Monetary Policy without an Anchor*

Luigi Bocola Alessandro Dovis Kasper Jørgensen
Rishabh Kirpalani

October 2025

Abstract

Policymakers often cite the risk that inflation expectations might "de-anchor" as a key reason for responding forcefully to inflationary shocks. We develop a model to analyze this trade-off and to quantify the benefits of stable long-run inflation expectations. In our framework, households and firms are imperfectly informed about the central bank's objective and learn from its policy choices. Recognizing this interaction, the central bank raises interest rates more aggressively after adverse supply shocks and accepts short-run output costs to secure more stable inflation expectations. The strength of this reputation channel depends on how sensitive long-run inflation expectations are to surprises in interest rates. Using high-frequency identification, we estimate these elasticities for emerging and advanced economies and find large negative values for Brazil. We fit our model to these findings and use it to quantify how reputation building motives affect monetary policy decisions, and the role of central bank's credibility in promoting macroeconomic stability.

Keywords: De-anchoring of inflation expectations, reputation, optimal policy

*First draft: February 2025. Bocola: Stanford University and NBER. Dovis: University of Pennsylvania and NBER. Jørgensen: European Central Bank. Kirpalani: University of Wisconsin. We thank Luca Fornaro, Jordi Gali, Patrick Kehoe, Jon Steinsson, and seminar participants at "Frontiers in Macroeconomics", Stanford, MIT, Federal Reserve Bank of Philadelphia, Central Bank of Chile, Sciences Po Summer Workshops in Economics, the MadMac 2025 conference, BSE Summer Forum, the SED meetings, and Yale for helpful comments. We thank Jens Christensen for kindly providing us some of the data used in this study. Tomas Tapak provided excellent research assistance. The views and opinions in this paper are those of the authors and do not necessarily reflect those of the European Central Bank or its staff. All errors are our own. Luigi Bocola thanks the National Science Foundation for generous funding.

1 Introduction

Large inflation spikes—such as those that followed the pandemic—can reshape how households and firms perceive the central bank's commitment to price stability. If that commitment is questioned, inflation expectations can rise and lead firms to adjust prices preemptively, amplifying inflationary pressures. Policymakers therefore view the risk of a de-anchoring in long-run inflation expectations as a central reason for strong actions. The experience of Brazil after the pandemic illustrates the point. In March 2021 its central bank began raising interest rates—more than a year ahead of similar moves in the United States or the euro area—explicitly mentioning concerns about a de-anchoring of inflation expectations.[1] How much did such concerns shape Brazil's policy response? How much "reputation" did the central bank build by acting swiftly, and why, by contrast, were the United States and the euro area able to delay tightening monetary policy?

We address these questions by analyzing monetary policy in an economy in which households and firms are imperfectly informed about the central bank's objectives. The central bank sets policy without commitment, and private agents infer whether it is "hawkish" or "dovish" from its actions. When a policymaker is perceived as hawkish—placing greater weight on inflation stabilization—long-run inflation expectations are lower and less sensitive to economic shocks. Recognizing this, the central bank has an incentive to keep real interest rates higher on average and to raise them more aggressively in response to adverse supply shocks, accepting short run economic costs to strengthen its reputation. We develop an empirical strategy to measure the strength of these reputation forces and show that they can be a key driving force for monetary policy.

We consider a standard New Keynesian framework. Firms' optimal pricing decisions deliver a canonical Phillips curve; households' consumption satisfies an Euler equation. The monetary authority sets nominal interest rates under discretion to minimize a loss function that weights deviation of output and inflation from their targets. The central bank can be one of two types—a Hawk or a Dove—with the former placing greater emphasis on inflation stabilization. Private agents do not observe the central bank's type and they update beliefs from observed policy choices using Bayes' rule. The economy is subject to three types of shocks: supply shocks (which shift the Phillips curve), productivity shocks, and monetary shocks (transitory disturbances to the relative preferences over inflation and the output gap).

[1] For example, the minutes of the March 2021 meeting of the Brazil's Monetary Policy Committee (Copom) report that "several members also emphasized that the inflationary pressures observed in 2021 could spill over into inflation expectations for 2022, creating a risk of de-anchoring of expectations over the monetary policy-relevant horizon." See https://www.bcb.gov.br/content/copom/atascopom/Copom237-not20210317237.pdf.

Under perfect information, optimal monetary policy in a Markov equilibrium reduces to a static trade-off between inflation and output stabilization (Clarida, Gali, and Gertler, 1999; Woodford, 2003; Galí, 2015). This trade-off is only relevant in response to supply shocks, as productivity shocks can be perfectly stabilized. When a positive supply shock occurs, the central bank raises interest rates to reduce output and offset the inflationary impact. The strength of this response depends on the central bank's preferences and on the slope of the Phillips curve, which governs the sensitivity of inflation to the output gap.

Under imperfect information about the central bank's type, monetary policy affects inflation both directly—via the slope of the Phillips curve—and indirectly, through inflation expectations. Because the Dove delivers a higher and more volatile inflation (reflecting classic inflation and stabilization bias),[2] gaining a reputation for being a Hawk lowers expected inflation and reduces its sensitivity to shocks, thereby helping stabilize current inflation. These gains are traded off against the short-run output costs of building reputation. Specifically, the central bank needs to keep real interest rates higher on average and respond more aggressively to inflationary supply shocks, in order to signal a strong commitment to inflation stabilization.

We leverage this trade-off to design a measurement strategy. Specifically, we show that the strength of the reputation channel depends on how sensitive inflation expectations are to surprises in interest rates. When this elasticity is large, deviating from the static trade-off is more valuable because it buys the policymaker more anchoring of inflation expectations. Using high frequency identification around policy meetings, we estimate these elasticities for emerging and advanced economies. We find large negative elasticities for Brazil: a 100-basis-point positive surprise in the Selic rate is associated with a 0.40-percentage-point decline in five-year, five-year forward market-based inflation compensation and a 0.18-percentage-point decline in survey-based expectations four years ahead. In other economies in our sample, long-run expectations appear more firmly anchored, with elasticities statistically indistinguishable from zero at five-year and longer horizons.

We fit the model to Brazil, choosing parameters to match both the measured elasticity as well as the unconditional behavior of output, inflation and nominal interest rates. Consistent with the data, long-run inflation expectations can be very sensitive to policy surprises, especially in periods in which the private sector is uncertain about the policymaker's type.[3] This property is the result of two key ingredients: a "slow" learning process, so that inno-

[2]The former arises in our model when the output target of the central bank is above the flexible-price level of output, while the latter reflects the central bank's imperfect ability to offset the effects of economic shocks on output and inflation.

[3]Because of Bayes' rule, the posterior distribution over the policymaker's type is less sensitive to new information when the private sector holds a dogmatic prior—that is, when agents are quite sure of facing the hawkish or dovish policymaker.

vations in the beliefs of the private sector persist in the long run; and an inflation bias, so the Hawk and the Dove policymaker implement different inflation levels. With these two ingredients, the model can match the observed elasticity.

We finally use the parametrized model to quantify the importance of the reputation channel. Because of the strong sensitivity of long-run inflation expectations, the central bank with low reputation tilts policy aggressively toward inflation stabilization in order to influence beliefs, accepting larger output costs in the short run. We show that these effects are sizable: to replicate the same static trade-off between output and inflation in an economy without reputation building motives, one would need to increase the slope of the Phillips curve by a factor of 4 times (from 0.075 to 0.306) on average and 9 (from 0.075 to 0.696) in periods of high uncertainty.

As reputation rises, inflation expectations fall and become more firmly anchored, something that allows the central bank to respond less forcefully to shocks while still maintaining price stability. We use the model to quantify this "reputation dividend" and to disentangle its sources. When the credibility of the central bank is low, gains in reputation lead to a drastic fall in the inflation bias at the cost of a larger stabilization bias—as the central bank responds forcefully to supply shocks in an attempt to build reputation. Eventually, as the central bank acquires credibility, the reputation building motives become less relevant. The central bank in those situations is able to stabilize inflation with smaller movements in the real interest rates.

Recent work by Nakamura, Riblier, and Steinsson (2025) shows that the ability of central banks to "look through" the Covid supply shocks correlate with measures of credibility, with central bank with stronger credibility tightening less aggressively than those with weaker credibility. Our framework provides a theory of this phenomenon and complementary evidence by quantifying the reputational channel that links credibility to the anchoring of expectations. Consistent with their findings, our model implies that central banks with low reputation respond aggressively to stabilize inflation, whereas those with higher credibility can look through shocks because expectations are more firmly anchored.

Related literature. Our model contributes to a large literature that studies reputation building and management of private sector expectations in monetary economies. We build on the work of Barro (1986), Backus and Driffill (1985), Lu, King, and Pasten (2016), Kostadinov and Roldán (2024), de Aguilar (2024), and Amador and Phelan (2024). In all these policy games, as in the classic work by Rogoff (1985), having a reputation of being a Hawk is beneficial for the central bank because it can either reduce the *inflation bias* coming from the incentive to increase output above the equilibrium level as in Kydland and Prescott (1977)

and Barro and Gordon (1983b) or allow for a better trade-off between inflation and output in response to supply shocks, the *stabilization bias*.[4] In our model, both of these margins are present.

Aside from technical differences, the main contribution of our paper relative to this literature is to show how data can be used to discipline the role of reputation in shaping monetary policy decisions. This emphasis on the measurement is shared also by King and Lu (2022) and the contemporaneous work of Caravello, Martinez-Bruera, and Carrasco (2025). King and Lu (2022) use variation in the term structure of inflation expectations to infer central bank reputation. Our approach departs from theirs by using high-frequency methods to study how inflation expectations at different horizons respond to surprises in interest rates.[5] As we show in the paper, this statistic is directly related to the incentives to build reputation in our dynamic economy and, as such, it provides critical information for quantifying the strength of these forces.

Caravello et al. (2025) study optimal monetary policy in an environment that shares some of the features present in our model. They develop a measure of central bank reputation based on the slope in a cross-sectional regression of output-gap and private sector inflation forecasts—a mapping derived under the assumption that private agents believe that the central bank acts myopically. By contrast, we rely on the sensitivity of long-run inflation expectations to monetary surprises, which provides a sufficient statistic for reputational considerations in our model where the private sector correctly perceives that the central bank chooses monetary policy strategically to manage inflation expectations.[6]

A number of studies have focused on the role of imperfect information and learning in shaping the private sector response to monetary policy changes. Erceg and Levin (2003), for example, emphasize the importance of these features in generating inflation persistence and high output costs during the Volcker disinflation. Bauer, Pflueger, and Sunderam (2024) consider a learning model where the private sector updates its view about the central bank's reaction function over time by looking at realized interest rates. Closer to our application, Bonomo, Carvalho, Eusepi, Perrupato, Abib, Ayres, and Matos (2024) study

[4] Afrouzi, Halac, Rogoff, and Yared (2023) study the Markov Perfect Equilibrium in a New Keynesian model with a dynamic Phillips curve and full information and show that it delivers an inflation bias because of the presence of markups that are not undone by wage subsidies (as typically assumed in the New-Keynesian literature).

[5] Morelli and Moretti (2023) also use a high-frequency identification strategy to quantify the contribution of reputation in the context of a sovereign default model.

[6] Relatedly, Christoffel and Farkas (2025) develop a two-regime model in which long-run inflation expectations can diverge from the central bank's target. When the probability of entering this regime is set as a function of the policy rate, the central bank has incentives to act more hawkish as in our model. In an application to the euro area with an exogenous policy rule, they find little evidence of de-anchoring during the post-Covid inflation surge but evidence of below-target expectations in the early 2010s

the impact that observed monetary policy changes in Brazil had on inflation expectations during the 2010s. In these contributions, the central bank follows an interest rate rule while private agents infer features of the rule (e.g. the inflation target) through signal extraction. By contrast, we study an optimizing central bank that internalizes how its policy choices shape the private sector's learning problem, and we focus on quantifying how these aspects shape the optimal policy.

Our paper is also related to a literature that measures the degree of de-anchoring of inflation expectations. Coibion and Gorodnichenko (2025) provide an overview of these methods. One class of measures relies on the level of long-run expectations—for example, their distance from the central bank's target or the degree of disagreement across forecasters (see Kumar, Afrouzi, Coibion, and Gorodnichenko 2015; Coibion, Gorodnichenko, and Kumar 2018; Bems, Caselli, Grigoli, and Gruss 2021). More recent work builds on the idea, pioneered by Bernanke (2007), of using the sensitivity of long-run inflation expectations to short-run news as an indicator of de-anchoring. Carvalho, Eusepi, Moench, and Preston (2023) and Gati (2023), for instance, embed adaptive learning mechanisms into New Keynesian models and treat expectations as anchored when long-run beliefs are relatively insensitive to short-run forecast errors. Our theory provides a framework for interpreting these measures. In our model, long-run inflation expectations are more at risk of de-anchoring when they are highly sensitive to policy surprises, and this greater sensitivity strengthens the incentive for the central bank to stabilize inflation in the face of shocks. Thus, through the lens of our model, measures based on the sensitivity of long-run expectations to short-run policy news are more informative about de-anchoring risk than those based solely on the level of long-run expectations.[7]

The empirical strategy in the paper builds on work using high-frequency identification to measure how monetary policy shocks affect interest rates and inflation expectations (see Cook and Hahn (1989), Kuttner (2001) and Cochrane and Piazzesi (2002), among many others). Evidence for long-horizon expectations is mixed. Before TIPS data, Gürkaynak, Sack, and Swanson (2005) viewed significant moves in long-term nominal forwards as shifts in the steady-state inflation rate, though subsequent authors pointed out that these may also reflect risk-premium changes (Hanson and Stein, 2015; Kekre, Lenel, and Mainardi, 2024). Using early TIPS data, Beechey and Wright (2009) find significant effects on far-forward

[7]For example, in Brazil long-run inflation expectations rose only modestly during the pandemic, broadly in line with developments in other emerging markets. Nevertheless, the Brazilian central bank perceived a high risk of de-anchoring and responded by raising interest rates aggressively in early 2021. Our model helps reconcile this apparent inconsistency: when long-run expectations are highly sensitive to policy changes, the risk of de-anchoring is elevated, giving the central bank a strong incentive to act swiftly. By tightening aggressively, the Brazilian central bank likely prevented expectations from drifting further, thereby keeping long-run inflation expectations anchored through the pandemic.

inflation compensation, whereas Gürkaynak, Levin, and Swanson (2010) and Nakamura and Steinsson (2018) find more limited evidence. Our estimates for the U.S. are more in line with these latter results.

Evidence for emerging markets is scarcer. De Pooter, Robitaille, Walker, and Zdinak (2014) show weaker anchoring in Brazil than in Mexico or Chile; Robitaille, Zhang, and Weisberg (2024) reach similar conclusions with newer data. Witheridge (2024) examines shorter- to medium-term horizons and finds that inflation expectations actually increase following monetary easing in some emerging markets. Bonomo et al. (2024) document that Brazil's 2011 policy shift significantly affected long-run expectations. Relative to these papers, our contribution is to show that these elasticities are directly linked to the monetary authority's incentives to build reputation within a benchmark model of optimal monetary policy. This connection allows us to use the elasticities to discipline the strength of reputation-building motives and to quantify the macroeconomic stabilization benefits of anchoring inflation expectations. This approach builds on our earlier work in Bocola, Dovis, Jørgensen, and Kirpalani (2024), where we combined high-frequency identification with a structural model to assess the macroeconomic implications of changes in the perceived reaction function of the Federal Reserve.

Layout. The paper is structured as follows. Section 2 introduces the model. In section 3 we use an analytically tractable special case of the model to discuss the reputation channel and our empirical strategy to measure it. Section 4 provides the estimates of the elasticity of long-run inflation expectations to monetary surprises for a set of emerging and advanced economies. Section 5 presents the main quantitative counterfactuals. Section 6 concludes.

2 Model

2.1 Environment

We consider a sticky price economy with Rotemberg adjustment costs. Time is discrete and indexed by $t = 0, 1, ...$ The economy is populated by a continuum of identical and infinitely lived households, final good producers, intermediate good producers, and a central bank.

Households have preferences over consumption, c_t, and hours worked, l_t given by

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t U(c_t, l_t), \tag{1}$$

where β is the discount factor. We assume that

$$U(c,l) = \frac{c^{1-\sigma}-1}{1-\sigma} - \chi \frac{l^{1+1/\nu}}{(1+1/\nu)}.$$

The final consumption good is produced by combining differentiated intermediate goods according to the CES technology

$$y_t = \left(\int_0^1 y_i^{\frac{1}{\mu_t}} di \right)^{\mu_t} \tag{2}$$

where $\mu_t > 1$ is related to the constant elasticity of substitution across varieties, ϵ_t, by the following relation: $\mu_t = \epsilon_t / [\epsilon_t - 1]$. Time-variation in μ_t will generate shifts in the Phillips curve. So, we will refer to μ_t as the *supply shock*.

The intermediate goods are produced using a linear technology,

$$y_{it} = z_t n_{it} \tag{3}$$

where z_t is an aggregate productivity shock and n_{it} denotes the labor utilized in the production of good i. Feasibility requires that $\int n_{it} di = l_t$.

Each intermediate good is supplied by a monopolistically competitive firm that hires labor in a competitive labor market and operates the linear technology (3). The firm faces quadratic adjustment costs when changing its prices,

$$\frac{\phi z_t}{2} \left[\frac{P_{it}}{P_{it-1}} - 1 \right]^2.$$

For tractability, we assume that the adjustment cost scales with aggregate productivity.[8]

We assume that the two exogenous shocks, μ_t and z_t, follow AR(1) processes (in log):

$$\log \mu_{t+1} = (1-\rho_\mu) \log \bar{\mu} + \rho_\mu \log \mu_t + \sigma_\mu \varepsilon_{\mu,t+1}$$
$$\log z_{t+1} = (1-\rho_z) \log \bar{z} + \rho_z \log z_t + \sigma_z \varepsilon_{z,t+1}$$

where the random variables $\varepsilon_{\mu,t+1}$ and $\varepsilon_{z,t}$ follow a standard normal distribution, $\bar{\mu} > 1$ and $\bar{z} > 0$ are the steady state value for the markup and productivity, respectively. We denote the implied transition probability for $s_t = (\mu_t, z_t)$ as $\Gamma(s_{t+1}|s_t)$.

[8] One justification for this assumption is that adjustment costs are in terms of wages that are proportional to the aggregate productivity shock z_t.

2.2 Private equilibrium conditions

In the Appendix, we define the competitive equilibrium for an arbitrary interest rate policy. The household's equilibrium behavior is summarized by an intertemporal Euler equation

$$\frac{1}{1+i_t} = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{1}{1+\pi_{t+1}} \right] \tag{4}$$

where π_{t+1} is the inflation rate, and the behavior of monopolistically competitive producers is summarized by the Phillips curve

$$\tilde{\pi}_t = \frac{y_t \left[\frac{\mu_t c_t^{-\sigma} \chi {l_t}^{1/\nu}}{z_t} - 1 \right]}{\phi \left[\mu_t - 1 \right]} + \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{z_{t+1}}{z_t} \tilde{\pi}_{t+1} \right] \tag{5}$$

where $\tilde{\pi}_t \equiv \pi_t (1+\pi_t)$.

Finally, the resource constraint is

$$c_t = z_t l_t - \frac{\phi z_t}{2} \pi_t^2 = y_t - \frac{\phi z_t}{2} \pi_t^2. \tag{6}$$

2.3 Monetary policy

The central bank chooses the nominal interest rates to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t R \left(\pi_t, y_t; \theta_t, \varepsilon_t, s_t \right) \tag{7}$$

where $R \left(\pi_t, y_t; \theta_t, \varepsilon_t, s_t \right)$ is the period utility function, which depends on inflation, output, the central bank weights over these two, and the state s_t. The central bank's weights depends on two components: a persistent preference shock θ_t and an i.i.d. disturbance ε_t. We will refer to θ_t as the central bank *type* and assume that it can take on two values, $\theta_t \in \{\theta_D, \theta_H\}$, where θ_D stands for *Dove* and θ_H for *Hawk* with $\theta_H > \theta_D$. A hawkish central bank cares more about inflation stabilization relative to output stabilization than a dovish central bank. We assume that θ_t follows a Markov chain with transition matrix P_θ, and that ε_t is normally distributed.

We consider the following specification of the R function,

$$R\left(\pi_t, y_t; s_t, \alpha_t, y_t^*\right) = -\frac{1}{2}\left[\alpha_t \pi_t^2 + (1-\alpha_t)\left(\frac{y_t - y^*\left(s_t\right)}{y^*\left(s_t\right)}\right)^2\right] \tag{8}$$

where the weight on inflation is

$$\alpha_t = \alpha\left(\theta_t, \varepsilon_t\right) \tag{9}$$

with α being a strictly increasing function with range $[0,1]$. Thus, for any realization of ε_t the weight on inflation stabilization is higher for the hawkish central bank than the dove, $\alpha\left(\theta_H, \varepsilon_t\right) > \alpha\left(\theta_D, \varepsilon\right)$ for all ε. The term $y^*\left(s_t\right)$ is a potentially state varying output target. For example, it could be the efficient level of output, $y^*\left(s_t\right) = z_t \chi^{-1/(\sigma+1/\nu)}$.

The output target allows for an inflation bias as in Kydland and Prescott (1977) and Barro and Gordon (1983a). If $y^*\left(s_t\right)$ is higher than the flexible price level of output – which itself is lower than the efficient level due to the intermediate good firm's markups that are not offset by a fiscal subsidy – the central bank will have an incentive to generate positive inflation to bring the equilibrium level of output closer to the target.

The type of the central bank is private information and private agents learn about the central bank's type over time by observing its action. The presence of the i.i.d. shock ε_t will ensure that actions do not fully reveal the central bank's type.[9] We denote by ρ_t the *reputation* of the central bank, which is defined as the probability that private agents assign to facing the hawkish central bank. The initial level of reputation common to all private agents is ρ_0.

2.4 Markov Perfect Equilibrium

Let $h^{t-1} = \left(s_0, x_0; s_1, x_1; ....; s_{t-1}, x_{t-1}\right)$ be the history of shocks that are public information, $s_t = \left(\mu_t, z_t\right)$, and outcomes $x_t = \left(i_t, y_t, c_t, \pi_t\right)$. The timing is as follows: in each period t, the central bank acts first and chooses the interest rate i_t to maximize (7) given their private history $\left(h^{t-1}, s_t, \theta_t, \varepsilon_r\right)$, then private agents simultaneously choose c_t, y_t and π_t according to (4)–(6) given the public history $\left(h^{t-1}, s_t, i_t\right)$.

We focus on a *Markov perfect equilibrium* where the public history h^{t-1} is summarized by the prior of facing the hawkish central bank ρ_t. The central bank strategy depends on s_t, its reputation ρ_t at the beginning of the period, and its private information $(\theta_t, \varepsilon_t)$. The equilibrium outcome from competitive agents are instead functions of the exogenous

[9] The shock ε_t functions similarly to an "implementation" shock, introducing noise that prevents private agents from perfectly observing the government's intended policy.

public state s_t, the initial reputation, and the observed nominal interest rate i_t.

