, along the balanced growth path, output divided by Z_t is stationary. Consequently, to make the model stationary, we divide the remaining variables by Z_t. We denote the stationary variables as $\tilde{Y}_t \equiv \frac{Y_t}{Z_t}$, $\tilde{W}_t \equiv \frac{W_t}{Z_t}$, $\mu_t \equiv ln(\frac{Z_t}{Z_{t-1}})$, and $M_t \equiv \frac{Z_t}{Z_{t-1}}$. The latter denotes the gross growth rate of TFP.

We divide (3) and (4) by Z_t to obtain the stationary versions of the new IS curve and labor-leisure equation

$$\frac{\tilde{W}_t}{P_t} = N_t^{\phi}\tilde{Y}_t. \tag{B.1}$$

$$\tilde{Y}_t^{-1} = \beta\,\mathbb{E}_t\,\tilde{Y}_{t+1}^{-1}M_{t+1}^{-1}\frac{R_t}{\Pi_{t+1}}. \tag{B.2}$$

Lowercase variables with hats denote the percent deviations of the stationary variables from their steady-state levels, which are given by

$$\hat{w}_t - \hat{P}_t = \phi\hat{n}_t + \hat{y}_t,$$

[20]Apart from steps 1 and 2, the algorithm consists of steps of disposing regimes that will never be reached, see Bianchi and Melosi (2016) for details.

$$\hat{y}_t = \mathbb{E}_t \, \hat{y}_{t+1} + \mathbb{E}_t \, \hat{\mu}_{t+1} - (\hat{r}_t - \mathbb{E}_t \, \hat{\pi}_{t+1}),$$

where $\hat{\mu}_{t+1} = \mu_{t+1} - \bar{\mu}$ is the log-deviation of the growth rate of TFP from its ergodic mean $\bar{\mu}$ as defined below.

B.2 Firms

Detrending the production function (5) and demand function for intermediary goods (6) yields

$$\tilde{Y}_t(i) = \left(\frac{P_t(i)}{P_t}\right)^{-\epsilon} \tilde{Y}_t. \tag{B.3}$$

The detrended production function is given by

$$\tilde{Y}_t(i) = N_t(i). \tag{B.4}$$

We assume that firms charge a constant markup $\mu = \frac{\epsilon}{\epsilon - 1}$ over their stationary nominal marginal costs. That is,

$$P_t(i) = \mu \tilde{W}_t. \tag{B.5}$$

The detrended firm profits at period t read

$$\tilde{\Omega}_t = P_t(i)\tilde{Y}_t(i) - \tilde{W}_t N_t(i).$$

Since the firms are owned by the households, future profits are discounted with the stochastic discount factor $Q_{t,t+k}$, which is given by the household's maximization problem

$$Q_{t,t+k} = \beta^k \frac{Y_t}{Y_{t+k}} \frac{P_t}{P_{t+k}}.$$

The detrended discount factor is

$$\tilde{Q}_{t,t+k} = \beta^k \frac{\tilde{Y}_t}{\tilde{Y}_{t+k}} \frac{Z_t}{Z_{t+k}} \frac{P_t}{P_{t+k}}.$$

When allowed to optimize its price $P_t^*(i)$ the firm maximizes:

$$\mathbb{E}_t \sum_{k=0}^{\infty} \theta^k \tilde{Q}_{t,t+k} \{P_t^*(i)\tilde{Y}_{t+k}(i) - \tilde{W}_{t+k} N_{t+k}(i)\}$$

subject to the stationary demand function (B.3) and production function (B.4).

Solving the optimization problem, log-linearizing the variables around their steady states, and rearranging results in the New Keynesian Phillips curve

$$\hat{\pi}_t = \beta\,\mathbb{E}_t\,\hat{\pi}_{t+1} + \frac{(1-\beta\theta)(1-\theta)(1+\phi)}{\theta}\hat{y}_t.$$

B.3 TFP Growth Rate

As stressed above, we express the TFP process (7) in growth rates to induce stationarity. That is,

$$\mu_t = \mu(\xi_t^Z).$$

Log-linearizing the TFP growth rate yields

$$\hat{\mu}_t = \hat{\mu}(\xi_t^Z),$$

with $\hat{\mu}_t$ and $\hat{\mu}(\xi_t^Z)$ denoting the percentage-point deviations of TFP growth and the TFP drift from their ergodic mean $\bar{\mu}$, respectively.

B.4 Monetary Policy

Log-linearizing equation (14) around the variables' steady states results in

$$\hat{r}_t = (1 - D_{\xi_t^{MP}})(\phi_\pi\hat{\pi}_t + \phi_y\hat{y}_t) - D_{\xi_t^{MP}}\log R_{ss}.$$

B.5 Steady States and Ergodic Mean

The ergodic mean of the TFP growth rate is given by

$$\bar{\mu} = (\bar{p}_1 + \bar{p}_2)\mu_h + (\bar{p}_3 + \bar{p}_4)\mu_l,$$

where $\bar{p}_i$ reflects the ergodic probability to be in regime i.

We fix the steady-state labor input N_{ss} to 0.3, and the price level P_{ss} and the inflation rate π_{ss} to 1. According to equation (B.5), the steady-state wage level is given by the inverse of the markup, that is

$$W_{ss} = \frac{\epsilon - 1}{\epsilon}.$$

Given N_{ss} and W_{ss}, we can then pin down the steady-state value of output according to the labor-leisure equation (B.1)

$$Y_{ss} = \frac{W_{ss}}{N_{ss}^{\phi}}.$$

The steady-state interest rate R can be pinned down by the New IS curve (B.2)

$$R_{ss} = \frac{\bar{M}}{\beta},$$

where $\bar{M} = exp(\bar{\mu})$.