To minimize the notation, for any function $f(s', \theta', \varepsilon')$ we define the central bank's expectations as $\mathbb{E}^{cb} f(s', \theta', \varepsilon') = \int \int \sum_{\theta'} f(s', \theta', \varepsilon') P_{\theta'\theta} d\Gamma(s'|s) \Phi(\varepsilon') d\varepsilon'$ and for the private agents as $\mathbb{E}^{\rho} f(s', \theta', \varepsilon') = \int \int [\rho f(s', \theta_H, \varepsilon') + (1-\rho) f(s', \theta_D, \varepsilon')] d\Gamma(s'|s) \Phi(\varepsilon') d\varepsilon'$, where Φ is the density of a standard normal random variable.

Definition. A Markov perfect equilibrium is a central bank's strategy $i(s, \rho, \theta, \varepsilon)$ and its associated value function $V(s, \rho, \theta, \varepsilon)$, allocation and price rules, $Y(s, \rho, i)$, $C(s, \rho, i)$, $\pi(s, \rho, i)$, and a law of motion for reputation such that i) the central bank's strategy and value function solve

$$V(s, \rho, \theta, \varepsilon) = \max_i R(\pi(s, \rho, i), Y(s, \rho, i); \alpha(\theta, \varepsilon)) + \beta \mathbb{E}^{cb} V(s', \rho'(s, \rho, i), \theta', \varepsilon') \tag{10}$$

ii) the allocation and price rules satisfy (4)–(6), or

$$\frac{1}{1+i} = \mathbb{E}^{\rho'} \left[\frac{C(s', \rho', i')^{-\sigma} / C(s, \rho, i)^{-\sigma}}{1 + \pi(s', \rho', i')} \right] \tag{11}$$

$$\tilde{\pi}(s, \rho, i) = \frac{Y(s, \rho, i) \left[\frac{\mu\chi}{z} \frac{(Y(s,\rho,i)/z)^{1/\nu}}{C(s,\rho,i)^{-\sigma}} - 1 \right]}{\phi z [\mu - 1]} + \beta \mathbb{E}^{\rho'} \left[\left(\frac{C(s', \rho', i')}{C(s, \rho, i)} \right)^{-\sigma} \frac{z'}{z} \tilde{\pi}(s', \rho', i') \right] \tag{12}$$

$$Y(s, \rho, i) = C(s, \rho, i) + \frac{\phi}{2} z \pi(s, \rho, i)^2 \tag{13}$$

where $i' = i(s', \rho', \theta', \varepsilon')$, and iii) the law of motion for reputation satisfies Bayes' rule whenever possible i.e. for all $i \in \text{range}(i(s, \rho, \theta, \cdot))$

$$\rho'(s, \rho, i) = \frac{\rho \Pr(i|\theta_H, s, \rho) P_{HH} + (1-\rho) \Pr(i|\theta_D, s, \rho) P_{DH}}{\rho \Pr(i|\theta_H, s, \rho) + (1-\rho) \Pr(i|\theta_D, s, \rho)}. \tag{14}$$

We will focus on a separating equilibrium in which the central bank chooses a different policy for any private state (θ, ε) and restrict attention to equilibria in which it is not possible to perfectly learn the central bank's type by observing its choices. [10]

Note that, given the allocation and price rules, the central bank is effectively choosing the current equilibrium outcome $x = (i, y, c, \pi)$ to maximize its value subject to the two conditions that characterize private behavior, (11)–(13), and the law of motion for reputation (14). Thus, abstracting from the zero lower bound (ZLB) on the nominal interest rate, it is convenient to characterize the equilibrium outcome by finding $\pi(s, \rho, \theta, \varepsilon)$, $Y(s, \rho, \theta, \varepsilon)$, $C(s, \rho, \theta, \varepsilon)$, and $V(s, \rho, \theta, \varepsilon)$ that solve a primal version of the central bank's problem where

[10] To guarantee this is the case, when numerically solving the model, we check that the range of $i(s, \rho, \theta, \cdot)$ does not depend on θ for any (s, ρ).

it chooses the allocation and prices subject to the equilibrium conditions:

$$V(s,\rho,\theta,\varepsilon) = \max_{\pi,y,c} R(\pi,y;\alpha(\theta,\varepsilon),s) + \beta \mathbb{E}^{cb} V(s',\rho',\theta',\varepsilon') \tag{15}$$

subject to

$$\begin{aligned} y &= c + \frac{\phi}{2} z \pi^2 \\ \tilde{\pi} &= \frac{y\left[\frac{\mu \chi (y/z)^{1/\nu} c^{\sigma}}{z} - 1\right]}{\phi z [\mu - 1]} + \beta \frac{c^{\sigma}}{z} \tilde{\Pi}(s,\rho,y) \\ \rho' &= p(s,\rho,y) \end{aligned}$$

taking as given the law of motion for its reputation, $p(s,\rho,y)$, and the evolution of private sector's inflation expectations captured by the term $\tilde{\Pi}(s,\rho,y)$. The term $\beta \frac{c^{\sigma}}{z} \tilde{\Pi}(s,\rho,y)$ is the expectation of inflation under the risk-neutral measure. Note that we dropped the Euler equation (11) because it simply defines the nominal interest rate i. In equilibrium, it must be that

$$\tilde{\Pi}(s,\rho,y) = \int\int \left[\rho' \frac{z' \tilde{\pi}(s',\rho',\theta_H,\varepsilon')}{C(s',\rho',\theta_H,\varepsilon')^{\sigma}} + (1-\rho') \frac{z' \tilde{\pi}(s',\rho',\theta_D,\varepsilon')}{C(s',\rho',\theta_D,\varepsilon')^{\sigma}}\right] d\Gamma(s'|s)\, \Phi(\varepsilon')\, d\varepsilon' \tag{16}$$

where $\rho' = p(s,\rho,y)$, and

$$p(s,\rho,y) = \frac{\rho \Phi(\varepsilon : y = Y(s,\rho,\theta_H,\varepsilon))\, P_{HH} + (1-\rho)\, \Phi(\varepsilon : y = Y(s,\rho,\theta_D,\varepsilon))\, P_{DH}}{\rho \Phi(\varepsilon : y = Y(s,\rho,\theta_H,\varepsilon)) + (1-\rho)\, \Phi(\varepsilon : y = Y(s,\rho,\theta_D,\varepsilon))}, \tag{17}$$

where Φ is the density of a standard normal random variable. Finding an equilibrium then boils down to finding a set of allocation rules $\{\pi(s,\rho,\theta,\varepsilon), C(s,\rho,\theta,\varepsilon), Y(s,\rho,\theta,\varepsilon)\}$ and a value function that solve (15) and satisfy (16) and (17).

Note that in the primal version of the central bank problem, (15), we dropped the Euler equation (11) from the constraint set and the nominal interest rate as a choice. This is because the Euler equation simply defines the nominal interest rate given the chosen allocation.

3 Reputation channel: theory and measurement

In this section, we study how reputation concerns shape the central bank's behavior and equilibrium outcomes. We then outline how to measure the strength of this channel in the data. To make the analysis transparent, we begin with a tractable special case of the

economy. Specifically, we assume there is no inflation bias: the central bank does not seek to stimulate output above the steady state flexible-price level, that is we set $y^*(s) = \left(\frac{z^{1+1/\nu}}{\bar{\mu}\chi}\right)^{1/(\sigma+1/\nu)}$ Given this assumption, both central bank types achieve $y_{ss} = y^*(\bar{s})$ where $\bar{s} = (\bar{\mu}, \bar{z})$ and $\pi_{ss} = 0$ in a deterministic steady state. Thus, we can consider a log-linear approximation of the private equilibrium conditions around this point. We denote log deviations from the output target with a hat, e.g. $\hat{y}_t \equiv \log y_t - \log y_t^* \approx (y_t - y_t^*)/y_t^*$. For all the other variables, a hat denotes deviation from their steady state value. Using this notation, the equilibrium behavior of the private sector can be characterized by an intertemporal Euler equation that replaces (4)

$$\hat{y}_t = -\frac{1}{\sigma}\left[i_t - \mathbb{E}_t^{\rho_{t+1}}\pi_{t+1}\right] + \mathbb{E}_t\hat{y}_{t+1} - \frac{1+1/\nu}{\sigma+1/\nu}(1-\rho_z)\hat{z}_t, \tag{18}$$

where the expected change in productivity, $\mathbb{E}_t(\hat{z}_{t+1} - \hat{z}_t) = -(1-\rho_z)\hat{z}_t$, act as a demand shock, and by a Phillips curve that replaces (5),

$$\pi_t = \kappa\hat{y}_t + \kappa_\mu\hat{\mu}_t + \beta\mathbb{E}_t^{\rho_{t+1}}\pi_{t+1} \tag{19}$$

where $\kappa = \frac{\left(\sigma+\frac{1}{\nu}\right)y_{ss}}{\phi[\mu-1]}$, $\kappa_\mu = \frac{\kappa}{\sigma+1/\nu}$, and $\hat{c}_t = \hat{y}_t + \frac{1+1/\nu}{\sigma+1/\nu}\hat{z}_t$ in place of (6).[11] While this environment misses some of the features present in the fully fledged economy—notably the possibility of an inflation bias—it is enough to isolate the core mechanism of our model.

Section 3.1 studies optimal monetary policy in this environment. When setting monetary policy, the central bank understands that its policy decisions affect the expectations formation process of firms and households. We show that this motive leads the central bank to raise real interest rates more aggressively after an inflationary shock; by doing so, the central banks acquires reputation, and higher reputation allows to "anchor" inflation expectations and to better smooth the effect of the shock on current inflation. Because of this *reputation channel*, the central bank adopts a more hawkish stance relative to a scenario in which the policymaker does not internalize expectation formation. Importantly, we show that these motives to build reputation are state dependent, and are particularly strong when private agents are uncertain about the central bank's type.[12]

We next turn to how these reputation effects can be measured. In our framework, the central bank's incentive to acquire reputation depends on how responsive inflation expecta-

[11] See the Appendix for a derivation.

[12] We abstract from an effective lower bound on the nominal interest rate, as our focus is on the post-Covid inflation surge. When policy is constrained by a lower bound—either because it is currently binding or expected to bind due to demand shocks—the central bank would have incentives to be perceived as more dovish in order to relax this constraint.

tions are to policy changes. Section 3.2 shows that this statistic can be estimated by looking at the response of inflation expectations to "monetary surprises"—defined as the private sector forecast updates of nominal interest rates. This result is key to our quantitative strategy, which uses high-frequency identification methods and indirect inference to discipline the importance of the reputation channel in our model.

Section 3.3 concludes by discussing how these insights extends to the fully fledged economy presented in Section 2.

3.1 The reputation channel

We consider a case where uncertainty about the central bank's type is resolved after the first period, and the type of the central bank is perfectly persistent. In period 1, the type is unknown and private agents form expectations about future outcomes according to their prior, ρ, and the observed policy choices of the central bank. In period 2, the central bank's type is credibly revealed to households and firms, and there are no shocks to the preferences of the policymaker from that point onward ($\varepsilon_t = 0$ for all $t \geq 2$).

We solve for the equilibrium allocation by backward induction, first deriving the equilibrium outcome from period 2 onward and then studying the optimal policy problem at date 1.

Long-run under full information The equilibrium outcome from period 2 onward is the equilibrium outcome in a Markov perfect equilibrium with public information and exogenous state $\hat{s}_t$. Abstracting from the ZLB constraint, there exists a unique Markov equilibrium outcome.[13] The central bank's problem is

$$V_2(\hat{s}, \theta) = \max_{\pi, \hat{y}} R(\pi, \hat{y}; \alpha(\theta, 0)) + \beta \mathbb{E}\left[V_2(\hat{s}', \theta) | \hat{s}\right] \tag{20}$$

subject to the log-linear Phillips curve (19). Since there is symmetric information between the private sector and the central bank, expectations are taken with respect to the exogenous state $\hat{s}'$ given the current realization $\hat{s}$.

Letting R_x denote the derivative of R with respect to argument x, we have that the equilibrium $(\pi, \hat{y})$ at any point in time must satisfy the static first order condition

$$R_\pi \kappa = -R_{\hat{y}} \tag{21}$$

[13] Armenter (2018) shows that there can be multiple Markov equilibria when respecting the ZLB. Furthermore, even abstracting from the ZLB, there are also multiple time-dependent equilibria that differs from the unique Markov perfect equilibrium and have the property that inflation goes to infinity as $t \to \infty$.

and the Phillips curve. Using the definition of R in (8), we can solve these two equations that characterize the equilibrium outcome to obtain allocation and pricing rules[14]

$$\hat{y}_2(\hat{s},\theta) = -\frac{\alpha(\theta,0)\kappa^2}{[1-\alpha(\theta,0)](1-\beta\rho_\mu)+\alpha(\theta,0)\kappa^2}\hat{\mu}$$
$$\pi_2(\hat{s},\theta) = \frac{[1-\alpha(\theta,0)]\kappa}{[1-\alpha(\theta,0)](1-\beta\rho_\mu)+\alpha(\theta,0)\kappa^2}\hat{\mu}.$$

After a positive supply shock, the monetary authority increases real interest rates and induces a fall in output, which mitigates the effect of the shock on inflation. This policy response is stronger the greater the weight on inflation stabilization $\alpha(\theta,0)$, and the steeper the Phillips curve—the larger is κ.

It follows that, conditional on a positive supply shock $\hat{\mu}$, the economy with a hawkish central bank has lower inflation and more negative output than an economy with a dovish central bank.

Short-run with reputational concerns We now consider the problem in period 1 given the publicly observable state, $(\hat{s},\rho)$, the value function from period 2 onward, $V_2(\hat{s},\theta)$, and the equilibrium level of inflation that will prevail in the second period, $\pi_2(\hat{s},\theta)$. We define the private sector's expectations of inflation in the next period *after* the central bank acts as

$$\Pi_2(\hat{s},\rho,\hat{y}) = p(\hat{s},\rho,\hat{y})\,\mathbb{E}\left[\pi_2\left(\hat{s}',\theta_H\right)|\hat{s}\right] + \left[1-p(\hat{s},\rho,\hat{y})\right]\mathbb{E}\left[\pi_2\left(\hat{s},'\theta_D\right)|\hat{s}\right] \tag{22}$$

where $p(\hat{s},\rho,\hat{y})$ satisfies (17) given the central bank's policy function $\hat{y}(\hat{s},\rho,\theta,\varepsilon)$.

We can then write the problem for the central bank in period 1 as

$$V_1(\hat{s},\rho,\theta,\varepsilon) = \max_{\pi,\hat{y}} R(\pi,\hat{y};\alpha(\theta,\varepsilon)) + \beta\mathbb{E}\left[V_2\left(\hat{s}',\theta\right)|\hat{s}\right] \tag{23}$$

subject to

$$\pi = \kappa\hat{y} + \kappa_\mu\hat{\mu} + \beta\Pi_2(\hat{s},\rho,\hat{y}) \tag{24}$$

taking the function $\Pi_2(\cdot)$ as given. Because the continuation value $V_2(\hat{s}',\theta)$ does not depend on the choices in period 1, the optimal choice is obtained by maximizing R subject to

[14]Note that realizations of the productivity shock $\hat{z}_t$ do not affect the behavior of inflation and the output gap. Productivity shocks only enter in the log-linearized Euler equations. The central bank can then attain zero inflation and a a zero output gap simply by changing the nominal interest rate to offset the term $-\frac{1+1/\nu}{\sigma+1/\nu}(1-\rho_z)\hat{z}_t$ in (18) with no trade-off between output and inflation stabilization.

the Phillips curve (24), or simply:

$$\hat{y}_1(\hat{s},\rho,\theta,\varepsilon) = \arg\max_{\hat{y}} R\left(\kappa\hat{y} + \kappa_\mu\hat{\mu} + \beta\Pi_2(\hat{s},\rho,\hat{y}),\hat{y};\alpha\right). \tag{25}$$

Finding an equilibrium then boils down to finding a policy $\hat{y}_1(\cdot)$ and expectations $\Pi_2(\cdot)$ that are consistent with (22) and (25). We focus on a *separating equilibrium* in which the central bank chooses a different policy for any $\alpha = \alpha(\theta,\varepsilon)$.[15] In fact, inspecting problem (25), the optimal $\hat{y}_1$ depends on θ and ε solely through their effect on $\alpha = \alpha(\theta,\varepsilon)$. We can thus write the equilibrium outcome as $\hat{y}_1(\hat{s},\rho,\alpha)$.

We now solve for the optimal policy of the central bank. Toward this end, we can write the first order condition from the central bank problem (25)

$$R_\pi\left[\kappa + \beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\right] = -R_{\hat{y}}. \tag{26}$$

Equation (26) is very similar to the optimality condition under full information, equation (21), with the exception that the monetary authority now also takes into account the effects of his policy actions at date 1 on inflation expectations—the term

$$\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}} \equiv -\mathbb{E}\left[\pi_2(\hat{s}',\theta_D) - \pi_2(\hat{s}',\theta_H)\,|\hat{s}\right]\frac{\partial p(\hat{s},\rho,\hat{y})}{\partial\hat{y}}.$$

As we show next, this term is non-negative. To understand why, consider a situation in which the economy faces an inflationary supply shock, $\hat{\mu} > 0$. From the previous analysis, we know that the dovish central banker will implement a higher level of inflation than the hawkish one when $\hat{\mu} > 0$, so $\mathbb{E}\left[\pi_2(\hat{s}',\theta_D) - \pi_2(\hat{s}',\theta_H)\,|\hat{s}\right] > 0$. Thus, the sign of $\partial\Pi_2(\hat{s},\rho,\hat{y})/\partial\hat{y}$ depends on the sign of $\partial p(\hat{s},\rho,\hat{y})/\partial\hat{y}$. This latter expression tells us in which direction the private sector updates its prior after seeing a marginally higher level of output in the economy—with a negative value implying that a higher level of output induces the private sector to update their prior toward the dovish policymaker.

Let $a(\hat{y};s,\rho)$ be the inverse function of $\hat{y}_1(s,\rho,\cdot)$ given s and ρ. For all $\hat{y}$ on path we can then rewrite (17) as

$$p(\hat{s},\rho,\hat{y}) = P(\alpha = a(\hat{y};\hat{s},\rho);\rho) \equiv \frac{\rho\Pr(\alpha = a(\hat{y};\hat{s},\rho)\,|\theta_H)}{\rho\Pr(\alpha = a(\hat{y};\hat{s},\rho)\,|\theta_H) + (1-\rho)\Pr(\alpha = a(\hat{y};\hat{s},\rho)\,|\theta_D)}.$$

[15] This economy could also admit pooling equilibria in which for all α the central bank chooses the same policy and there is no updates in the type of the central bank, $\rho' = \rho$. These pooling equilibria are supported by off-equilibrium path beliefs that assign a reputation lower than ρ for deviations that entail more tightening than along the equilibrium path. These beliefs do not satisfy the intuitive criterion.

Moreover, if $\hat{y}_1(\hat{s}, \rho, \alpha)$ is differentiable in α, we have

$$\frac{\partial p(\hat{s}, \rho, \alpha)}{\partial \hat{y}} = \frac{P'(a(\hat{y}; \hat{s}, \rho))}{\partial \hat{y}_1(\hat{s}, \hat{\rho}, \alpha) / \partial \alpha}.$$

Because higher values of α are more likely when the monetary authority is of type θ_H, $P' \geq 0$, the sign of $\partial p(\hat{s}, \rho, \hat{y}) / \partial \hat{y}$ depends solely on the sign of $\partial \hat{y}_1(\hat{s}, \rho, \alpha) / \partial \alpha$. For $\hat{\mu} > 0$, output strictly decreases in α because the central bank is more willing to sacrifice output when the weight on inflation stabilization is greater. Thus, a higher realization of $\hat{y}$ signals that it is more likely that the central bank is a dove, $\partial p(\hat{s}, \rho, \hat{y}) / \partial \hat{y} < 0$. This, in turn, will push up inflation expectations, since the dovish policymaker is expected to implement higher inflation in the future, $\frac{\partial \Pi_2(\hat{s}, \rho, \hat{y})}{\partial \hat{y}} \geq 0$.[16]

Thus, when facing an inflationary supply shock, the central bank understands that it can moderate the effect of the shock on current inflation by increasing its reputation, as this leads to inflation expectations that are more "anchored" to those of the hawkish policymaker. *Ceteris paribus*, this channel incentivizes the central bank to stabilize more inflation at the cost of a larger fall in output. We show this in the next proposition, which compares the equilibrium of our economy to one with perfect information.

Proposition 1. *The central bank behaves in a more hawkish manner when its type is unknown by the private sector. That is, if $\mu > 0$ ($\mu < 0$) then $y_1(s, \rho, \alpha)$ is more negative (positive) than in the Markov equilibrium under perfect information under the Hawk, strictly so if $\rho \in (0, 1)$.*

Under perfect information, reducing output helps stabilize inflation directly through the Phillips curve. With imperfect information, the same policy also has an indirect effect: it raises the central bank's reputation, which lowers expected future inflation and thereby moderates current inflation. This reputation channel induces the central bank to follow a tighter policy than in the perfect information case.

Interestingly, the reputation channel is state-dependent. First, it operates only when the private sector is uncertain about the policymaker's type. When the prior is dogmatic, that is when $\rho = 0$ or $\rho = 1$, current monetary policy decisions have no effect on future expectations, $\partial \Pi_2(\hat{s}, \rho, \hat{y}) / \partial \hat{y} = 0$, and the reputation channel will be muted.

Second, in this example it operates only when the economy faces supply shocks. When $\hat{\mu} = 0$, there is no inflation/output trade-off and, given the functional form used for R, both types of policymakers will want to implement the same allocation ($\hat{y} = 0, \pi = 0$). Therefore, the private sector's beliefs about the central bank's type equal their prior irrespective

[16] Symmetrically, if $\hat{\mu} < 0$ the equilibrium output is strictly increasing in α. So a higher realization of $\hat{y}$ increases the government reputation. Also in this case $\partial \Pi_2(\hat{s}, \rho, \hat{y}) / \partial \hat{y} \geq 0$ since $\mathbb{E}[\pi_2(\hat{s}', \theta_D) - \pi_2(\hat{s}', \theta_H) | \hat{s}] < 0$ when $\hat{\mu} < 0$.

of the observed policy choices ($\rho' = \rho$), and the monetary authority will not have an incentive to choose policy in order to change its reputation. In those cases, the equilibrium allocations with perfect and imperfect information coincide.

Macroeconomic effects of the reputation channel What are the macroeconomic implications of these reputational incentives? To answer this question, it is useful to compare the equilibrium outcome of our economy to that of an economy with a "myopic" central bank—a monetary authority that ignores the consequences of its actions on inflation expectations. These two economies are equivalent in all other respects, including the expectation formation process of firms and households. Therefore, comparing these two economies will isolate the effects that reputational incentives have on monetary policy and macroeconomic outcomes.

Specifically, we assume that the myopic central bank chooses output and inflation according to the static trade-off

$$R_\pi\left(\pi_1^{myopic}, \hat{y}_1{}^{myopic}; \alpha\right)\kappa = -R_{\hat{y}}\left(\pi_1^{myopic}, \hat{y}_1{}^{myopic}; \alpha\right) \tag{27}$$

while the private sector forms expectations as in the benchmark economy, with inflation being determined by the Phillips curve $\pi_1^{myopic} = \kappa \hat{y}_1{}^{myopic} + \kappa_\mu \hat{\mu} + \beta \Pi_2\left(\hat{s}, \rho, \hat{y}_1{}^{myopic}\right)$. The next proposition characterizes the differences in output and inflation between this economy and the benchmark.

Proposition 2. *Consider a myopic central bank that sets policy according to (27). Then, for small $\hat{\mu} > 0$, we have*

$$\hat{y}_1^{myopic}(\hat{s}, \rho, \alpha) = \hat{y}_1(\hat{s}, \rho, \alpha) + \frac{\alpha\beta \frac{\partial \Pi_2(\hat{s}, \rho, \hat{y}_1)}{\partial \hat{y}_1} \pi(\hat{s}, \rho, \alpha)}{\left[\alpha\kappa^2 + \alpha\kappa\beta \frac{\partial \Pi_2(\hat{s}, \rho, \hat{y}_1)}{\partial \hat{y}_1} + (1-\alpha)\right]} > \hat{y}_1(\hat{s}, \rho, \alpha) \tag{28}$$

$$\pi_1^{myopic}(\hat{s}, \rho, \alpha) = \pi_1(\hat{s}, \rho, \alpha) + \frac{\left[\kappa + \beta \frac{\partial \Pi_2(\hat{s}, \rho, \hat{y}_1)}{\partial \hat{y}_1}\right] \alpha\beta \frac{\partial \Pi_2(\hat{s}, \rho, \hat{y}_1)}{\partial \hat{y}_1} \pi(\hat{s}, \rho, \alpha)}{\left[\alpha\kappa^2 + \alpha\kappa\beta \frac{\partial \Pi_2(\hat{s}, \rho, \hat{y}_1)}{\partial \hat{y}_1} + (1-\alpha)\right]} > \pi_1(\hat{s}, \rho, \alpha). \tag{29}$$

Equations (28) and (29) show that reputational incentives lead to tighter monetary policy following an inflationary supply shock, something that leads to a larger recession, as $\hat{y}_1^{myopic}(\hat{s}, \rho, \alpha) > \hat{y}_1(\hat{s}, \rho, \alpha)$, but less severe impact of the shock on inflation, $\pi_1^{myopic}(\hat{s}, \rho, \alpha) > \pi_1(\hat{s}, \rho, \alpha)$.

These expressions also show that the magnitudes of these differences crucially depend on the elasticity $\partial \Pi_2(\hat{s}, \rho, \hat{y}_1) / \partial \hat{y}_1$: the more responsive inflation expectations are to the

policy actions of the central bank, the larger the incentives to use interest rates to manage inflation expectations, and the larger the impact of reputation on output and inflation.

Numerical example We now illustrate the workings of the model using a numerical example. The solid lines in Figure 1 plots the equilibrium outcome in period 1 as a function of the supply shock $\hat{\mu}$, and compares it to the equilibrium outcome of the economy with full information (circled lines). For this numerical experiment, we assume that the policymaker's type is the Hawk ($\theta = \theta_H$), that the private sector initially holds a low level of reputation $\rho = 0.5$, and that $\varepsilon = 0$. Because we are conditioning on $\theta = \theta_H$, the full information economy also corresponds to our benchmark economy with an initial reputation of $\rho = 1$. Therefore, the comparison between these two cases is also informative about the role of inherited reputation for monetary policy. The panels plot the equilibrium outcome as a function of the markup shock $\hat{\mu}$.

As discussed above, when private agents are uncertain about the central banker's type, the monetary authority responds to a positive supply shock by increasing interest rates more aggressively and implementing a larger fall in output. Despite the more restrictive policy stance, inflation outcomes in this example are worse in the economy with imperfect information. This happens because inflation expectations respond more strongly to supply shocks in the economy with imperfect information.

To see that, let's rewrite the Phillips curve as

$$\pi_1 - \kappa \hat{y}_1 = \kappa_\mu \hat{\mu} + \beta \left\{ \mathbb{E}\left[\pi_2\left(\hat{s}', \theta_D\right) | \hat{s} \right] - p(\hat{s}, \rho, \hat{y}_1) \mathbb{E}\left[\pi_2\left(\hat{s}', \theta_D\right) - \pi_2\left(\hat{s}', \theta_H\right) | \hat{s} \right] \right\}.$$

For fixed $\hat{y}_1$, a lower ρ reduces $p(\hat{s}, \rho, \hat{y}_1)$. Because $\mathbb{E}\left[\pi_2\left(\hat{s}', \theta_D\right) | \hat{s}\right] > \mathbb{E}\left[\pi_2\left(\hat{s}', \theta_H\right) | \hat{s}\right]$ when $\hat{\mu} > 0$, a reduction in ρ increases inflation for any level of $\hat{y}_1$. In other words, the central bank in the imperfect information economy faces a worse inflation-output trade-off than the central bank in the perfect information economy, and achieves a higher level of inflation even if it implements a more restrictive policy.

These reputational considerations, however, contribute to reduce the inflationary effects of the supply shocks in the economy. This is illustrated by comparing the solid line to the dashed line—the equilibrium outcome in the economy with a myopic central bank. Ignoring the effect that monetary policy has on inflation expectations, the myopic central bank chooses a less restrictive response, and output is higher than in the baseline economy. This leads to an increase in current inflation driven both by the direct output effect from the Phillips curve and by higher inflation expectations.

As mentioned earlier, the comparison between our economy and the perfect information

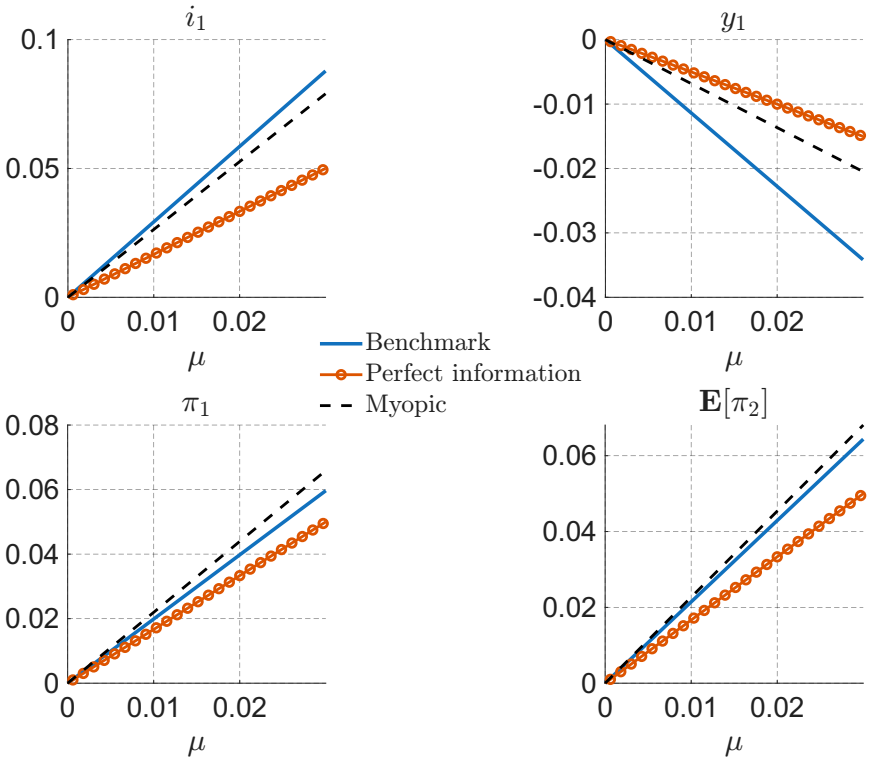


Figure 1: Equilibrium response to a supply shock

Notes: For this numerical illustration, we set $\kappa = 0.2$, $\beta = 0.94$, $\theta_H = 0.6$, $\theta_D = 0.2$, $\sigma_\mu = 0.02$, $\sigma_\varepsilon = 2$. The policy functions are constructed setting $\theta = \theta_H$ and $\varepsilon = 0$. The solid line reports the policy function for $\rho = 0.5$, while the circled line for $\rho = 1$.

case is also informative about how optimal policy and outcomes vary according to the level of reputation inherited by the central bank at date 1. Specifically, we have seen that a central bank with $\rho = 1$ implements better outcomes (inflation closer to zero/output closer to target) than a central bank with lower reputation, $\rho = 0.5$. To further understand the role of reputation, we plot in Figure 2 the policy function for output, inflation and expected inflation as a function of ρ, conditional on an inflationary supply shock $\hat{\mu} > 0$.

Consistent with the above discussion, higher reputation reduces the effects of the supply shock on expected inflation, with $\Pi_2(\hat{s}, \rho, \hat{y})$ declining toward $\mathbb{E}[\pi_2(\hat{s}', \theta_H) | \hat{s}]$ as ρ increases.

The relationship between $\hat{y}_1$ and ρ is, instead, more complex. When ρ is close to zero, $\hat{y}_1$ decreases as a function of ρ: at $\rho = 0$, the central bank's actions have no effect on its reputation, so the optimal choice of $(\hat{y}_1, \pi_1)$ follows the static rule of equation (20). As ρ increases from zero, the central bank has an incentive to act in a more hawkish fashion to build reputation, which reduces output. At higher levels of ρ, however, inflation expecta-

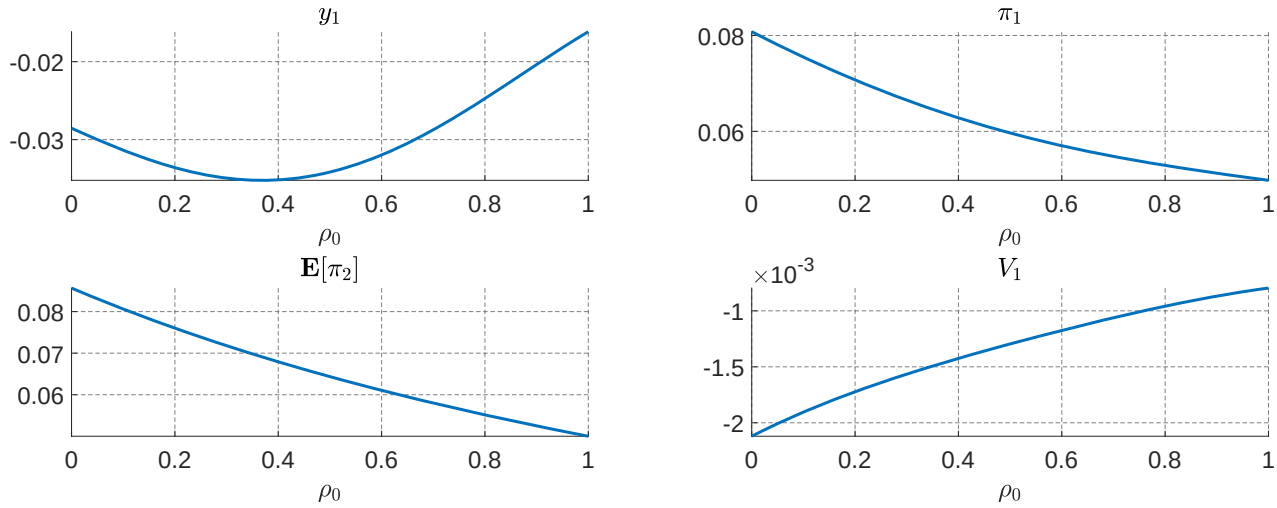


Figure 2: Equilibrium response to reputation

Notes: For this numerical illustration, we set $\kappa = 0.2$, $\beta = 0.94$, $\theta_H = 0.6$, $\theta_D = 0.2$, $\sigma_\mu = 0.02$, $\sigma_\varepsilon = 2$. The policy functions are constructed setting $\theta = \theta_H$, $\varepsilon = 0$ and $\mu = 0.03$.

tions become more firmly anchored, allowing the central bank to stabilize inflation with a smaller fall in output. This explains why $\hat{y}_1$ eventually rises with ρ. This implies that the central bank can achieve better inflation and output outcomes today, as both $\hat{y}_1$ and π_1 get closer to the central bank's target when ρ increases toward 1.

Therefore, the value function of the central bank increases in ρ, as the bottom-right quadrant of the figure shows. This illustrates the idea of a "credibility dividend", as higher reputation allows the central bank to achieve better inflation and output outcomes when the economy is hit by a supply shock.

3.2 Quantifying the reputation channel

As discussed above, the strength of the reputation channel and its implications for macroeconomic outcomes depend on how responsive inflation expectations are to the policy decisions of the central bank, the term $\partial \Pi_2 (\hat{s}, \rho, \hat{y}) / \partial \hat{y}$ in equation (21). In this section, we discuss how we can identify this term using data on nominal interest rates and inflation expectations.

Let $\iota (\hat{y}; \hat{s}, \rho)$ to the interest rate that implements a level of output $\hat{y}$ given $(\hat{s}, \rho)$, obtained from the Euler equation (18). We can then define a monetary surprise to be

$$\Delta \mathbb{E} i (\hat{s}, \rho, \hat{y}) = \iota (\hat{y}; \hat{s}, \rho) - \left[\rho \int \iota \left(\hat{y}_1 (\hat{s}, \rho, \alpha(\theta_H, \varepsilon)) \right) \phi (\varepsilon) \, d\varepsilon + (1 - \rho) \int \hat{y}_1 (\hat{s}, \rho, \alpha(\theta_D, \varepsilon)) \, \phi (\varepsilon) \, d\varepsilon \right].$$

That is, $\Delta\mathbb{E}i$ is the difference between the realized nominal interest rate that implements $\hat{y}$ and the ex-ante expectations of the private sector given the state of the economy $\hat{s}$ and the prior ρ. Similarly, we can denote by $\Delta\mathbb{E}\pi_2(\hat{s},\rho,\hat{y})$ the update in expected second period inflation after the private sector observes the nominal interest rate chosen by the central bank.

The ratio of these two objects defines the elasticity of inflation expectations to monetary surprises: the percent change in expected future inflation following a 1% surprise in nominal interest rates. The following proposition characterizes this elasticity in the simplified economy of Section 3.1.

Proposition 3. *The elasticity of expected long-run inflation to a monetary policy surprise is*

$$b_2(\hat{s},\rho,\hat{y}) \equiv \frac{\Delta\mathbb{E}\pi_2(\hat{s},\rho,\hat{y})}{\Delta\mathbb{E}i(\hat{s},\rho,\hat{y})} \approx -\frac{\Pi_2'(\hat{s},\rho,\bar{y})}{\sigma - \left(1-\sigma\frac{\alpha\kappa^2}{(1-\alpha)}\right)\Pi_2'(\hat{s},\rho,\bar{y})}, \tag{30}$$

where $\Pi_2'(\hat{s},\rho,\bar{y}) \equiv \frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\Big|_{\hat{y}=\bar{y}}$ *and* $\bar{y}$ *is such that* $\mathbb{E}(\pi_2|\hat{s},\rho) = \Pi_2(\hat{s},\rho,\bar{y})$.

Equation (30) shows that there is a tight connection in our model between the elasticity of inflation expectations to monetary policy surprises and $\partial\Pi_2(\hat{s},\rho,\hat{y})/\partial\hat{y}$. Specifically, $b_2(\hat{s},\rho,\hat{y})$ is increasing (in absolute value) in the size of $\partial\Pi_2/\partial\hat{y} \geq 0$, and it equals zero when there is no uncertainty about the policymaker's type, $\rho \in \{0,1\}$. This result is important because it links an inherently unobservable object—the reputational wedge relative to the perfect information benchmark—to a potentially observable one: the sensitivity of inflation expectations to monetary surprises.

A large literature in monetary economics has focused on identifying monetary surprises and their effects on the economy via high-frequency identification. This result shows that these statistics—and specifically the response of expected inflation to a monetary surprise—provide important information on $\partial\Pi_2(\hat{s},\rho,\hat{y})/\partial\hat{y}$, the key term that regulates the role of reputation in our economy. In our quantitative analysis, we will build on this insight by first estimating $b_2(\hat{s},\rho,\hat{y})$ using high-frequency data and then using this estimate as empirical target in an indirect inference procedure to discipline the structural parameters of the model.

3.3 Discussion

So far we have described how reputation shapes monetary policy decisions and how this mechanism can be disciplined using data on inflation expectations and interest rates. For

that purpose, we have made two simplifications to our model. First, we have focused on an environment in which the policymaker's type is revealed after one period, which allowed us to study the equilibrium in a tractable "two-period" framework. Second, we have considered a special version in which there is no inflation bias—something that allowed us to log-linearize the private sector equilibrium conditions around the zero-inflation steady state. In what follows, we discuss how relaxing these two assumptions affects the key insights of this section.

Starting with the first point, we can show that the trade-off faced by the central bank in the infinite horizon economy is very similar to the one we studied in Section 3.1. To see that, consider the environment described in this section but assume that the central bank's type is not revealed in period 2, and that the type is not perfectly persistent.

The first order condition of the central bank is

$$-\left(R_{\pi}\kappa + R_{\hat{y}}\right) = \beta\left\{R_{\pi}\frac{\partial \mathbb{E}^{\rho'}\left[\pi(\hat{s}',\rho',\theta',\varepsilon')\right]}{\partial\rho'} + \mathbb{E}^{cb}\left[\frac{\partial V(\hat{s}',\rho',\theta',\varepsilon')}{\partial\rho'}\right]\right\}\frac{\partial p\left(\hat{s},\rho,\hat{y}\right)}{\partial\hat{y}}. \quad (31)$$

If we were to ignore reputation-building considerations, the optimal choice of inflation and output would be such that $R_{\pi}\kappa = -R_{y}$, as in equation (21). With imperfect information, instead, there is a wedge between these two terms, represented by the left-hand side of equation (31). This wedge has two components, shown on the right-hand side of the above expression.

The first component captures the fact that a more hawkish stance allows the central bank to moderate the impact of supply shocks on *current* inflation through its effect on inflation expectation, the term $\partial \mathbb{E}^{\rho'}\left[\pi'\right]/\partial\rho'$. This is the core mechanism discussed in simplified model of Section 3.1.

In addition to this term, there is a second motive to distort the static first-order condition: by acquiring reputation, the central bank can improve the inflation-output trade-off in the *future*, an effect that is summarized by the term $\mathbb{E}^{cb}\left[\partial V'/\partial\rho'\right]$. This second term did not appear in the simplified model because, in that example, the central bank's type was revealed in period 2 and so reputation had no value from that point onward.

Applying the envelope theorem to a log-linearized version of problem (15) we have

$$\frac{\partial V(\hat{s},\rho,\theta,\varepsilon)]}{\partial\rho} = R_{\pi}\beta\frac{\partial \mathbb{E}^{\rho'}\left[\pi'|\hat{s}\right]}{\partial\rho'}\frac{\partial\rho'}{\partial\rho} + \beta\mathbb{E}^{cb}\frac{\partial V\left(\hat{s}',\rho',\theta',\varepsilon'\right)}{\partial\rho'}\frac{\partial\rho'}{\partial\rho}.$$

Thus, we can then rewrite equation (31) as

$$-\left(R_{\pi,t}\kappa + R_{y,t}\right) = \sum_{k=1}^{\infty} \beta^k \mathbb{E}_t^{cb}\left[R_{\pi,t+k-1}\frac{\partial \mathbb{E}^{\rho_{t+k}}\pi_{t+k}}{\partial \rho_{t+k}}\frac{\partial \rho_{t+k}}{\partial \rho_t}\right]\frac{\partial \rho_t}{\partial \hat{y}_t} \tag{32}$$

where $k \geq 1$ indexes decisions k periods ahead, $\partial\rho_k/\partial\rho_1 = \partial\rho_k/\partial\rho_{k-1} \times \partial\rho_{k-1}/\partial\rho_1$ for all $k \geq 1$.

Comparing equation (32) with equation (26), we can see that trade-off that the central bank faces in the infinite horizon economy essentially is the same as the one of the two-period model of Section 3.1—with the exception that in the former higher reputation has the benefit of stabilizing the present discounted value of inflation expectations in the far future, and not just those next period.

Regarding the second point, in this section we have considered an economy in which the two types of central banks implement the same level of inflation in steady state. The full model, instead, allows the output target in the central bank's objective function to be above the flexible-price one. Under that parametrization, the model would generate positive levels of inflation in steady state. In addition, average inflation would be higher for the dovish policymaker than for the hawkish one, as the former places more weight on keeping output close to its target than the latter.

This feature does not conceptually change the trade-off faced by the central bank. However, it gives the model more flexibility to match the potentially large long-run elasticities of inflation expectations to monetary policy surprises. To see that, note that the term $\partial\Pi_2\left(\hat{s},\rho,\hat{y}\right)/\partial\hat{y}$ in the model of Section 3.1 is closely tied to the supply shocks, as it equals zero when $\hat{\mu} = 0$ or when the persistence of the supply shock goes to zero. In the fully fledged model, we will instead be able to obtain a sizable elasticity irrespective of the stochastic process for the supply shock.

As we will see in Section 5, this flexibility will be important to match the evidence on the sensitivity of inflation expectations to monetary surprises in Brazil.

4 Monetary policy and inflation expectations

The model predicts that the responsiveness of long-run inflation expectations to policy surprises—the elasticity $b(\hat{s},\rho,\hat{y})$ —is informative about the strength of the reputation channel. We now estimate this elasticity empirically. Section 4.1 describes the data and our empirical specification, while Section 4.2 reports our estimates for a group of advanced and emerging market economies.

4.1 Data and empirical specification

We now estimate the elasticity of inflation expectations to monetary policy surprises. We follow Cook and Hahn (1989), Kuttner (2001) and Cochrane and Piazzesi (2002), among many others and identify monetary policy surprises as variation in short term nominal interest rates in a window around a monetary policy meeting. We measure monetary surprises as the change in short-term interest rate swap rates linked to the key interbank rate in a two-day window around policy announcements, balancing the need to avoid confounding news with the need to give markets sufficient time to react and process new information.

We use financial market data on inflation compensation—the difference between the yields on nominal and inflation-linked bonds—or the inflation-linked swaps rates as a proxy for inflation expectations. The advantage of these financial market measures is that they are available at a daily frequency and reflect the views of investors who have strong incentives to be accurate, as they are effectively placing bets on their inflation outlooks. In addition, the term structure of these instruments allows us to examine how monetary policy surprises affect inflation expectations at different forward horizons. The disadvantage is that the financial market measures reflect both inflation expectations as well as compensation for holding inflation risk and differential liquidity premia for holding nominal and inflation protected-bonds. These risk and liquidity premia can be large and volatile, especially for countries with a history of high and volatile inflation. To guard against the possibility that inflation risk premia are driving our results, we also use survey-based measures of inflation expectations. These survey-based measures are generally considered free from inflation risk premia. However, their lower frequency—typically monthly or quarterly—reduces the statistical power of our event study approach, as revisions to expectations may reflect a wide range of other shocks. Moreover, unlike financial market participants, survey respondents often lack strong incentives to invest in forecast accuracy.

We collect these variables for a range of advanced and emerging market economies, see Appendix D.1 for a detailed discussion of data sources and definitions. A key data limitation is that relatively few countries have well-developed markets for inflation-protected bonds or inflation-linked swaps, which restricts the availability of high-frequency market-based measures. Among advanced economies, we include the Euro Area (EA), the United Kingdom (UK), and the United States (US); among emerging markets, we focus on Brazil, Chile, and Mexico. For all these countries, we obtain a consistent sample going from 2010 up to 2024.[17] For each country, we estimate

[17]We drop observations between January and August 2020 from the analysis due to the market turmoils associated with the global outbreak of Covid-19.

$$\Delta \mathbb{E}_t[\pi_{t+k,t+k+s}] = a + b_{(k,s)} \Delta i_t + \eta_t, \tag{33}$$

where $\Delta \mathbb{E}_t[\pi_{k,k+s}]$ is the revision in expected average inflation between k and $k+s$ periods around the monetary policy meeting and Δi_t is the monetary policy surprise. Both variables are expressed in annualized percentages, so the coefficient $b_{(k,s)}$ measures the elasticity of inflation expectations to monetary policy surprises at different forecasting horizons.

There are two remarks that are important to make at this stage. First, as we discussed in Section 3.2, the elasticity in our model is state-dependent and varies over time. In the data, instead, we identify just one value for $b_{(k,s)}$ for the forward horizon $[k, k+s]$, which we can interpret as an average elasticity over the sample period of analysis.

Second, in our benchmark specification we will focus on a "long-run" measure of inflation expectations, five to ten years ahead. The reasons why we focus our analysis on the long-run is twofold. First, inflation expectations may respond to monetary surprises for reasons that are different from the reputation motives discussed in this paper, and this may confound our measurement. For example, in a New Keynesian model with capital, interest rate changes will affect capital accumulation and, by changing productive capacity, will have an impact on expectations of future inflation. To the extent that reputational effects are more persistent than these propagation mechanisms, looking at the long run will allow us to better isolate the former. Second, as equation (32) shows, in our economy the monetary authority cares about the present discounted value of the elasticity of inflation expectations to policy actions. Detecting a sizable long-run elasticity will then be a strong signal that this object is large—as we expect the elasticity to be continuous in k.

4.2 Results

Table 1 reports estimates of $b_{(k,s)}$ from equation (32) for the different countries. Negative coefficients indicate that monetary tightening lowers long-term inflation expectations. Panel A reports the information when using inflation compensation as a measure of long-run inflation expectations, while Panel B reports the results when using surveys.

We can see that long-run inflation expectations are very sensitive to monetary policy surprises for the Brazilian economy, with unexpected interest rate hikes being associated with significant reductions in expectations about inflation in the long-run. These effects are statistically significant at the conventional levels, and they are economically sizable: a 100 basis points surprise in the three month interest rate swap rate causes, on average, a reduction of roughly 50 basis points when using inflation compensation and approximately

Table 1: Estimates of equation (33)

Panel A: Market-based results

	Brazil	Chile	Mexico	Euro area	UK	US
$b_{(5y5y)}$	-0.48^{***}	-0.05	0.02	0.03	-0.05	-0.09
	(0.18)	(0.05)	(0.07)	(0.06)	(0.07)	(0.09)
R^2	0.08	0.03	0.00	0.00	0.01	0.00
Sample	2010-2024	2010-2023	2010-2024	2010-2024	2010-2024	2010-2024
# obs.	92	104	109	133	143	114

Panel B: Survey-based results

	Brazil	Chile	Mexico	Euro area	UK	US
b	-0.18^{**}	-0.00	0.03	0.11	0.16	-0.04
	(0.09)	(0.03)	(0.06)	(0.08)	(0.09)	(0.06)
R^2	0.22	0.00	0.00	0.05	0.01	0.00
Sample	2010-2024	2010-2023	2010-2024	2010-2024	2010-2024	2010-2024
# obs.	92	117	106	56	47	54

Notes: The table reports regression coefficients from regressions of long-horizon inflation expectations onto monetary policy shocks, equation (33). Panel A reports results for market-based measures of inflation expectations based on either the difference between nominal and real bond yields or based on inflation-linked swap rates. The market-based inflation expectations are 5yr-5yr forward rates. Panel B reports results for survey-based measures of inflation expectations. See Appendix D.1 for details about the exact horizons of the survey-based inflation expectations for the different countries.

20 basis points when using surveys.[18] For the other countries, we do not find robust evidence that long-run inflation expectations respond to monetary policy changes. The coefficient b for those countries, while often negative, is not significantly different from zero at conventional levels, and the R^2 of these regressions are very close to zero in almost all cases. These results hold regardless of the measure of long-run inflation expectations used in the analysis—whether based on financial markets or on survey data.

This result can be further illustrated by plotting the estimates for $b_{(k,s)}$ at different horizons for Brazil and the US. This experiment is reported in Figure 3. Following a 100 basis point surprise in the OIS rate in the US, we observe a hump-shaped response for inflation expectations: they are roughly insensitive in the near term ([0,2] years), they are significantly negative in the medium term, declining by 20 basis points over the [2,5] years horizon, and negative but not significantly different from zero at conventional levels in

[18] These difference in the point estimates could be due to multiple factors. First, market-based measures of inflation expectations contain a risk-premium component that may itself be sensitive to monetary policy surprises, and this could drive a wedge between the market-based measures and the surveys. Second, it is worth pointing out that not all participants in the Focus survey update their forecasts immediately after a monetary policy meeting. This unresponsiveness bias downward the sensitivity of inflation expectations to monetary surprises.

the longer run. These results are consistent with previous work that has focused on the relationships between monetary surprises and market-based inflation expectations data. For example, Hanson and Stein (2015) and Nakamura and Steinsson (2018) find a negative but statistically insignificant effect of monetary surprises on inflation break-evens five to ten years out, while Acosta, Ajello, Bauer, Loria, and Miranda-Agrippino (2025) recently document an hump-shaped pattern, with break-even inflation being more responsive to monetary surprises at intermediate maturities. In the case of Brazil, instead, a surprise tightening is on average associated with a very persistent and sizable drop in inflation expectations, with little evidence of mean reversion.

Why do inflation expectations in Brazil respond differently to a monetary shock relative to other countries? Through the lens of our model, one plausible interpretation is that bond market investors and professional forecasters during the sample period of analysis are less certain about the monetary policy framework of Brazil than they are for the other countries in our analysis. As we have seen in Section 3.1, the elasticity of inflation expectations to changes in interest rates depend on the underlying level of reputation ρ. Countries with high levels of reputation have more anchored inflation expectations relative to a country with intermediate levels of reputation. This interpretation is consistent with narrative evidence in Bonomo et al. (2024), which documented major shifts in monetary policy in Brazil over the period of analysis; to the extent that these policy shifts have raised uncertainty about the policy objectives of the central bank, they may have made reputational factors more relevant in shaping long-run inflation expectations.

While the estimated long-run elasticities are close to zero for the other countries, this should not be interpreted as evidence that reputational forces are unimportant in those contexts. In fact, for the United States we find that monetary policy surprises affect inflation expectations at medium-term horizons, and our model suggests that reputation forces would be relevant even if long-run expectations are firmly anchored. See also the complementary evidence in Caravello et al. (2025) regarding this point. Rather, the logic of our model suggests that these reputational motives should be stronger in environments such as Brazil, where monetary surprises have a larger and more persistent effect on inflation expectations.

One important concern about the results is that they may be affected by the so-called central bank information effect (see Nakamura and Steinsson, 2018). The issue is that, when the central banks announce policy rate decisions, they may at the same time reveal "private" information about the state of the economy, which investors may take into account and update their own assessment of the economy. This could bias our results, because they rely on the changes in interest rate swap rates around monetary events to solely reflect

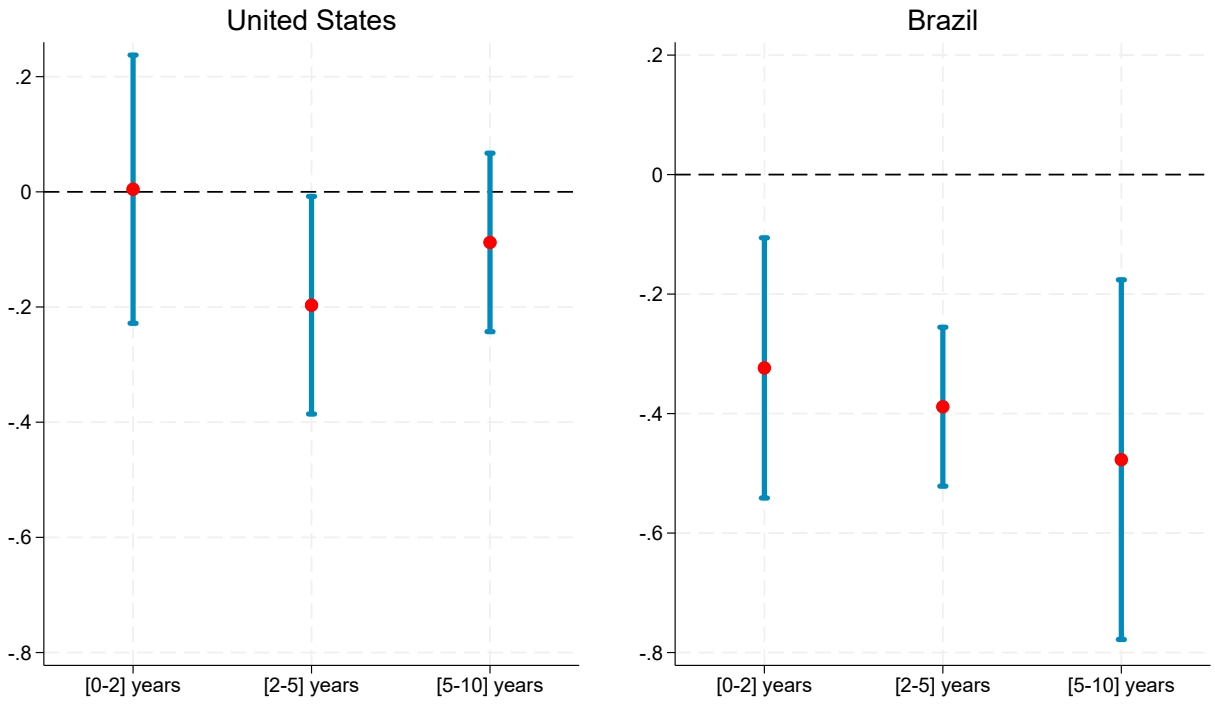


Figure 3: Monetary shocks and inflation expectations: US vs. Brazil

Notes: The figure reports estimates of $b_{(k,s)}$(point estimates and 90% confidence intervals) in (33) for various horizons for the US and Brazil. The confidence intervals are constructed using robust standard errors. Inflation expectations are market-based and calculated from the difference between nominal and real bond yields. The sample covers 2010 through 2024.

monetary policy surprises.

We think this issue is less relevant in our setting because we are studying the response of inflation expectations at far-forward horizon (five to ten years out), which should not be particularly sensitive to information about the current state of the economy. To verify this intuition, we follow Jarociński and Karadi (2020) and control for the central bank information effect by focusing on days where interest rate swap rates and stock prices move in opposite direction. The results are available in Appendix D.1. The key take away from the exercise is that our empirical results are robust to this modification.

5 Quantitative analysis

Having documented that long-run inflation expectations in Brazil respond strongly to monetary policy surprises, we now use the model to quantify how this sensitivity affects monetary policy decisions. The goal is to translate the empirical elasticity estimated in Section

4 into an economically meaningful measure of the reputational forces that arises when the central bank acts to manage expectations. To do so, we calibrate the model to Brazilian data, choosing parameters so that the model matches, among other targets, the implied elasticity of expected inflation to interest-rate surprises matches we measured.

The analysis proceeds in three steps. First, in Section 5.1, we calibrate the model to the Brazilian economy Second, in Section 5.2, we use our framework to measure the importance of reputation building motives in the conduct of monetary policy by comparing equilibrium outcomes in our economy to those of an economy in which the central bank is myopic. Finally, in Section 5.3, we use the model to quantify how reputation affects the inflation–output trade-off faced by the central bank.

5.1 Calibration and model fit

We calibrate the model at a quarterly frequency using Brazilian data. We define the central bank output target to be proportional to the efficient level of output, $y_t^* = \delta \left(z_t^{1+1/\nu} / \chi \right)^{1/(\sigma+1/\nu)}$. We let $\alpha(\theta, \varepsilon) = \theta / \left[\theta + (1-\theta) e^{-\varepsilon} \right]$, and we assume that the Markov chain governing the transition between types is symmetric, with P_θ denoting the probability of remaining in a regime.

The model parameters are $\left[\phi, \sigma, \nu, \beta, \bar{\mu}, \chi, \rho_z, \sigma_z, \rho_\mu, \sigma_\mu, \delta, \theta_H, \theta_D, P_\theta, \sigma_\varepsilon \right]$. We fix a subset of these parameters to standard values. We set the intertemporal elasticity of substitution and the Frisch elasticity of labor supply to 1, the discount rate to 0.99, $\bar{\mu} = 1.2$, and set $\chi = 1/\bar{\mu}$, so that the flexible-price level of output is normalized to 1 in steady state. Given these parameters, we choose ϕ so that the slope of the log-linearized Phillips curve, $\kappa = \frac{\sigma + \frac{1}{\nu}}{\phi(\bar{\mu}-1)}$, is equal to 0.075, a value in line with estimates for the Brazilian economy after the adoption of inflation targeting (dos Santos Fernandes and Portugal, 2023). We set $\rho_z = 0.95$ and $\sigma_z = 0.0065$, consistent with previous business cycle research.

The remaining parameters, $\left[\rho_\mu, \sigma_\mu, \delta, \theta_H, \theta_D, P_{HH}, \sigma_\varepsilon \right]$, govern the stochastic process of the supply shock μ_t and of the objective function of the central bank. We choose these parameters so that the model replicates a set of moments as closely as possible.

The moments are computed using quarterly data on hp-filtered real gross domestic product per capita (in logs), annualized inflation and the annualized policy rate controlled by the Brazilian central bank (Selic rate)—see Appendix D.1 for data sources and definitions. Our sample goes from 2010:q1 to 2019:q4. We target the sample standard deviation, first-order auto-correlation and cross-correlation of these three variables. As argued in Bonomo et al. (2024), this period was characterized by two different "regimes" for monetary policy: a regime in which the government influenced the central bank into maintaining low in-

Table 2: Model parameters

Parameter	Value	Note
ϕ	133.333	Adjustment cost, inflation
$1/\sigma$	1.000	Intertemporal elasticity of substitution
ν	1.000	Frisch elasticity of labor supply
β	0.990	Discount factor
$\bar{\mu}$	1.200	Average markup
χ	0.833	Disutility of labor
ρ_z	0.950	Persistence, productivity shocks
σ_z	0.065	Standard deviation, productivity shocks
ρ_μ	0.900	Persistence, supply shocks
σ_μ	0.013	Standard deviation, supply shocks
δ	0.706	Output target
θ_H	0.500	Weight on inflation, Hawk
θ_D	0.150	Weight on inflation, Dove
P_θ	0.995	Probability of remaining in a policy regime
σ_ε	2.400	Standard deviation, monetary shocks

terest rates despite inflation running persistently above target, from 2010:q1 up to 2016:q2; and a regime in which the central bank pursued stricter policies to make sure inflation returned to target, from 2016:q3 onward. We identify the former as a "dovish" regime and the latter as a "hawkish" one, and we include in the set of empirical targets average inflation and nominal interest rates over these two sub-samples. Finally, we include the sensitivity of long-run inflation expectations to monetary surprises estimated in Section 4.2—specifically, the one computed using the Focus survey at a four year forecasting horizon (panel B of Table 1). It is worth noting that the model is over-identified, since we are fitting 13 moments with only 7 parameters.

We compute the model implied moments via a long simulation ($T = 150,000$). The standard deviations, auto-correlations and cross-correlation are computed at the ergodic distribution. We compute average inflation and nominal interest rates when the private sector is almost certain to be facing the Dove ($\rho_t \leq 0.05$) and similarly when it is almost certain to be facing the Hawk ($\rho_t \geq 0.95$), and equate those to average inflation and nominal interest rates in the sub-samples.[19]

Finally, we follow the discussion in Section 3.1 and compute the elasticity of long-run inflation expectations to monetary surprises, for each period t, as $b_t^{(k)} \equiv \Delta\mathbb{E}_t[\pi_{t+k}]/\Delta\mathbb{E}_t[i_t]$, where $\Delta\mathbb{E}_t[x_{t+k}]$ denotes the revision in the private sector expectations of variable x in

[19]Results do not change meaningfully if we make the cut-offs equal to 0.1 and 0.9, respectively, or if we use the actual realization of the types rather than reputation to compute the conditional means.

period $t+k$ after agents observe the interest rate chosen by the monetary authority at time t. In our model $b_t^{(k)}$ is state-dependent and varies over time, while in our empirical analysis we obtain only an average elasticity over the sample of analysis. Given that this period was characterized by large monetary policy shifts in Brazil, we compute the model-implied counterpart as the average elasticity conditioning on high policy uncertainty, which we define to be periods in which $\rho_t \in (0.4, 0.6)$. Consistent with the data, we consider a forward horizon of 16 quarters.

Table 2 reports the value of the parameters while Table 3 reports the empirical and model implied moments. Despite being over-identified, the model captures remarkably well the behavior of output, inflation and nominal interest rates in the sample. With a few exceptions, the model matches accurately the moments that describe the volatility, auto-correlation and cross-correlation patterns of these three variables, and it reproduces well the average behavior of inflation and nominal interest rates over periods of low and high reputation. In our simulations, average inflation and nominal interest rates are, respectively, 6.45% and 10.73% in periods in which the central bank has a low reputation, and they are 4.60% and 7.06% when reputation is high—numbers comparable to the Brazilian experience over the sample.

Importantly, the model produces an empirically plausible elasticity of long-term inflation expectation to monetary innovations, with inflation expectations three-years ahead falling on average -0.15% after a 1% unexpected increase in nominal interest rates. This average value masks a sizable degree of state-dependence in the model. Figure 4 explores this aspect. The left panel of the figure plots the distribution of $b_t^{(16)}$ in our simulations. While the elasticity typically hovers around zero, there is a long left tail, indicating that the model can create episodes in which long-run inflation expectations are at risk of de-anchoring.

The right panel of the figure shows that policy uncertainty is a key driver of this state-dependence. The illustration plots the conditional mean of $b_t^{(16)}$ against ρ_t. We can see that the elasticity on average is close to zero when either ρ_t is close to zero or when it is close to one: in those instances, the private sector has little uncertainty about the policymaker's type, so innovations to nominal interest rates have little effects on long-run inflation expectations. Long-run inflation expectations are, instead, more sensitive to unexpected changes in nominal interest rates when the private sector is more uncertain about the policymaker's type, that is when ρ_t takes intermediate values.

The model produces a sizable elasticity of long-run inflation expectations to monetary innovations via a combination of a slow learning process and different underlying levels of inflation for the two types. To understand why, we can rewrite the expectations of inflation

Table 3: Output, inflation and nominal interest rates: model vs data

Moment	Data	Model
Stdev$(\log Y_t)$	1.95	2.02
Stdev(π_t)	2.01	1.82
Stdev(i_t)	2.78	3.27
Acorr$(\log Y_t)$	0.90	0.84
Acorr(π_t)	0.92	0.91
Acorr(i_t)	0.95	0.45
Corr$(\log Y_t, \pi_t)$	-0.23	0.00
Corr$(\log Y_t, i_t)$	-0.45	-0.39
Corr(π_t, i_t)	0.81	0.54
Mean$(\pi_t \vert \text{Dove})$	6.68	6.45
Mean$(i_t \vert \text{Dove})$	10.83	10.73
Mean$(\pi_t \vert \text{Hawk})$	4.23	4.60
Mean$(i_t \vert \text{Hawk})$	8.42	7.06
$b^{(16)}$	-0.18	-0.15

Notes: The sample standard deviations, autocorrelations and cross-correlation are computed over the entire sample period, 2010:q1-2019:q4. Average inflation and nominal interest rates are computed over the 2010:q1-2016:q2 sub-sample for the dovish regime, and the 2016:q3-2019:q4 sub-sample for the hawkish regime. The moments in the model are computed on a long simulation, T=150,000. Inflation and nominal interest rates are expressed in annualized percent, while output Y_t is reported in logs and multiplied by 100. Standard deviations, autocorrelations and cross-correlations are computed using all observations. Average inflation and nominal interest rates are computed conditioning on $\rho_t \leq 0.05$ for the dovish regime and $\rho_t \geq 0.95$ for the hawkish regime. The average elasticity in the model, $b^{(16)}$, is computed conditioning on $\rho_t \in (0.4, 0.6)$.

k period forward as

$$\mathbb{E}_t[\pi_{t+k}] = \Pr{}_t(\theta_{t+k} = \theta_H)\mathbb{E}_t[\pi_{t+k}|\theta_{t+k} = \theta_H] + [1 - \Pr{}_t(\theta_{t+k} = \theta_D)]\,\mathbb{E}_t[\pi_{t+k}|\theta_{t+k} = \theta_D]$$

First, in our calibration, the monetary authority targets a level of output above the flexible-price one, $\delta > \bar{\mu}^{-1/(\sigma+1/\nu)}$. This implies that the two types of policymakers target different underlying levels of inflation: the Dove, who in our calibration places on average a higher weight on output stabilization, tolerates higher levels of inflation than the Hawk, as it tries to push output closer to the target. This feature implies that $\mathbb{E}_t[\pi_{t+k}|\theta_{t+k}]$ varies with θ_{t+k} even for large k. If δ was equal to $\bar{\mu}^{-1/(\sigma+1/\nu)}$– as in our simplified model in Section 3 this conditional expectation would converge to the unconditional mean as k grows, at the same speed as the convergence of the supply shock. Given our estimates for ρ_μ, the half-life of a supply shock is approximately 6.5 quarters. Thus, without an inflation bias, the model cannot produce a lasting effect of reputations on long-term (16 quarters) inflation expectations..

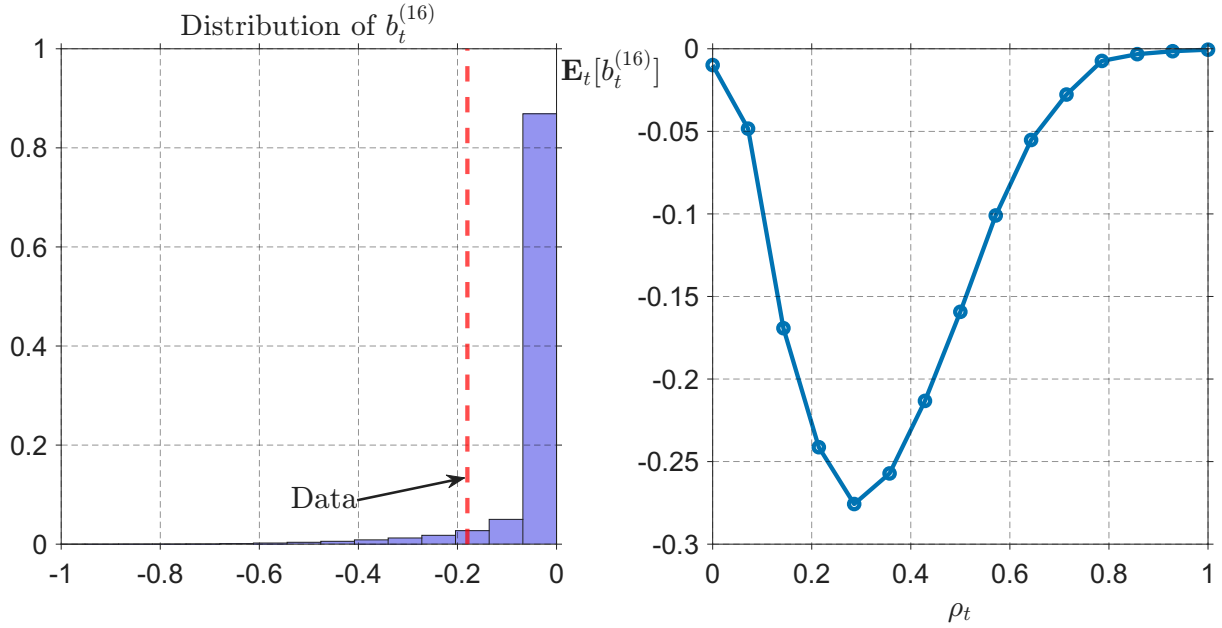


Figure 4: Model-implied distribution of the elasticity of long-run inflation expectations to monetary policy surprises

Notes: We simulate our economy for $T = 150,000$ periods and for each t we compute $b_t^{(16)}$ as explained in the text. The left panel reports the histogram of $b_t^{(16)}$ over this simulation while the vertical dotted line depicts the estimated elasticity in the data. The right panel reports the mean of $b_t^{(16)}$ conditioning on different reputation levels ρ, $\mathbb{E}\left[b_t^{(16)} | \rho_t \in (\rho - a, \rho + a)\right]$, with $a = 0.05$.

Second, in our calibration θ is very persistent ($P_\theta = 0.995$) and monetary innovations are large, $\sigma_\varepsilon = 2.4$. Because of this, types switch infrequently and agents learn slowly over time, so a sufficiently large monetary innovation can have long-lasting effects on reputation—that is, it can meaningfully move $\Pr_t(\theta_{t+k} = \theta_H)$ for k large.[20]

The combination of these two features—a slow learning process and different underlying inflation levels for the two types—imply that long-run inflation expectations in our model can move realistically. This effect is more important when the private sector is uncertain about the policymaker's type, as in those state beliefs are more responsive.

5.2 Quantifying the reputation channel

We now assess the importance of reputation considerations when setting monetary policy. For this purpose, we follow the analysis of Section 3 and compare the equilibrium outcome

[20] Both of these ingredients are necessary. If the Markov process was not very persistent, forward expectations would quickly converge to the unconditional mean. If the monetary shocks were not sizable, agents would learn the policymaker's type quickly over time, and forecast revisions would not affect expectations of inflation far in the future.

Table 4: The reputation channel: low vs. high policy uncertainty

	Benchmark	Myopic	Benchmark	Myopic
Moment	High uncertainty		Low uncertainty	
$\text{Mean}(\pi_t)$	4.42	5.61	5.12	5.30
$\text{Stdev}(\pi_t)$	1.54	1.83	1.83	1.95
$\text{Mean}(\hat{y}_t)$	-0.84	-0.41	-0.54	-0.30
$\text{Stdev}(\hat{y}_t)$	0.68	0.13	0.70	0.13
$\text{Mean}\left(\frac{R_t^{bench}-R_t}{\text{abs}(R_t^{bench})}\right)$	0.00	-0.37	0.00	-0.13

Notes: We simulate the benchmark economy and the one with a myopic central bank for $T = 150,000$ periods. The first two columns report statistics conditioning on $\rho_t \in (0.4, 0.6)$, while the last two columns condition on $\rho_t \notin (0.4, 0.6)$. Inflation is reported in annualized percentages while output is reported in percentage deviation from the central bank's target.

of our economy to one in which the central bank is myopic. The expectation formation process is the same as in the benchmark economy: households and firms observe monetary policy decisions and update their prior on the policymaker's type using Bayes' rule. The only difference is that the myopic central bank does not internalize the effects that monetary policy decisions have on reputation, while these are fully internalized in our model.

Table 4 reports the mean and standard deviation for inflation and output, the latter expressed in percentage deviation from the central bank's target, $\hat{y}_t = (y_t - y_t^*)/y_t^*$. The first two columns report these moments conditioning on $\rho_t \in (0.4, 0.6)$ ("High uncertainty") and on $\rho_t \notin (0.4, 0.6)$ ("Low uncertainty"): as we showed in Proposition 2, the difference in outcomes between these two economies depend on how sensitive inflation expectations are, and the analysis in the previous sub-section has showed that the current level of reputation is the key determinant of this sensitivity in our calibration. The table also reports the average value of the policymaker's objective function R_t in the myopic economy relative to the one in the benchmark.

Consistent with the discussion of the simplified model of Section 3, reputation building motives induce the central bank to react more to inflationary shocks, resulting in a less volatile inflation process. This comes at the costs of larger deviations of output from the central bank's target. These differences can be quite sizable, especially in periods of high uncertainty: the standard deviation of inflation is 12% larger in the economy in the economy with a myopic central bank, while output is approximately five times more volatile in the benchmark than in the myopic economy.

In addition to these differences in volatilities, the two economies also differ in their average behavior due to the inflation bias. When uncertainty is high, the benchmark economy

features substantially lower inflation on average than the economy with a myopic central bank, 4.42% vs 5.61%. In the benchmark economy, the central bank has an incentive to build a hawkish reputation and, due to this incentive, it accumulates reputation over time. This moderates inflation on average because higher reputation anchors expectations toward those of the hawkish policymaker, which in our calibration implements a lower inflation rate on average.

When aggregating those numbers using the policymaker's objective function, we can see that reputation building motives lead to considerable gains for the central banks: the average value of the period objective function R_t in the myopic economy is 37% below that of the benchmark in periods of high uncertainty and 13% below in normal times.

An alternative way of quantifying the importance of the reputation channel is to compute the implicit slope of the Phillips curve that rationalizes the observed relation between output and inflation in a model without reputation building motives. That is, given a path $\{\hat{y}_t, \pi_t\}$ simulated from our model, we can compute

$$\kappa_t = \frac{(1-\alpha_t)\hat{y}_t}{\alpha_t \pi_t}.$$

From the analysis in Section 3.1, we know that κ_t will be larger than the actual slope of the Phillips curve, κ, as the reputation channel pushes the central bank to favor inflation stabilization over output stabilization. This exercise allows us to quantify by how much.

In our simulations, the median value of κ_t is 0.303 unconditionally, and 0.697 if we condition in periods during which uncertainty is high, $\rho_t \in (0.4, 0.6)$. These numbers are substantially larger than the actual slope of the Phillips, which is 0.075 in our calibration. That is, one would need to increase the slope of the Phillips curve by a factor of four (or by a factor of nine in periods of high uncertainty) to rationalize—in a model without reputation—the incentives to stabilize inflation due to the reputation channel.

5.3 The value of reputation

The previous experiment has showed that reputational considerations can have quantitatively important effects on monetary policy. As we have discussed in Section 3, policymakers have an incentive to build reputation because higher ρ_t moderates future inflation expectations and make them less responsive to economy shocks, something that helps the policymaker achieve a better inflation-output trade-off over time. We now use our model to quantify the value of reputation for the policymaker.

To this end, consider the expected value of the central bank's objective function next

period, $\mathbb{E}_t^{cb}[R_{t+1}|\theta, s, \rho]$ for type θ in state s. For P_θ large enough (as in our calibration), this expression can be approximated by

$$\begin{aligned}\mathbb{E}_t^{cb}[R_{t+1}|\theta, s_t, \rho_t] \approx & -\mathbb{E}_t\left\{[1-\alpha(\theta,\varepsilon_{t+1})]\,\mathbb{E}_t\left[\hat{y}_{t+1}|\varepsilon_{t+1}\right]^2 + \alpha(\theta,\varepsilon_{t+1})\mathbb{E}_t\left[\pi_{t+1}|\varepsilon_{t+1}\right]^2\right\} \\ & - \mathbb{E}_t\left\{[1-\alpha(\theta,\varepsilon_{t+1})]\,Var_t\left(\hat{y}_{t+1}|\varepsilon_{t+1}\right) + \alpha(\theta,\varepsilon_{t+1})Var_t\left[\pi_{t+1}|\varepsilon_{t+1}\right]\right\} \quad (34)\end{aligned}$$

where expectations are taken with respect to the shocks that are public information, s_{t+1}, and ε_{t+1}. The first term is related to the "inflation bias", as it capture the loss of value for the central bank that comes from trying to push output above the flexible-price level resulting in average inflation above zero.[21] The second term is related to the "stabilization bias", and it captures the loss in value arising from the inability of the central bank to keep inflation and output at their target in response to shocks.

The left panels of Figure 5 plot these conditional moments as a function of ρ. The top panels of the figure plot the conditional mean and variance of output—in deviation from the central bank's target—while the bottom panels report the same statistic for inflation. When constructing this figure, we set $\theta = \theta_H$ and $\hat{\mu} = 0$. As in the simplified model of Section 3.1, we can identify two regions in these plots, depending on ρ.

At low reputation levels, an increase in ρ is associated with lower output on average and with higher output volatility. This happens because the sensitivity of inflation expectations is increasing in ρ in this region of the state space—see Figure 4—and so do the incentives to build reputation. Therefore, as ρ increases, the central bank implements a more hawkish policy, setting higher real interest rates both on average and in response to inflationary supply shocks.

The benefits of these reputation gains can be seen mostly in the reduction of average inflation, which drops from 6.75% to 4.75% when reputation goes from 0 to 0.5. The volatility of inflation, instead, moderately increases in this region. The effect of reputation on inflation volatility is a priori ambiguous in this region of the state space. On the one hand, a higher reputation induces the central bank to respond more strongly to supply shocks, something that reduces the volatility of inflation. On the other hand, as we have discussed previously, the elasticity of inflation expectations to monetary surprises increases in this region of the state space, and this makes inflation more sensitive to monetary surprises. In our calibration, these two effects mostly offset each other, with the second effect being slightly stronger.

[21] This term would be equal to zero when $\hat{\mu} = 0$ if the output target of the central bank was equal to the flexible-price level of output ($\delta = \bar{\mu}^{-1/(\sigma+1/\nu)}$).

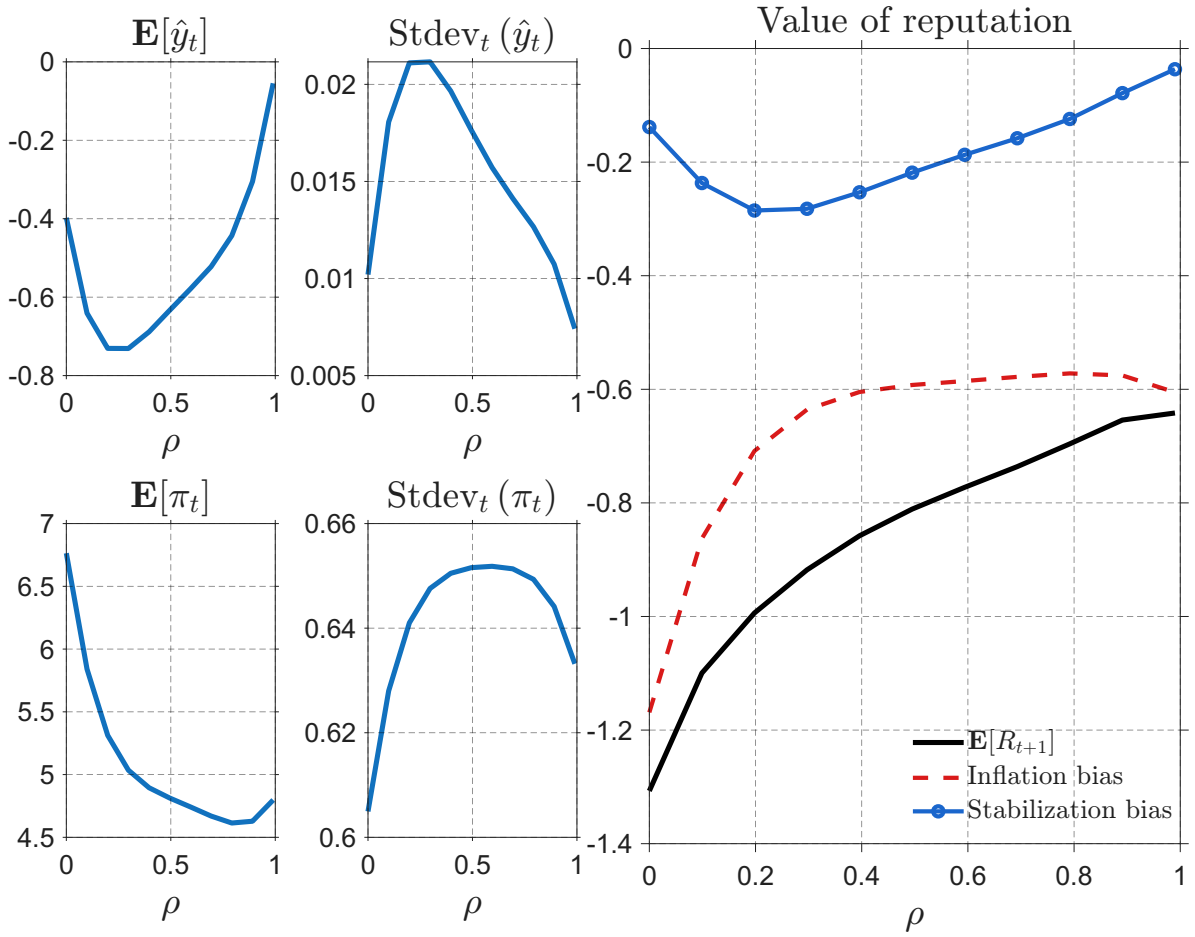


Figure 5: The value of reputation and its components

Notes: The panels on the left plot the conditional mean and standard deviation of output and inflation, for different levels of reputation. Inflation is reported in annualized percentages and output in percentage deviation from the central bank's target. These conditional moments are computed via simulations. We set $\theta = \theta_H$, $\hat{\mu} = 0$. Then, for each ρ and ε' in a grid of values, we simulate $K = 2000$ different realizations of the supply shock next period, and compute output and inflation for each of these realizations. The solid lines in those figures plots the mean and standard deviation across these different realizations as a function of ρ, for $\varepsilon' = 0$. The right panel aggregates these moments according to equation (34) and plots$\mathbb{E}_t[R_{t+1}]$ and its decomposition into the inflation bias component and the stabilization bias component.

At high reputation levels, an increase in ρ is associated with higher levels of output and with a reduction in its volatility. As the private sector becomes increasingly confident that the policymaker is a Hawk, the central bank's incentive to build reputation weakens. Additionally, inflation expectations become more firmly anchored at lower levels as ρ increases, which allows the central bank to achieve a better inflation-output trade-off. The combination of these two forces explains why the economy in this region is characterized by output and inflation that are increasingly close to the central bank's target and that are less sensitive to supply shocks.

The right panel of the figure aggregates these conditional moments to plot the objective function of the central bank and its decomposition in equation (34) as a function of ρ. The central bank achieves a better inflation-output trade-off as ρ goes from zero to 1. Initially, the gains come exclusively from the reduction in the inflation bias, as the stabilization bias component deteriorates. As the central bank graduates into a high-reputation institution, the inflation bias component stabilizes and the value of reputation is reflected in a better ability of the central bank to stabilize output and inflation against economic shocks.

Overall, our analysis uncovers a sizable *reputation dividend* for the central bank.

Two additional observations. Figure 5 highlights two additional points. First, it shows why inflation expectations alone are not a good measure of de-anchoring risk. Compare, for example, the outcomes in our economy at $\rho = 0.5$ and $\rho = 1$. From Figure 5, we can see that in these two cases inflation expectations are quite similar, mostly because the central bank at $\rho = 0.5$ optimally chooses a tighter policy due to strong reputation-building motives. Yet, the central bank is more at risk of a de-anchoring at $\rho = 0.5$ than it is at $\rho = 1$, in the sense that inflation expectations would increase much more drastically in the former case if the central bank unexpectedly implemented a more dovish policy. As we have argued in the paper, the elasticity of inflation expectations to monetary surprises is a more appropriate indicator of de-anchoring risk in this circumstance.

Second, the mechanisms underlying Figure 5 can help rationalize why different countries had different monetary policy responses to the post-pandemic inflation episode, and why we observed quite similar inflation outcomes despite these differences.[22] As documented by Nakamura et al. (2025), countries with a recent history of high inflation (a proxy for reputation in our model) tightened monetary policy much more aggressively and earlier, yet attained inflation paths similar to those of countries with a history of low inflation (high reputation). This outcome is consistent with our model.

To clarify this last point, Figure 6 plots the impulse response function to an inflationary supply shock under low reputation (solid line) and high reputation (dashed line), assuming that the central bank is a Hawk. With low reputation, strong incentives to build credibility lead the central bank to raise interest rates very aggressively, with nominal interest rates increasing twice as much relative to the high reputation case. Consistently, the low-reputation economy features a more severe output contraction than the high-reputation economy, as in the latter, the central bank can adjust interest rates more moderately since expectations are already well anchored.

[22] A model featuring private-sector learning about stochastically evolving policy-response coefficients would fail to account for these observations.

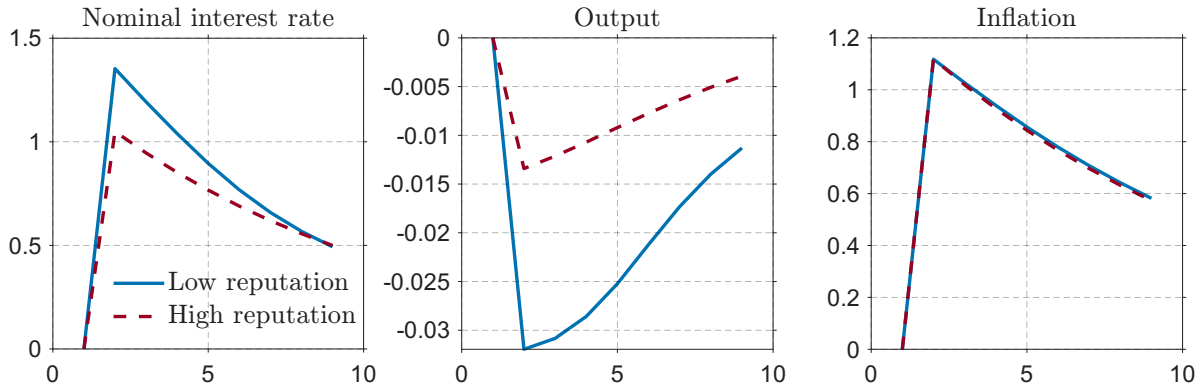


Figure 6: Impulse response functions to a supply shock

Notes:Impulse response functions are computed as follows. We set $\theta_t = \theta_H$ and $\varepsilon_t = 0$ for all t. For the low-reputation case, we set $\rho_1 = 0.25$ (solid line), while we set $\rho_1 = 1.00$ for the high-reputation case (dashed line). We consider two paths for the supply shock. In the first path, we set $\varepsilon_{\mu,t} = 2$ for $t = 2$ and zero otherwise. In the second path, we set $\varepsilon_{\mu,t} = 0$ for all t. The lines in the figure plot the difference in $\log(y_t/z_t)$, π_t and i_t between these two paths. Output is reported in percentages, while nominal interest rates and inflation are expressed in annualized percentage points.

6 Conclusion

This paper has studied monetary policy in an environment in which long-run inflation expectations are imperfectly anchored and can be influenced by the central bank's actions. We show that when agents are uncertain about the central bank's long-run objectives, policymakers face stronger incentives to respond aggressively to inflationary supply shocks in order to signal their commitment to price stability. This process of reputation-building helps stabilize long-run inflation expectations and improves the inflation-output trade-off.

We show theoretically that the strength of the reputation channel in the model is closely linked to the sensitivity of long-run inflation expectations to monetary surprises. We estimate these elasticities for a set of emerging and advanced economies, finding that they are especially large for Brazil. Through the lens of our model, this finding suggests large gains from building reputation—a core insight that can help explain why central banks with a recent history of high inflation quickly and aggressivly raised interest rates after the pandemic.

Our analysis takes as given the objective function of the central bank and how it varies over time, and points out that these movements introduce a non-trivial inference problem for the private sector. In practice, those movements are the result of delicate institutional interactions between the monetary and the fiscal authority, and how those interactions

change with economic conditions. In Bocola, Chaumont, Dovis, and Kirpalani (2025) we are currently studying these and related issues.

References

Acosta, M., A. Ajello, M. Bauer, F. Loria, and S. Miranda-Agrippino (2025): "The Financial Market Reaction to FOMC Communication: Evidence from a New Event-Study Database," Tech. rep., Unpublished manuscript. 4.2

Afrouzi, H., M. Halac, K. S. Rogoff, and P. Yared (2023): "Monetary policy without commitment," Tech. rep., National Bureau of Economic Research. 4

Amador, M. and C. Phelan (2024): "Central Bank Reputation With Noise," . 1

Armenter, R. (2018): "The perils of nominal targets," *The Review of Economic Studies*, 85, 50–86. 13

Backus, D. and J. Driffill (1985): "Inflation and reputation," *The American Economic Review*, 75, 530–538. 1

Barro, R. J. (1986): "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, 17, 3–20. 1

Barro, R. J. and D. B. Gordon (1983a): "A positive theory of monetary policy in a natural rate model," *Journal of political economy*, 91, 589–610. 2.3

——— (1983b): "Rules, discretion and reputation in a model of monetary policy," *Journal of monetary economics*, 12, 101–121. 1

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024): "Perceptions about monetary policy," *The Quarterly Journal of Economics*, 139, 2227–2278. 1

Beauregard, R., J. H. E. Christensen, E. Fischer, and S. Zhu (2024): "Inflation Expectations and Risk Premia in Emerging Bond Markets: Evidence from Mexico," *Journal of International Economics*, 151, 103961. D.1

Beechey, M. J. and J. H. Wright (2009): "The High-frequency Impact of News on Long-term Yields and Forward Rates: Is it Real?" *Journal of Monetary Economics*, 56, 535–544. 1

Bems, R., F. Caselli, F. Grigoli, and B. Gruss (2021): "Expectations' anchoring and inflation persistence," *Journal of International Economics*, 132, 103516. 1

Bernanke, B. (2007): "Inflation expectations and inflation forecasting," Tech. rep., Board of Governors of the Federal Reserve System (US). 1

Bocola, L., G. Chaumont, A. Dovis, and R. Kirpalani (2025): "Accounting for credibility: Monetary-fiscal interactions and the credibility of the central bank mandate," Tech. rep., Unpublished manuscript. 6

Bocola, L., A. Dovis, K. Jørgensen, and R. Kirpalani (2024): "Bond market views of the fed," Tech. rep., National Bureau of Economic Research. 1

Bonomo, M., C. Carvalho, S. Eusepi, M. Perrupato, D. Abib, J. Ayres, and S. Matos (2024): "Abrupt monetary policy change and unanchoring of inflation expectations," *Journal of Monetary Economics*, 145, 103576. 1, 4.2, 5.1, D.1

Caravello, T., P. Martinez-Bruera, and A. Carrasco (2025): "Hawkish Dove or Dovish Hawk? Optimal Monetary Policy with Reputational Concerns," *Optimal Monetary Policy with Reputational Concerns (April 30, 2025)*. 1, 4.2

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023): "Anchored inflation expectations," *American Economic Journal: Macroeconomics*, 15, 1–47. 1

Ceballos, L., J. H. E. Christensen, and D. Romero (2025): "A Post-Pandemic New Normal for Interest Rates in Emerging Bond Markets? Evidence from Chile," *Journal of International Money and Fincance*, 150, 103234. D.1

Christoffel, K. P. and M. Farkas (2025): "Managing the risks of inflation expectation de-anchoring," . 6

Clarida, R., J. Gali, and M. Gertler (1999): "The science of monetary policy: a new Keynesian perspective," *Journal of economic literature*, 37, 1661–1707. 1

Cochrane, J. H. and M. Piazzesi (2002): "The Fed and Interest Rates – A High-frequency Identification," *American Economic Review*, 92, 90–95. 1, 4.1

Coibion, O. and Y. Gorodnichenko (2025): "Inflation, Expectations and Monetary Policy: What Have We Learned and to What End?" Tech. rep., National Bureau of Economic Research. 1

Coibion, O., Y. Gorodnichenko, and S. Kumar (2018): "How do firms form their expectations? New survey evidence," *American Economic Review*, 108, 2671–2713. 1

Cook, T. and T. Hahn (1989): "The effect of changes in the federal funds rate target on market interest rates in the 1970s," *Journal of Monetary Economics*, 24, 331–351. 1, 4.1

de Aguilar, A. R. (2024): "Debt, Inflation, And Government Reputation," . 1

De Pooter, M., P. Robitaille, I. Walker, and M. Zdinak (2014): "Are Long-term Inflation Expectations Well Anchored in Brazil, Chile, and Mexico?" *International Journal of Central Banking*, 10, 337–400. 1

dos Santos Fernandes, G. and M. S. Portugal (2023): "Flattening of the Phillips curve in Brazil after inflation targeting: evidence from a multi-region model with time-varying parameters," Tech. rep., Unpublished manuscript. 5.1

Erceg, C. J. and A. T. Levin (2003): "Imperfect credibility and inflation persistence," *Journal of monetary economics*, 50, 915–944. 1

Gaglianone, W. P., R. Giacomini, J. V. Issler, and V. Skreta (2022): "Incentive-driven inattention," *Journal of Econometrics*, 231, 188–212. D.1

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press. 1

Gati, L. (2023): "Monetary policy and anchored expectations—An endogenous gain learning model," *Journal of Monetary Economics*, 140, S37–S47. 1

Gürkaynak, R. S., A. Levin, and E. Swanson (2010): "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-term Bond Yields in the U.S., U.K., and Sweden," *Journal of European Economic Association*, 8, 1208–1242. 1

Gürkaynak, R. S., B. Sack, and E. Swanson (2005): "The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models," *American economic review*, 95, 425–436. 1

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The US Treasury yield curve: 1961 to the present," *Journal of monetary Economics*, 54, 2291–2304. D.1

Hanson, S. G. and J. C. Stein (2015): "Monetary Policy and Long-term Real Rates," *Journal of Financial Economics*, 115, 429–448. 1, 4.2

Jarociński, M. and P. Karadi (2020): "Deconstructing monetary policy surprises—the role of information shocks," *American Economic Journal: Macroeconomics*, 12, 1–43. 4.2

Kekre, R., M. Lenel, and F. Mainardi (2024): "Monetary policy, segmentation, and the term structure," Tech. rep., National Bureau of Economic Research. 1

King, R. G. and Y. K. Lu (2022): "Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation," Tech. rep., National Bureau of Economic Research. 1

Kostadinov, R. and F. Roldán (2024): "Reputation and the Credibility of Inflation Plans," *Available at SSRN 4567776*. 1

Kumar, S., H. Afrouzi, O. Coibion, and Y. Gorodnichenko (2015): "Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand," *Brookings papers on economic activity*, 2015, 151–225. 1

Kuttner, K. N. (2001): "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of monetary economics*, 47, 523–544. 1, 4.1

Kydland, F. E. and E. C. Prescott (1977): "Rules rather than discretion: The inconsistency of optimal plans," *Journal of political economy*, 85, 473–491. 1, 2.3

Lu, Y. K., R. G. King, and E. Pasten (2016): "Optimal reputation building in the New Keynesian model," *Journal of Monetary Economics*, 84, 233–249. 1

Morelli, J. M. and M. Moretti (2023): "Information frictions, reputation, and sovereign spreads," *Journal of Political Economy*, 131, 3066–3102. 5

Nakamura, E., V. Riblier, and J. Steinsson (2025): "Beyond the Taylor Rule," Tech. rep., National Bureau of Economic Research. 1, 5.3

Nakamura, E. and J. Steinsson (2018): "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 133, 1283–1330. 1, 4.2, 4.2

Robitaille, P., T. Zhang, and B. Weisberg (2024): "How Well-Anchored Are Long-term Inflation Expectations in Latin America?" *FEDS Notes, Washington: Board of Governors of the Federal Reserve System, December 20,*. 1

Rogoff, K. (1985): "The optimal degree of commitment to an intermediate monetary target," *The quarterly journal of economics*, 100, 1169–1189. 1

Witheridge, W. (2024): "Monetary Policy and Fiscal-led Inflation in Emerging Markets," *Working Paper*. 1

Woodford, M. (2003): *Interest and prices: Foundations of a theory of monetary policy*, Princeton University Press. 1

Appendix

A Equilibrium definition

In this section, we fully define the equilibrium for the private agents. We start by defining a Recursive Competitive Equilibrium (RCE) for the private agents given a set of interest rate rules followed by the central bank of each types and the evolution of reputation consistent with Bayes' rule. In particular, let $S = (P_-, s, \rho)$ and denote the policy rules as $\iota(S, \theta, \varepsilon)$. The evolution of reputation follows Bayes' rule as in the text: $\rho' = \rho'(S, i)$.

Given $\iota(S, \theta, \varepsilon)$ and $\rho'(S, i)$, a RCE is a set of decision rule and a value function for the stand-in household, c, l, b' and v; a decision rule and value function for the intermediate good firm, p', y, and J; price functions P, W, and Q; aggregate allocation rules C, L, and Y; and a transfer rule T such that i) c, l, b' and v solve the stand-in household's problem:

$$v(S, b, i) = \max_{c,l,b'} U(c, l) + \beta \mathbb{E}^{\rho'} v\left(S', b', i'\right) \tag{35}$$

subject to the budget constraint

$$P(S, i)\, c + \frac{b'}{1+i} = W(S, i)\, l + b + T(S, i),$$

the NPG, and the evolution of the aggregate states, $S' = (s', \rho')$ and $i' = \iota(s', \rho', \theta, \varepsilon')$ with $\rho' = \rho'(S, i)$; ii) p, y and J solve the intermediate-good firm's problem:

$$J(S, i, p_-) = \max_{p,y} \left(p - \frac{W(S, i)}{z}\right) y - \frac{\phi z}{2} \left(\frac{p}{p_-} - 1\right)^2 P(S, i) + \mathbb{E}^{\rho'} \left[Q\left(S, S', i'\right) J\left(S', i', p\right)\right] \tag{36}$$

subject to the demand function

$$y = \left(\frac{p}{P(S, i)}\right)^{\frac{\mu}{1-\mu}} Y(S, i)$$

the evolution of the aggregate states, $S' = (s', \rho')$ and $i' = \iota(s', \rho', \theta, \varepsilon')$ with $\rho' = \rho'(S, i)$; iii) representativeness conditions $c(S, i, 0) = C(S, i)$, $l(S, i, 0) = L(S, i)$, $b'(S, i, 0) = 0$ and $p(S, i, P_-) = P(S, i)$, $y(S, i, P_-) = Y(S, i)$; and iv) market clearing

$$\begin{aligned} zL(S, i) &= Y(S, i) \\ Y(S, i) &= C(S, i) + \frac{\phi z}{2} \left(\frac{P(S, i)}{P_-} - 1\right)^2 \end{aligned}$$

where the nominal stochastic discount factor is given by

$$Q\left(S,S',i'\right)=\beta\frac{U_C\left(S',i'\right)/P\left(S',i'\right)}{U_C\left(S,i\right)/P\left(S,i\right)}$$

and the nominal aggregate transfers from intermediate good firms to the households are given by

$$T\left(S,i\right)=\left(P\left(S,i\right)-\frac{W\left(S,i\right)}{z}\right)Y\left(S,i\right)-\frac{\phi z}{2}\left(\frac{P\left(S,i\right)}{P_{-}}-1\right)^{2}.$$

We can derive the inflation rate as

$$\pi\left(S,i\right)=\frac{P\left(S,i\right)}{P_{-}}-1.$$

Note that this formulation implies that the policy, allocation and price rules induce a continuation competitive equilibrium for every after any (S,i), also for the interest rates that do not arise along the equilibrium path.

A Markov perfect equilibrium is a central bank policy rule $\iota\left(S,\theta,\varepsilon\right)$ and its associated value function $V\left(S,\theta,\varepsilon\right)$, the evolution of reputation $\rho'\left(S,i\right)$, and the associated recursive competitive equilibrium such that that i) ι and V solve the central bank's problem

$$V\left(s,\rho,\theta,\varepsilon\right)=\max_{i}R\left(\pi\left(s,\rho,i\right),Y\left(s,\rho,i\right);\alpha\left(\theta,\varepsilon\right)\right)+\beta\mathbb{E}^{cb}V\left(s',\rho'\left(s,\theta,i\right),\theta',\varepsilon'\right)$$

ii) ρ' satisfies Bayes' rule whenever possible[23]

$$\rho'\left(S,i\right)=\frac{\rho\Pr\left(i|\theta_H,S\right)P_{HH}+\left(1-\rho\right)\Pr\left(i|\theta_D,S\right)P_{DH}}{\rho\Pr\left(i|\theta_H,S\right)+\left(1-\rho\right)\Pr\left(i|\theta_D,S\right)}. \tag{37}$$

and iii) given ι and p, the allocation and pricing rules are a RCE.

We can further simplify the definition of the equilibrium by noticing that allocation and price rules only have to satisfy (11), (12), and (13).

Let's first argue that (11)–(13) are necessary. Consider first the stand-in household's problem, (35). The first order necessary conditions after imposing representativeness are

$$\frac{1}{1+i}=\mathbb{E}^{\rho}\left[\frac{C\left(S',i'\right)^{-\sigma}/C\left(S,i\right)^{-\sigma}}{1+\pi\left(S',i'\right)}\right]$$

[23]Here the assumption is that range(ι) does not depend on θ. For $i\notin$ range (ι) then we can set $\rho'=0$.

and

$$\frac{W(S,i)}{P(S,i)} = \frac{\chi L(S,i)^{1/\nu}}{C(S,i)^{-\sigma}}.$$

Thus, the Euler equation is necessary. Consider next the intermediate good firm's problem, (36). The necessary first order condition is

$$\begin{aligned}
0 &= \left(\frac{p}{P}\right)^{\frac{\mu}{1-\mu}} Y + \frac{\mu}{1-\mu}\left(p - \frac{W}{z}\right)\left(\frac{p}{P}\right)^{\frac{\mu}{1-\mu}-1}\frac{1}{P}Y - \phi z\left(\frac{p}{p_-} - 1\right)\frac{P}{p_-} + \mathbb{E}^{\rho'}\left[Q'\frac{\partial J(S',i',p)}{\partial p}\right] \\
&= \left(\frac{p}{P}\right)^{\frac{\mu}{1-\mu}} Y + \frac{\mu}{1-\mu}\left(p - \frac{W}{z}\right)\left(\frac{p}{P}\right)^{\frac{\mu}{1-\mu}-1}\frac{1}{P}Y - \phi z\left(\frac{p}{p_-} - 1\right)\frac{P}{p_-} + \mathbb{E}^{\rho'}\left[Q'\phi z'\left(\frac{p'}{p} - 1\right)\frac{p'}{p^2}P'\right]
\end{aligned}$$

where the second line follows from the envelope condition. Imposing representativeness (symmetry in pricing decision), $p = P$ and $p' = P'$, we obtain

$$\begin{aligned}
0 &= Y + \frac{\mu}{1-\mu}\left(1 - \frac{W/P}{z}\right)Y - \phi z\pi(1+\pi) + \mathbb{E}^{\rho'}\left[Q'\phi z'\pi'(1+\pi')^2\right] \\
&= \frac{1}{1-\mu}\left(1 - \mu\frac{W/P}{z}\right)Y - \phi z\pi(1+\pi) + \mathbb{E}^{\rho'}\left[\beta\left(\frac{C'}{C}\right)^{-\sigma}\phi z'\pi'(1+\pi')\right]
\end{aligned}$$

where in the second line we used the definition of the nominal SDF. Dividing by ϕz and rearranging we obtain

$$\pi(1+\pi) = \frac{Y}{\phi z(1-\mu)}\left(1 - \mu\frac{W/P}{z}\right) + \mathbb{E}^{\rho'}\left[\beta\left(\frac{C'}{C}\right)^{-\sigma}\frac{z'}{z}\pi'(1+\pi')\right]$$

which is the Phillips curve in (12) once we substitute the expression for the real wage from the household's problem. Finally, the resource constraint (13) also has to hold because it is an equilibrium requirement.

To show that (11)–(13) are sufficient, we need to verify that the appropriate transversality conditions are satisfied. For the Euler equation, noting that $B = 0$, the transversality is satisfied. For the firm's problem, we need $\lim_{t\to\infty}\mathbb{E}_0^{\rho'}\left[\beta^t\left(\frac{C_t}{C_0}\right)^{-\sigma}\phi z_t\tilde{\pi}_t\right] = 0$ which is satisfied if $\tilde{\pi}_t$ and marginal utility of consumption are bounded.

Note that P_- does not matter and only inflation does so we can drop it from the state as we are doing in the text so $S = (s,\rho)$.

B Derivations for Section 3

B.1 Log-linearized economy

Assume that $y^*(s) = \left(\frac{z^{1+1/\nu}}{\bar{\mu}\chi}\right)^{1/(\sigma+1/\nu)}$.

Steady state We first show that there exists a steady state with $y_{ss} = c_{ss} = y^*$ and $\pi_{ss} = 0$ that is independent of the central bank's type. If $\alpha = (\theta, 0)$ is known and $\bar{y}^* = y^*(\bar{s})$ is the flexible price output in steady state, the primal version of the central bank's problem is

$$V = \max_{\pi,y,c} -\frac{1}{2}\left[\alpha\pi^2 + (1-\alpha)\left(\frac{y-\bar{y}^*}{\bar{y}^*}\right)^2\right] + \beta V \tag{38}$$

subject to

$$y = c + \frac{\phi}{2}\bar{z}\pi^2$$
$$\tilde{\pi} = \frac{y\left[\frac{\bar{\mu}\chi y^{1/\nu}c^{\sigma}}{\bar{z}^{1+1/\nu}} - 1\right]}{\phi\left[\mu - 1\right]} + \beta\frac{c^{\sigma}}{\bar{z}}\tilde{\Pi}_{ss}$$

with

$$\tilde{\Pi}_{ss} = \frac{\bar{z}\tilde{\pi}_{ss}}{c_{ss}}$$

We will proceed by guess and verify. If $y_{ss} = c_{ss} = \bar{y}^*$ and $\pi_{ss} = 0$ is a steady state, then $\tilde{\Pi}_{ss} = 0$. By inspecting problem (38) with $\tilde{\Pi}_{ss} = 0$, it is clear that the central bank – no matter α— can attain the maximal static value of $R = 0$ by setting $\pi = 0$ and $y = c = y^*$ because by construction $\frac{\bar{\mu}\chi y^{\sigma+1/\nu}}{\bar{z}^{1+1/\nu}} = 1$ and the constraints in (38) are satisfied. Thus, the guess is verified: $y_{ss} = c_{ss} = \bar{y}^*$ and $\pi_{ss} = 0$ is a steady state.

Log-linearization We next log-linearize the equilibrium conditions (4)-(6) around the zero-inflation steady state. We obtain a log-linearized Euler equation,

$$-i_t = -\sigma\left(E_t\hat{c}_{t+1} - \hat{c}_t\right) - E_t\pi_{t+1}$$

where $\hat{c}_t = \log c_t - \log c_{ss}$, and a linearized resource constraint[24]

$$\log y_t - \log y_{ss} = \hat{c}_t.$$

Since we defined $\hat{y}_t$ as deviation of output from its target, $\hat{y}_t \equiv \log y_t - \log y_t^*$, and

$$\log y_t^* = \log y_{ss} + \frac{1+1/\nu}{\sigma+1/\nu}\hat{z}_t, \tag{39}$$

we can write

$$\hat{c}_t = \hat{y}_t + \frac{1+1/\nu}{\sigma+1/\nu}\hat{z}_t$$

and the log-linearized Euler equation as

$$-i_t = -\sigma\left(E_t\hat{y}_{t+1} - \hat{y}_t\right) - E_t\pi_{t+1} - \sigma\frac{1+1/\nu}{\sigma+1/\nu}E_t\left(\hat{z}_{t+1} - \hat{z}_t\right)$$

Using $\log y - \log y_{ss} = \hat{c}_t$, the log-linearized Phillips curve is[25]

$$\begin{aligned}
\pi_t = & \left[\frac{\left[\bar{\mu}\frac{\chi y_{ss}^{\sigma+1/\nu}}{\bar{z}^{1+1/\nu}} - 1\right]}{\phi\left[\bar{\mu}-1\right]} + \frac{y_{ss}\left(\frac{1}{\nu}+\sigma\right)\bar{\mu}\frac{\chi y_{ss}^{\sigma+1/\nu-1}}{\bar{z}^{1+1/\nu}}}{\phi\left[\bar{\mu}-1\right]}\right]\left(y_t - y_{ss}\right) \\
& - \frac{y_{ss}\left(1+\frac{1}{\nu}\right)\frac{\bar{\mu}\chi y_{ss}^{\sigma+1/\nu}}{\bar{z}^{1+1/\nu}}}{\phi\left[\bar{\mu}-1\right]}\frac{\left(z_t-\bar{z}\right)}{\bar{z}} + \frac{y_{ss}\frac{\bar{\mu}\chi y_{ss}^{\sigma+1/\nu}}{\bar{z}^{1+1/\nu}}}{\phi\left[\mu-1\right]}\frac{\left(\mu_t-\mu\right)}{\bar{\mu}} + \beta E_t\pi_{t+1} \\
= & \frac{y_{ss}\left(\sigma+\frac{1}{\nu}\right)}{\phi\left[\bar{\mu}-1\right]}\frac{\left(y_t-y_{ss}\right)}{y_{ss}} - \frac{y_{ss}\left(1+\frac{1}{\nu}\right)}{\phi\left[\bar{\mu}-1\right]}\frac{\left(z_t-\bar{z}\right)}{\bar{z}} + \frac{y_{ss}}{\phi\left[\bar{\mu}-1\right]}\frac{\left(\mu_t-\mu\right)}{\bar{\mu}} + \beta E_t\pi_{t+1} \\
= & \frac{y_{ss}\left(\sigma+\frac{1}{\nu}\right)}{\phi\left[\bar{\mu}-1\right]}\left(\log y_t - \log y_{ss}\right) - \frac{y_{ss}\left(1+\frac{1}{\nu}\right)}{\phi\left[\bar{\mu}-1\right]}\hat{z}_t + \frac{y_{ss}}{\phi\left[\bar{\mu}-1\right]}\hat{\mu}_t + \beta E_t\pi_{t+1} \\
= & \frac{\left(\sigma+\frac{1}{\nu}\right)y_{ss}}{\phi\left[\mu-1\right]}\hat{y}_t + \frac{y_{ss}}{\phi\left[\mu-1\right]}\hat{\mu}_t + \beta E_t\pi_{t+1} \\
= & \kappa\hat{y}_t + \kappa_\mu\hat{\mu}_t + \beta E_t\pi_{t+1}
\end{aligned}$$

where the second equality follows from $\bar{\mu}\chi y_{ss}^{\sigma+1/\nu}/\bar{z}^{1+1/\nu} = 1$ by definition of y_{ss} , the third equality is algebra and it uses that $\hat{z}_t \approx \frac{(z_t-\bar{z})}{\bar{z}}$, $\hat{\mu}_t \approx \frac{(\mu_t-\bar{\mu})}{\bar{\mu}}$, and (39), and in the last line we

[24]Since the steady state has zero inflation, small deviations from the steady state do not generate first order output costs and $\hat{c} = \hat{y}$. This is the analog of having no first order costs from price dispersion in a Calvo model around a zero inflation steady state.

[25]Also in deriving this expression we are again using zero inflation steady state to have no effect from changes in discounting.

define

$$\kappa = \frac{\left(\frac{1}{\nu} + \sigma\right) y_{ss}}{\phi\left[\mu - 1\right]}, \quad \kappa_\mu = \frac{\kappa}{\left(\sigma + \frac{1}{\nu}\right)}.$$

We can also write the objective as

$$\begin{aligned} R\left(\pi, y, \theta, \varepsilon, s\right) &= -\frac{1}{2}\left[\alpha\left(\theta, \varepsilon\right)\pi^2 + \left(1 - \alpha\left(\theta, \varepsilon\right)\right)\left(\frac{y - y^*\left(s\right)}{y^*\left(s\right)}\right)^2\right] \\ &\approx -\frac{1}{2}\left[\alpha\left(\theta, \varepsilon\right)\pi^2 + \left(1 - \alpha\left(\theta, \varepsilon\right)\right)\hat{y}^2\right] \end{aligned}$$

B.2 Simplified model

Long-run: Markov equilibrium under full information Letting $\alpha = \alpha\left(\theta, 0\right)$, the Markov equilibrium under full information solves

$$V_2\left(\hat{s}, \theta\right) = \max_{\pi, \hat{y}} -\frac{1}{2}\left[\alpha\left(\theta, \varepsilon\right)\pi^2 + \left(1 - \alpha\left(\theta, \varepsilon\right)\right)\hat{y}^2\right] + \beta\mathbb{E}V_2\left(\hat{s}', \theta\right)$$

subject to the Phillips curve,

$$\pi = \kappa\hat{y} + \kappa_\mu\hat{\mu} + \beta\mathbb{E}\pi\left(\hat{s}', \theta\right)$$

given inflation expectations $E\pi\left(\hat{s}', \theta\right)$. In equilibrium, inflation expectations must be consistent with the central bank's policy.

Lemma 4. *The allocation and price rule in the Markov equilibrium for the log-linear economy under full information are*

$$\pi_2\left(\hat{s}, \theta\right) = \frac{\left(1 - \alpha\right)\kappa_\mu}{\left(1 - \alpha\right)\left(1 - \beta\rho\right) + \alpha\kappa^2}\hat{\mu} \tag{40}$$

$$\hat{y}_2\left(\hat{s}, \theta\right) = -\kappa^2\frac{\alpha\kappa_\mu}{\left(1 - \alpha\right)\left(1 - \beta\rho\right) + \alpha\kappa^2}\hat{\mu}. \tag{41}$$

To see this, note that the optimality condition is

$$\alpha\pi\kappa + \left(1 - \alpha\right)\hat{y} = 0 \to \hat{y} = -\kappa^2\frac{\alpha}{1 - \alpha}\pi \tag{42}$$

Thus, the allocation and price rule $\hat{y}\left(\hat{s}, \theta\right)$ and $\pi\left(\hat{s}, \theta\right)$ solve the optimality condition above

and the Phillips curve. Combining these two, we obtain

$$\pi = \frac{1-\alpha}{1-\alpha+\alpha\kappa^2}\left[\kappa_\mu\hat{\mu} + \beta\mathbb{E}\pi\left(\hat{s}', \theta\right)\right] = b\left[\kappa_\mu\hat{\mu} + \beta\mathbb{E}\pi\left(\hat{s}', \theta\right)\right]$$

with $b \equiv \frac{1-\alpha}{1-\alpha+\alpha\kappa^2}$. Solving it forward, we obtain

$$\begin{aligned}\pi &= b\kappa_\mu\hat{\mu}_0 + b^2\beta\left[\kappa_\mu\rho\hat{\mu}_0 + \beta\mathbb{E}_0\pi_1\right] = ... = \kappa_\mu b\sum_{t=0}^{\infty}\left(\beta\rho b\right)^t\hat{\mu}_0 = \frac{\kappa_\mu b}{1-\beta\rho b}\hat{\mu}_0 \\ &= \frac{\kappa_\mu\left(1-\alpha\right)}{\left(1-\alpha\right)\left(1-\beta\rho\right)+\alpha\kappa^2}\hat{\mu}_0\end{aligned}$$

which gives the pricing rule (40). Using the optimality condition (42) in (40) gives the allocation rule (41).

Short-run with reputational concerns As explained in the text, we focus on a *separating equilibrium* in which the central bank chooses a different policy for any $\alpha = \alpha\left(\theta, \varepsilon\right)$. In fact, inspecting problem (43), the optimal $\hat{y}_1$ depends on θ and ε solely through their effect on $\alpha = \alpha\left(\theta, \varepsilon\right)$. We can thus write the equilibrium outcome as $\hat{y}_1\left(\hat{s}, \rho, \alpha\right)$ that solves

$$\begin{aligned}\hat{y}_1\left(\hat{s}, \rho, \alpha\right) &= \arg\max_{\hat{y}} R\left(\kappa\hat{y} + \kappa_\mu\hat{\mu} + \beta\Pi_2\left(\hat{s}, \rho, \hat{y}\right), \hat{y}; \alpha\right) \\ &= \arg\max_{\hat{y}} -\frac{1}{2}\left[\alpha\left(\kappa\hat{y} + \kappa_\mu\hat{\mu} + \beta\Pi_2\left(\hat{s}, \rho, \hat{y}\right)\right)^2 + \left(1-\alpha\right)\hat{y}^2\right]\end{aligned} \tag{43}$$

taking the function $\Pi_2\left(\cdot\right)$ as given. Private sector's expectations of inflation in the next period *after* the central bank acts as

$$\Pi_2\left(\hat{s}, \rho, \hat{y}\right) = p\left(\hat{s}, \rho, \hat{y}\right)\mathbb{E}\left[\pi_2\left(\hat{s}', \theta_H\right)|\hat{s}\right] + \left[1 - p\left(\hat{s}, \rho, \hat{y}\right)\right]\mathbb{E}\left[\pi_2\left(\hat{s}, ' \theta_D\right)|\hat{s}\right] \tag{44}$$

where $p\left(\hat{s}, \rho, \hat{y}\right)$ satisfies (17) given the central bank's policy function $\hat{y}\left(\hat{s}, \rho, \theta, \varepsilon\right)$. Finding an equilibrium then boils down to find a policy $\hat{y}_1\left(\cdot\right)$ and expectations $\Pi_2\left(\cdot\right)$ that are consistent with (44) and (43).We

We are going to describe the procedure we are using to find an equilibrium where $\hat{y}_1\left(\hat{s}, \rho, \alpha\right)$ is strictly decreasing (increasing) in α, and $\Pi_2\left(\hat{s}, \rho, \hat{y}\right)$ is strictly increasing (decreasing) in α if $\hat{\mu} > 0$ ($\hat{\mu} < 0$).

Note that we can find the equilibrium $\left(\hat{y}_1, \Pi_2\right)$ independently for all $\left(\hat{s}, \rho\right)$. Thus, fix an arbitrary observable state $\left(\hat{s}, \rho\right)$ with $\hat{\mu} > 0$. Let $\pi_{2D} = \mathbb{E}\left[\pi_2\left(\theta_D, \hat{s}'\right)|\hat{s}\right]$ and $\pi_{2H} = \mathbb{E}\left[\pi_2\left(\theta_H, \hat{s}'\right)|\hat{s}\right]$ be the expected inflation rate next period if the private agents know they are facing the Dove or the Hawk for sure, respectively.

Let X be the space of functions $\Pi_2 : Y \to [\pi_{2H}, \pi_{2D}]$ that are increasing, continuous, differentiable, bounded over some bounded set Y.[26] Elements in this set represent the possible levels for equilibrium inflation expectations given a realized y.

Define the operator $T : X \to X$ as

$$T(\Pi_2)(y) = \pi_{2D} - P(a(\Pi_2)(y))\Delta\pi_2 \tag{45}$$

where $\Delta\pi_2 \equiv \pi_{2D} - \pi_{2H} > 0$ is the difference in second period inflation under a Dove and a Hawk, $a(\Pi_2)(y) = y^{-1}(\Pi_2)(\alpha)$ with

$$y(\Pi_2)(\alpha) = \arg\max_{y \in Y} R(\kappa y + \beta\Pi_2(y) + \mu, y; \alpha) \tag{46}$$

and, as in the text,

$$P(\alpha) \equiv \frac{\rho \Pr(\alpha|\theta_H)}{\rho \Pr(\alpha|\theta_H) + (1-\rho)\Pr(\alpha|\theta_D)}.$$

Note that without loss of generality we are imposing that $y \leq 0$ since we are considering $\mu > 0$.

Finding an equilibrium reduces to finding a fixed point of T, $\Pi_2 = T(\Pi_2)$, and the associated $y(\Pi_2)(\alpha)$ is the equilibrium outcome $y(\alpha)$.

We first show that T maps elements in X into X. Let's first show that $T(\Pi_2)(y)$ is increasing in y. Let $\Pi_2 \in X$ and consider the maximization problem in (46). We can write the objective function as

$$F(y;\alpha) = R(\kappa y + \beta\Pi_2(y) + \mu, y; \alpha).$$

Note that the objective function is submodular as

$$\frac{\partial F}{\partial y} = -\alpha\left[\kappa y + \beta\Pi_2(y) + \mu\right]\left[\kappa + \beta\Pi_2'(y)\right] - (1-\alpha)y$$
$$\frac{\partial F}{\partial y \partial \alpha} = y - \left[\kappa y + \beta\Pi_2(y) + \mu\right]\left[\kappa + \beta\Pi_2'(y)\right] < 0$$

where the last inequality follows from $y \leq 0$, $\Pi_2(y) > 0$, $\mu > 0$ and $\Pi_2' \geq 0$ since Π_2 is increasing. Thus, by Topkis theorem it follows that $y(\Pi_2)(\alpha)$ is strictly decreasing in α. This implies that it is invertible, and $a(\Pi_2)(y) \equiv y^{-1}(\Pi_2)(y)$ is strictly decreasing in y. Using this observation in (45), since $\Delta\pi_2 > 0$ and $P(\alpha)$ is increasing, we have that $T(\Pi_2)(y)$ is increasing in y. A symmetric argument applies for the case where $\hat{\mu} < 0$.

[26] We will later show that we can restrict to the set $Y = \left[\underline{\hat{y}}, 0\right]$ with $\underline{\hat{y}} = -\frac{\beta\pi_{2H}(\hat{s}) + \kappa_\mu \hat{\mu}}{\kappa}$.

We can establish other properties of the monotone equilibrium. In particular, if $\hat{\mu} > 0$, for any $\rho \in (0,1)$, the range of $\hat{y}_1(\hat{s}, \rho, \cdot)$ is $\left[\underline{\hat{y}}, 0\right]$ where $\underline{\hat{y}}$ is defined as

$$0 = \kappa \underline{\hat{y}} + \beta \pi_{2H}(\hat{s}) + \kappa_\mu \hat{\mu} \rightarrow \underline{\hat{y}} = -\frac{\beta \pi_{2H}(\hat{s}) + \kappa_\mu \hat{\mu}}{\kappa}$$

Thus, we have that if $\alpha = 0$,

$$\hat{y}_1(\hat{s}, \rho, 0) = 0, \quad p(\hat{s}, \rho, \hat{y} = 0) = 0, \quad \pi_1(\hat{s}, \rho, 0) = \frac{\beta \pi_{2H}(\hat{s}) + \kappa_\mu \hat{\mu}}{\kappa}$$

while if $\alpha = 1$,

$$\hat{y}_1(\hat{s}, \rho, 1) = \underline{\hat{y}}, \quad p\left(\hat{s}, \rho, \hat{y} = \underline{\hat{y}}\right) = 1, \quad \pi_1(\hat{s}, \rho, 1) = 0.$$

To see this, note that $\hat{y}_1(\hat{s}, \rho, \cdot)$ is a continuous and strictly increasing function over $[0,1]$. Thus, the image of $\hat{y}_1(\hat{s}, \rho, \cdot)$ is a compact set $[y_{min}, y_{max}]$. Since we assumed that ε is normally distributed, we have that

$$p(\hat{s}, \rho, y_{min}) = \frac{\rho \Pr(\alpha = 1|\theta_H)}{\rho \Pr(\alpha = 1|\theta_H) + (1-\rho)\rho \Pr(\alpha = 1|\theta_L)} = 1$$
$$p(\hat{s}, \rho, y_{max}) = \frac{\rho \Pr(\alpha = 0|\theta_H)}{\rho \Pr(\alpha = 0|\theta_H) + (1-\rho)\rho \Pr(\alpha = 0|\theta_L)} = 0$$

Next, we show that $y_{max} = 0$. If $\alpha = 0$, the first-order condition for problem (43) is simply $0 = -\hat{y}$. Thus, $y_{max} = 0$. Since we know that $p(\hat{s}, \rho, \hat{y} = 0) = 0$, substituting into the Phillips curve we have that $\pi_1(\hat{s}, \rho, 0) = \frac{\beta \pi_{2H}(\hat{s}) + \kappa_\mu \hat{\mu}}{\kappa}$.

Finally, we show that $y_{min} = \underline{y}$. If $\alpha = 1$, the first-order condition for problem (43) is

$$0 = -\left(\kappa + \beta \frac{\partial \Pi_2(\hat{s}, \rho, y)}{\partial y}\right) \pi$$

Thus, since $\kappa > 0$ and $\partial \Pi_2(\hat{s}, \rho, y)/\partial y \geq 0$, it must be that $\pi = 0$. Using this in the Phillips curve together with the observation that $p(\hat{s}, \rho, \hat{y} = y_{max}) = 1$ gives that $y_{max} = \underline{y}$ for all $\rho \in (0,1)$.

B.3 Proof of Proposition 1

Solving the optimality condition (26) using our functional form we have

$$\hat{y}_1(\hat{s},\rho,\alpha) = -\frac{\alpha\left[\kappa+\beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\right]}{(1-\alpha)+\alpha\kappa\left[\kappa+\beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\right]}\left[\kappa\hat{\mu}+\beta\Pi_2(\hat{s},\rho,\hat{y})\right]$$

If the central bank is a Hawk and there is full information then the Markov equilibrium outcome is given by

$$\hat{y}_1^{full}(\hat{s},\alpha) = -\frac{\alpha\kappa}{(1-\alpha)+\alpha\kappa^2}\left[\kappa\hat{\mu}+\beta\pi_2(\hat{s},\theta_H)\right]$$

Note that for any $\rho\in(0,1)$ and $\hat{\mu}>0$, we have that $\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}>0$ as argued above in the text. Therefore,

$$\frac{\alpha\left[\kappa+\beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\right]}{(1-\alpha)+\alpha\kappa\left[\kappa+\beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}\right]} = \frac{\alpha\kappa}{(1-\alpha)/\left[1+\beta\frac{\partial\Pi_2(\hat{s},\rho,\hat{y})}{\partial\hat{y}}/\kappa\right]+\alpha\kappa^2} > \frac{\alpha\kappa}{(1-\alpha)+\alpha\kappa^2} > 0$$

If $\hat{\mu}>0$ we also know that

$$\Pi_2(\hat{s},\rho,\hat{y}) = p(\hat{s},\rho,\hat{y})\,\pi_2(\hat{s},\theta_H)+(1-p(\hat{s},\rho,\hat{y}))\,\pi_2(\hat{s},\theta_D) > \pi_2(\hat{s},\theta_H) > 0$$

Using these two observation in the expressions for $\hat{y}_1(\hat{s},\rho,\alpha)$ and $\hat{y}_1^{full}(\hat{s},\alpha)$ we have

$$\hat{y}_1(\hat{s},\rho,\alpha) < \hat{y}_1^{full}(s,\alpha)$$

as wanted. A specular argument holds for $\hat{\mu}<0$. *Q.E.D.*

B.4 Proof of Proposition 2

Fixing the state $(\hat{s},\rho,\alpha)$, the myopic allocation solves

$$\begin{aligned}\alpha\pi^{myopic}\kappa+(1-\alpha)\,\hat{y}^{myopic} &= 0\\ \pi^{myopic} &= \kappa\hat{y}^{myopic}+\kappa_\mu\hat{\mu}+\beta\Pi_2\left(\hat{y}^{myopic}\right)\end{aligned}$$

while the optimal policy solves

$$\alpha\pi\left[\kappa+\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}\right]+\left(1-\alpha\right)\hat{y}=0$$

$$\pi=\kappa\hat{y}+\kappa_\mu\hat{\mu}+\beta\Pi_2\left(\hat{y}\right)$$

Thus,

$$\alpha\kappa\left(\pi-\pi^{myopic}\right)+\left(1-\alpha\right)\left(\hat{y}-\hat{y}^{myopic}\right)=-\pi\alpha\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}$$

or

$$\alpha\kappa\left(\left[\kappa\hat{y}+\beta\Pi_2\left(\hat{y}\right)\right]-\left[\kappa\hat{y}^{myopic}+\beta\Pi_2\left(\hat{y}^{myopic}\right)\right]\right)+\left(1-\alpha\right)\left(\hat{y}-\hat{y}^{myopic}\right)=-\pi\alpha\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}$$

Using a first-order Taylor approx for $\Pi_2\left(\hat{y}^{myopic}\right)\approx\Pi_2\left(\hat{y}\right)-\frac{\partial\Pi_2(\hat{y})}{\partial\hat{y}}\left(\hat{y}-\hat{y}^{myopic}\right)$, we can write

$$\begin{aligned}-\pi\alpha\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}&\approx\alpha\kappa\left(\kappa\left(\hat{y}-\hat{y}^{myopic}\right)+\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}\left(\hat{y}-\hat{y}^{myopic}\right)\right)+\left(1-\alpha\right)\left(\hat{y}-\hat{y}^{myopic}\right)\\&=\left[\alpha\kappa^2+\alpha\kappa\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}+\left(1-\alpha\right)\right]\left(\hat{y}-\hat{y}^{myopic}\right)\end{aligned}$$

Rearranging,

$$\hat{y}^{myopic}-\hat{y}\approx\frac{\pi\alpha\beta\frac{\partial\Pi_2(\hat{y})}{\partial\hat{y}}}{\left[\alpha\kappa^2+\alpha\kappa\beta\frac{\partial\Pi_2(\hat{y})}{\partial\hat{y}}+\left(1-\alpha\right)\right]}\geq0$$

and similarly,

$$\begin{aligned}\pi-\pi^{myopic}&\approx\left[\kappa+\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}\right]\left(\hat{y}-\hat{y}^{myopic}\right)\\&=\left[\kappa+\beta\frac{\partial\Pi_2\left(\hat{y}\right)}{\partial\hat{y}}\right]\frac{\pi\alpha\beta\frac{\partial\Pi_2(\hat{y})}{\partial\hat{y}}}{\left[\alpha\kappa^2+\alpha\kappa\beta\frac{\partial\Pi_2(\hat{y})}{\partial\hat{y}}+\left(1-\alpha\right)\right]}\end{aligned}$$

Q.E.D.

B.5 Proof of Proposition 3

Consider first $\Delta\mathbb{E}\pi_2\left(\hat{s},\rho,\hat{y}\right)$. Taking a first-order approximation at $\bar{y}$, we obtain

$$\Delta\mathbb{E}\pi_2\left(\hat{s},\rho,\hat{y}\right)=\Pi_2\left(\hat{s},\rho,\hat{y}\right)-\mathbb{E}\left(\pi_2|\hat{s},\rho\right)\approx\Pi_2'\left(\hat{s},\rho,\bar{y}\right)\left(\hat{y}-\bar{y}\right).$$

From period 2 onward, (40) and (41) imply that

$$\hat{y}_2(\hat{s},\theta) = -\frac{\alpha\kappa^2}{(1-\alpha)}\hat{\pi}_2(\hat{s},\theta)$$

hence

$$\mathbb{E}\left(\hat{y}_2|\hat{s},\rho'\right) = -\frac{\alpha\kappa^2}{(1-\alpha)}\Pi_2(\hat{s},\rho,\hat{y})$$

so the Euler equation in period 1 can be written as

$$\begin{aligned}\iota(\hat{s},\rho,\hat{y}) &= -\sigma\left[\hat{y} + \frac{\alpha\kappa^2}{(1-\alpha)}\Pi_2(\hat{s},\rho,\hat{y})\right] + \Pi_2(\hat{s},\rho,\hat{y}) + \gamma \\ &= -\sigma\hat{y} + \left(1 - \sigma\frac{\alpha\kappa^2}{(1-\alpha)}\right)\Pi_2(\hat{s},\rho,\hat{y}) + \gamma,\end{aligned}$$

where $\gamma = -\frac{1+1/\nu}{\sigma+1/\nu}(1-\rho_z)\hat{z}$. Then we have that

$$\Delta\mathbb{E}i(\hat{s},\rho,\hat{y}) \approx -\left[\sigma - \left(1 - \sigma\frac{\alpha\kappa^2}{(1-\alpha)}\right)\Pi_2'(\hat{s},\rho,\bar{y})\right](\hat{y}-\bar{y}).$$

Thus, combining the last expression with the first two, we obtain

$$\frac{\Delta\mathbb{E}\pi_2(\hat{s},\rho,\hat{y})}{\Delta\mathbb{E}i(\hat{s},\rho,\hat{y})} \approx -\frac{\Pi_2'(\hat{s},\rho,\bar{y})}{\sigma - \left(1 - \sigma\frac{\alpha\kappa^2}{(1-\alpha)}\right)\Pi_2'(\hat{s},\rho,\bar{y})}$$

Q.E.D.

C Numerical solution

In this section we presents the algorithm for the numerical solution of the non-linear model in Section 4. We first show that, by assuming that the IES is 1, consumption and output are linear in productivity and inflation is independent from productivity and so we can solve for a normalized problem where we drop z as a state variable. We then describe the algorithm we used to solve the model.

Under our parameterization, the primal problem in (15) reduces to

$$V(s,\rho,\theta,\varepsilon) = \max_{\pi,y,c} -\frac{1}{2}\left[\alpha(\theta,\varepsilon)\pi^2 + (1-\alpha(\theta,\varepsilon))\left(\frac{y-zy^*}{zy^*}\right)^2\right] + \beta\mathbb{E}^{cb}V\left(s',\rho',\theta',\varepsilon'\right) \quad (47)$$

subject to

$$
\begin{aligned}
y &= c + \frac{\phi}{2} z \pi^2 \\
\tilde{\pi} &= \frac{(y/z)\left[\mu \chi (y/z)^{1/\nu}(c/z) - 1\right]}{\phi[\mu - 1]} + \beta \frac{c}{z} \tilde{\Pi}(s, \rho, y) \\
\rho' &= p(s, \rho, y)
\end{aligned}
$$

where $y^*(s) = z y^* = z\delta\chi^{-1/(1+1/\nu)}$. The inflation expectations are given by

$$
\tilde{\Pi}(s, \rho, y) = \iint \left[\rho' \frac{\tilde{\pi}(s', \rho', \theta_H, \varepsilon')}{C(s', \rho', \theta_H, \varepsilon')/z'} + (1 - \rho') \frac{\tilde{\pi}(s', \rho', \theta_D, \varepsilon')}{C(s', \rho', \theta_D, \varepsilon')/z'} \right] d\Gamma(s' \mid s) \Phi(\varepsilon') d\varepsilon' \tag{48}
$$

and

$$
p(s, \rho, y) = \frac{\rho \Phi(\varepsilon : y = Y(s, \rho, \theta_H, \varepsilon)) P_{HH} + (1 - \rho) \Phi(\varepsilon : y = Y(s, \rho, \theta_D, \varepsilon)) P_{DH}}{\rho \Phi(\varepsilon : y = Y(s, \rho, \theta_H, \varepsilon)) + (1 - \rho) \Phi(\varepsilon : y = Y(s, \rho, \theta_D, \varepsilon))}
$$

Note that, in the constraint set and the objective function of (47), output and consumption only enters relative to the productivity level z; the same is true in (48). Thus, we can conjecture that in equilibrium output and consumption are linear in z and inflation is independent from z,

$$
C((\mu, z), \rho, \theta, \varepsilon) = z C((\mu, 1), \rho, \theta, \varepsilon) = z \bar{C}(\mu, \rho, \theta, \varepsilon), \tag{49}
$$

$$
Y((\mu, z), \rho, \theta, \varepsilon) = z Y((\mu, 1), \rho, \theta, \varepsilon) = z \bar{Y}(\mu, \rho, \theta, \varepsilon), \tag{50}
$$

$$
\pi((\mu, z), \rho, \theta, \varepsilon) = \pi((\mu, 1), \rho, \theta, \varepsilon) = \bar{\pi}(\mu, \rho, \theta, \varepsilon), \tag{51}
$$

the value function is independent from z, $V((\mu, z), \rho, \theta, \varepsilon) = V((\mu, 1), \rho, \theta, \varepsilon) = \bar{V}(\mu, \rho, \theta, \varepsilon)$, and

$$
\begin{aligned}
\tilde{\Pi}((\mu, z), \rho, y) &= \bar{\Pi}(\mu, \rho, y/z) \qquad (52) \\
&\equiv \iint \left[\rho' \frac{\tilde{\pi}(\mu', \rho', \theta_H, \varepsilon')}{\bar{C}(\mu', \rho', \theta_H, \varepsilon')} + (1 - \rho') \frac{\tilde{\pi}(s', \rho', \theta_D, \varepsilon')}{\bar{C}(\mu', \rho', \theta_D, \varepsilon')} \right] d\Gamma(\mu' \mid \mu) \Phi(\varepsilon') d\varepsilon'
\end{aligned}
$$

where $\rho' = \bar{p}(\mu, \rho, \cdot)$ given by

$$
\bar{p}(\mu, \rho, y) = \frac{\rho \Phi(\varepsilon : y = \bar{Y}(\mu, \rho, \theta_H, \varepsilon)) P_{HH} + (1 - \rho) \Phi(\varepsilon : y = \bar{Y}(\mu, \rho, \theta_D, \varepsilon)) P_{DH}}{\rho \Phi(\varepsilon : y = \bar{Y}(\mu, \rho, \theta_H, \varepsilon)) + (1 - \rho) \Phi(\varepsilon : y = \bar{Y}(\mu, \rho, \theta_D, \varepsilon))} \tag{53}
$$

Under this conjecture, we can rewrite the primal problem in terms of the normalized vari-

ables and dropping z as a state variable as

$$\bar{V}(\mu,\rho,\theta,\varepsilon) = \max_{\pi,y,c} -\frac{1}{2}\left[\alpha(\theta,\varepsilon)\pi^2 + (1-\alpha(\theta,\varepsilon))\left(\frac{y-y^*}{y^*}\right)^2\right] + \beta\mathbb{E}^{cb}\bar{V}(\mu',\rho',\theta',\varepsilon') \quad (54)$$

subject to

$$y = c + \frac{\phi}{2}\pi^2 \quad (55)$$

$$\tilde{\pi} = \frac{y\left[\mu\chi y^{1/\nu}c - 1\right]}{\phi[\mu-1]} + \beta c\bar{\Pi}(\mu,\rho,y) \quad (56)$$

$$\rho' = \bar{p}(\mu,\rho,y) \quad (57)$$

Finding an equilibrium then boils down to finding a set of allocation rules $\bar{Y}(\mu,\rho,\theta,\varepsilon)$, $\bar{C}(\mu,\rho,\theta,\varepsilon)$, $\bar{\pi}(\mu,\rho,\theta,\varepsilon)$, a value function $\bar{V}(\mu,\rho,\theta,\varepsilon)$, and an updating rule $\bar{p}(\mu,\rho,y)$ that solve (15) and satisfy (52)-(17). We can then find the policy rules for our original economy using (49)–(51).

The numerical solution of the model consists in approximating the functions $f = \{\bar{Y},\bar{C},\bar{\pi},\bar{V}\}$ via polynomials. Let $\mathbf{S} = (\hat{\mu},\rho,a)$ where $\hat{\mu}$ is the deviation from average markup shock, ρ is the level of reputation, a is the weight on inflation in the central bank's objective function. $\mathbf{T}(\mathbf{S})$ is a vector of Chebyshev polynomials evaluated at $\mathbf{S}$, and $\gamma^f(\theta)$ is the vector of Chebyshev coefficients that depends on the central banker's type θ. Then, we will approximate f as

$$f(\theta)(\hat{\mu},\rho,a) = \left(\gamma^f(\theta)\right)^\top \mathbf{T}(\mathbf{S}).$$

Before discussing the details of the algorithm, we present some specifics regarding the construction of the grid and of the polynomials. We construct a tensor grid of collocation points, denoted by $\mathcal{T}$, on the space of inflation weights α, supply shocks $\hat{\mu}$, and reputation ρ.[27] The bounds for this space are set to $[0.01, 0.99]$ for α, $\left[\frac{-2.5\sigma_\mu}{\sqrt{1-\rho_\mu^2}}, \frac{2.5\sigma_\mu}{\sqrt{1-\rho_\mu^2}}\right]$ for the supply shock, and $[0,1]$ for the reputation. We consider Chebyshev polynomials up to the fifth order on each dimensions.

The expectations with respect to monetary and markup shocks are computed using Gauss-Hermite quadrature with N_{GH} points, where we set $N_{GH} = 5$, for each shock separately, since the shocks are independent. For the purpose of evaluating expectations, we also construct a tensor grid $\mathcal{T}_{\mu,\rho}$ that is on the space of markup shock and reputation only. For the optimization step, we define the log-output grid y_{grid} as 5001 equally spaced points

[27] The policy function for output is very non-linear in α for α close to 1. In our experience, approximating the policy and value functions over $h = \log(1-\alpha)$ rather than α helps with the convergence of the algorithm.

on the interval $[-0.05, 0.05]$, and the weight grid a_{grid} as 5001 equally spaced points on $[0.01, 0.99]$.

Algorithm

We solve for the decision problem of the central bank via a backward induction algorithm. Taking as given the value function and the policies for iteration, $o-1$, the *inner loop* of the algorithm solves the decision problem of the central bank and obtains the value and policy functions for iteration o. The *outer loop* of the algorithm keeps iterating backward until the value function and policy functions between the last two iterations converge.

Step 1: Initialize Set up the grid of collocation points and provide the initial guess for the Chebyshev coefficients of policies and value function for the Hawk and the Dove: $\left\{\gamma^f_{o-1}(\theta)\right\}_{\theta\in\{\theta_H,\theta_D\}} = \left\{\gamma^Y_{o-1}(\theta), \gamma^\pi_{o-1}(\theta), \gamma^C_{o-1}(\theta), \gamma^V_{o-1}(\theta)\right\}_{\theta\in\{\theta_H,\theta_D\}}$ for $o=0$.

Step 2: Outer loop At iteration $o \geq 1$, taking as inputs the Chebyshev coefficients for the policy and value functions, $\left\{\gamma^f_{o-1}(\theta)\right\}_{\theta\in\{\theta_H,\theta_D\}}$, we compute the updated coefficient $\left\{\gamma^f_{o}(\theta)\right\}_{\theta\in\{\theta_H,\theta_D\}}$. To do so, we:

Construct the CB's expectations for every $(\hat{\mu}, \rho')$ Let $\mathcal{T}_{\mu,\rho}$ denote the tensor grid of collocation points for $(\hat{\mu}, \rho)$. Using $\left\{\gamma^f_{o-1}(\theta)\right\}$, we calculate expectations for the continuation value V and $\tilde{\pi}C^{-\sigma}$ conditional on current markup shock $\hat{\mu}$ and future reputation ρ' for every $(\hat{\mu}, \rho') \in \mathcal{T}_{\mu,\rho}$.[28] These expectations are represented by a collection of Chebyshev coefficients $\left\{\gamma^{EV}_o(\theta), \gamma^{\mathbb{E}\tilde{\pi}C^{-\sigma}}_o(\theta)\right\}$.

Inner loop The inner loop solves for the updated policy functions and value function coefficients,

$$\left\{\gamma^Y_o(\theta), \gamma^\pi_o(\theta), \gamma^C_o(\theta), \gamma^V_o(\theta)\right\}_{\theta\in\{\theta_H,\theta_D\}}.$$

[28] Let $\hat{\mu}, \rho' \in \mathcal{T}_{\mu,\rho}$. Given ρ' and $\hat{\mu}$, get a collection of values

$$\left\{\rho_\mu \hat{\mu} + \sqrt{2}\sigma_\mu \epsilon_j, \rho', \alpha\left(\theta, \sqrt{2}\sigma_\varepsilon \epsilon_\ell\right)\right\}_{(i,j)\in\{1,\ldots,N_{GH}\}^2}$$

where $\epsilon_\ell, \epsilon_j$ are the nodes of the Gauss-Hermite quadrature. Using the approximated functions from the outer loop $\hat{V}_{o-1}, \hat{\pi}_{o-1}, \hat{C}_{o-1}$, obtain the matrix of values of $V, \tilde{\pi}C^{-\sigma}$ at these points. Then take a weighted sum with the weights from the Gauss-Hermite quadrature. This gives us a conditional expectation of $V, \tilde{\pi}C^{-\sigma}$ given current markup shock $\hat{\mu}$ and future reputation ρ'. Compute this for every point on the tensor grid of collocation points $\mathcal{T}_{\mu,\rho}$.

Let $j = 0, 1, 2, ...$ index the iteration for the inner loop.

1. **Initialization (only for $j = 0$)**: For $j = 0$, set $\gamma_j^Y(\theta) = \gamma_{o-1}^Y(\theta)$ for both $\theta \in \{\theta_H, \theta_D\}$. Set the adjustment weight w_0 to 0.95.

2. **Invert the $j-1$ output policy function to recover the monetary shocks:** For any $j \geq 1$, for each central bank's type, for a given point $(\hat{\mu}, \rho)$ on the tensor grid of the collocation points $\mathcal{T}_{\mu,\rho}$, and for a given y on the log-output y_{grid}, we want to obtain the monetary shock $\varepsilon(\theta)(\hat{\mu}, \rho, y)$ such that

$$y = \left(\gamma_{j-1}^Y(\theta)\right)^\top \mathbf{T}\left(\hat{\mu}, \rho, \alpha\left(\theta, \varepsilon(\theta)(\hat{\mu}, \rho, y)\right)\right).$$

   To do so, we first find a weight $\tilde{\mathcal{A}}(\theta)(\hat{\mu}, \rho, y)$ on a weight grid a_{grid} that minimizes the squared distance between approximated log-output policy and the given value y, i.e.

$$\tilde{\mathcal{A}}(\theta)(\hat{\mu}, \rho, y) = \underset{a \in a_{\text{grid}}}{\operatorname{argmin}} \left(y - \left(\gamma_{j-1}^Y(\theta)\right)^\top \mathbf{T}(\hat{\mu}, \rho, a)\right)^2,$$

   We set the weight to 0.01 if the minimized squared distance exceeds the 10^{-5} threshold, i.e.

$$\mathcal{A}(\theta)(\hat{\mu}, \rho, y) = \begin{cases} 0.01 & \text{if } \min\limits_{a \in a_{\text{grid}}} \left(y - \left(\gamma_{j-1}^Y(\theta)\right)^\top \mathbf{T}(\hat{\mu}, \rho, a)\right)^2 > 10^{-5} \\ \tilde{\mathcal{A}}(\theta)(\hat{\mu}, \rho, y) & \text{otherwise} \end{cases}.$$

   Then we invert the weight to obtain the monetary shock,

$$\varepsilon_j(\theta)(\hat{\mu}, \rho, y) = \alpha^{-1}\left(\theta, \mathcal{A}(\theta)(\hat{\mu}, \rho, y)\right).$$

3. **Get the updated beliefs:** Use the monetary shocks values $\left\{\varepsilon_j(\theta)(\hat{\mu}, \rho, y)\right\}_{(\hat{\mu},\rho,y) \in \mathcal{T}_{\mu,\rho} \times y_{\text{grid}}}$ for $\theta \in \{\theta_H, \theta_D\}$ and the belief update equation (53) to evaluate the belief update function $\bar{p}_j(\hat{\mu}, \rho, y)$ for every $(\hat{\mu}, \rho, y)$ on $\mathcal{T}_{\mu,\rho} \times y_{\text{grid}}$.

4. **Get expectations as a function of y:** Compute the market expectations for inflation $\tilde{\Pi}_j(\hat{\mu}, \rho, y)$ for every $(\hat{\mu}, \rho, y)$ on $\mathcal{T}_{\mu,\rho} \times y_{\text{grid}}$ using the Chebyshev coefficients $\left\{\gamma_o^{\mathbb{E}\tilde{\pi}C^{-\sigma}}(\theta)\right\}$, the belief update function values $\left\{\bar{p}_j(\hat{\mu}, \rho, y)\right\}_{(\hat{\mu},\rho,y) \in \mathcal{T}_{\mu,\rho} \times y_{\text{grid}}}$ and the

equation (52) as[29]

$$\tilde{\Pi}_j(\hat{\mu},\rho,y) = \bar{p}_j(\hat{\mu},\rho,y)\left(\gamma_o^{\mathbb{E}\tilde{\pi}C^{-\sigma}}(\theta_H)\right)^{\top}\mathbf{T}\left(\hat{\mu},\bar{p}_j(\hat{\mu},\rho,y)\right) + (1-\bar{p}j(\hat{\mu},\rho,y))\left(\gamma_o^{\mathbb{E}\tilde{\pi}C^{-\sigma}}(\theta_D)\right)^{\top}\mathbf{T}\left(\hat{\mu},\bar{p}_j(\hat{\mu},\rho,y)\right).$$

5. **Search over y_{grid} to maximize the central bank's objective and compute the policies and value function:** For each central bank's type, we search over the y_{grid} to find the optimal output policy. We first use the normalized resource constraint (55), Phillips curve (56), and the market's inflation expectation $\tilde{\Pi}_j(\hat{\mu},\rho,y)$ to recover inflation for every $(\hat{\mu},\rho,y)$ on $\mathcal{T}_{\mu,\rho}\times y_{\text{grid}}$. Since in our calibration $\sigma=\nu=1$, this amount to solving a quadratic equation:

$$\pi_j(\theta)(\hat{\mu},\rho,y) = \left(\frac{-1+\sqrt{1-4a_\theta(\hat{\mu},\rho,y)\,b_\theta(\hat{\mu},\rho,y)}}{2a_\theta(\hat{\mu},\rho,y)}\right)$$

where

$$a_\theta(\hat{\mu},\rho,y) = 1+\frac{\phi}{2}\left[\frac{\hat{\mu}+\bar{\mu}}{\phi(\hat{\mu}+\bar{\mu}-1)}\chi\exp(y)^{1+\frac{1}{\nu}}+\tilde{\Pi}_j(\hat{\mu},\rho,y)\right],$$
$$b_\theta(\hat{\mu},\rho,y) = -\exp(y)\left(\frac{\hat{\mu}+\bar{\mu}}{\phi(\hat{\mu}+\bar{\mu}-1)}\chi\exp(y)^{1+\frac{1}{\nu}}-1\right)-\beta\exp(y)\,\tilde{\Pi}_j(\hat{\mu},\rho,y).$$

Compute and store $\pi_j(\theta)(\hat{\mu},\rho,y)$ for every $(\hat{\mu},\rho,y)$ on $\mathcal{T}_{\mu,\rho}\times y_{\text{grid}}$. Then, using these values, the approximated value function expectation,[30] compute the objective function for every $(\alpha,\hat{\mu},\rho)$ on $\mathcal{T}\times y_{\text{grid}}$ using the right side of (54). Search over y_{grid} to maximize the objective function of the central bank for every point on $\mathcal{T}$ and store the optimal values of y as $\{Y_{\text{new}}(\theta)(\mathbf{S})\}_{\mathbf{S}\in\mathcal{T}}$.

Then, for every collocation point, use the index for the optimal point on the y_{grid} and select the values for the optimal inflation policy and the value function. Finally, obtain the consumption policy values on collocation points using the goods market clearing condition (55).

6. **Update the output policy:** For each $\theta\in\{\theta_H,\theta_D\}$, and for every $\mathbf{S}$ on the tensor grid of collocation points $\mathcal{T}$, compute

$$Y_j(\theta)(\mathbf{S}) = \log\left((1-w_{j-1})\exp(Y_{\text{new}}(\theta)(\mathbf{S}))+w_{j-1}\exp\left(\left(\gamma_{j-1}^{Y}(\theta)\right)^{\top}\mathbf{T}(\mathbf{S})\right)\right)$$

[29] We are implicitly dropping the cross products with the weight dimension when we write $\mathbf{T}(\hat{\mu},p(\hat{\mu},\rho,y))$.
[30] Which corresponds to $\left(\gamma_o^{\mathbb{E}V}(\theta)\right)^{\top}\mathbf{T}(\hat{\mu},p(\hat{\mu},\rho,y))$.

where w_{j-1} is an adjustment weight.

7. **Check for convergence:** Compute the norm of the residuals for the iteration j as

$$\text{dist}_j = \max_{\mathbf{S}\in\mathcal{T},\theta\in\{\theta_H,\theta_D\}} \left| Y_{\text{new}}(\theta)(\mathbf{S}) - \left(\gamma_{j-1}^{Y}(\theta)\right)^{\top} \mathbf{T}(\mathbf{S}) \right|^2.$$

Check if dist_j is below the tolerance level, set to 10^{-5}. If not, update the weight w_j as

$$w_j = \begin{cases} \min\left\{1.01 w_{j-1}, 0.995\right\} & \text{if } \text{dist}_j > \text{dist}_{j-1} \\ w_{j-1} & \text{otherwise} \end{cases}$$

and go back to step 2 of the inner loop with updated $\gamma_j^{Y}(\theta)$, starting the $(j+1)$-th iteration of the inner loop.

If it is below the tolerance, compute the Chebyshev coefficients for inflation, consumption and the value function using their values computed at the collocation points from step 5 and the output policy function using the values at the collocation points from step 6. Return the new Chebyshev coefficients $\gamma_o^{\pi}(\theta)$, $\gamma_o^{C}(\theta)$, $\gamma_o^{V}(\theta)$ and $\gamma_o^{Y}(\theta)$.

Step 3: Check for convergence Compute the maximum of the squares of the differences between old and new output policies,

$$\text{dist}_o^{out} = \max_{\mathbf{S}\in\mathcal{T},\theta\in\{\theta_H,\theta_D\}} \left| \left(\gamma_o^{Y}(\theta)\right)^{\top} \mathbf{T}(\mathbf{S}) - \left(\gamma_{o-1}^{Y}(\theta)\right)^{\top} \mathbf{T}(\mathbf{S}) \right|^2.$$

If dist_o^{out} is below the tolerance level, set to 10^{-5}, stop. Otherwise, go back to the step 2 starting the $(o+1)$-th iteration of the outer loop.

D Empirical analysis

D.1 Data

Monetary policy surprises. We primarily rely on overnight index swap (OIS) rates to construct monetary policy shocks. In particular, we use OIS rates for Chile, the EA, the UK, and the US. In the cases of Brazil and Mexico, where OIS rates are not available, we follow the approach of Bolhuis, Das, and Yao (2024) and use alternative interest rate swaps with the shortest available tenor. For Brazil, this is the BRZ PRE-DI swap, which references the

overnight interbank rate (DI), and for Mexico, it is the MXN MTH swap, which is based on the 28-day repo rate (TIIE). For emerging economies, the money market rates have a maturity of three months.[31] For the EA, UK, and US, we use twelve-month OIS rates, as these countries faced effective lower bound (ELB) constraints on their policy rates for a significant portion of our sample period. Shorter-term rates—such as three-month OIS rates—were similarly constrained by the ELB and therefore provided limited information about monetary policy surprises during those episodes. The twelve-month rates display more promising variation around policy meetings in those cases. In all cases, the data source is Bloomberg.

Inflation compensation. As our focus is on long-horizon inflation expectations, we use a 5-year, 5-year (5y5y) forward inflation compensation. Setting aside inflation risk premia, this measure should reflect investors' average inflation expectations over the five-year period beginning five years from now. The EA and the UK inflation compensation is based on inflation-linked swap rates tied to the HICP and RPI indices, respectively. These are downloaded from Bloomberg. For Brazil, Chile, Mexico, and the US, we obtain inflation compensation measures as the difference between nominal and inflation-linked bond yields. The US measures are the estimates in Gürkaynak, Sack, and Wright (2007), the Brazilian measures are computed by ANBIMA (the Brazilian Financial and Capital Markets Association) and downloaded via Bloomberg, while Jens Christensen kindly provided us with his estimates for Chile and Mexico (see Beauregard, Christensen, Fischer, and Zhu (2024) and Ceballos, Christensen, and Romero (2025)).

Survey-based inflation expectations. Ideally, we would like to have survey-based expectations that cover a 5y5y horizon similar to the inflation compensation measures. This measure is not available for all countries, but we use responses to the survey questions that elicit inflation expectations at the longest available horizon for each country. For Chile and Mexico we use surveys conducted by or on behalf of the countries' central banks. These surveys are conducted monthly and report average inflation expectations over different future horizons: one year starting three years from now for Chile,[32] and between five and eight years from now for Mexico. We use the Survey of Professional Forecasters for both the EA and the US, which asks for inflation expectations covering horizons five years ahead

[31] Our results for emerging market economies are robust to constructing monetary policy surprises as the difference between the realized policy rate announced after each central bank meeting and the expected rate, measured by the median economist forecast in the Bloomberg survey conducted prior to the meeting instead of the changes in the three months interest swap rates around the policy meetings.

[32] Until November 2019, the longest horizon available in the Chilean survey was the year ending two years from now, which we use until the longer horizon becomes available.

and the five years starting in five years respectively. Those surveys are only conducted quarterly.[33] For the UK, we use the survey conducted by Consensus Economics, which asks for inflation expectations covering the five years that starts five years from now. That survey was only conducted bi-annually until 2014 and quarterly after that. For all of the surveys, we date the surveys based on the last day to respond to the particular survey. We then aggregate all monetary policy surprises occurring between two consecutive editions of the surveys to implement our regression analysis. For monthly surveys, this typically involves a single policy surprise, while for quarterly surveys, multiple surprises often fall within the interval.

The central bank of Brazil conducts an interesting survey for our event study analysis; the "Focus survey". All participants—which are mostly banks and asset managers, consistent with the notion of a "professional forecaster" survey as we have been using for the other countries—can update their forecasts at any time and the central bank publishes the results on a daily basis. There is no requirement to update on a particular frequency, but the central bank encourages regular updating by running frequent forecast accuracy contests with a published ranking of the best forecasters. In addition, if a participant makes no update to their expectations in a six month period, the participant automatically drops out of the panel. This should also help make the responses representative at any point in time. In our analysis, we will use the longest forecast horizon available, the fourth fiscal year forward. For additional details on the survey see Gaglianone, Giacomini, Issler, and Skreta (2022) and Bonomo et al. (2024).

D.2 Central Bank Information Effect

We control for central bank information effects by focusing on the monetary event days where the interest rate swaps move in the opposite direction of stock prices. The idea is that monetary policy shocks would be expected to move the two in opposite directions, whereas a growth news shock would be expected to move them in the same directions. We use the S&P500 index for the US and the BOVESPA index for Brazil. We then re-run our regression using only those days.

The results are broadly consistent with our baseline results. We continue to see that long-run inflation expectations are very sensitive to monetary policy surprises in Brazil, with

[33] We could have used the responses to the Survey of Monetary Analysts and the "Desk survey" conducted by the ECB and the New York Fed, respectively. Those are both conducted eight times a year, just before policy meetings, but there is no, or almost no, variation in the long-horizon inflation expectations reported in those surveys. That indicates that inflation expectations are seen as well anchored, and that we should expect regression coefficients close to zero for those countries, if those surveys are representative.

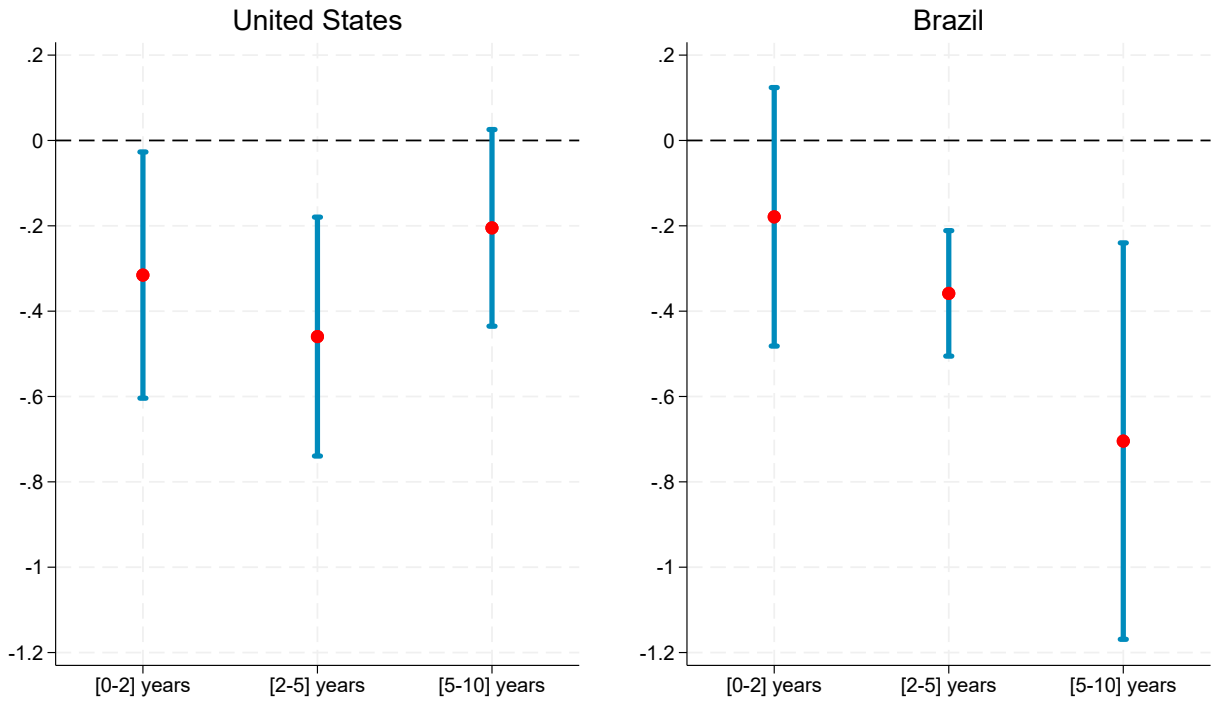


Figure 7: Controlling for central bank information effects: US vs. Brazil

Notes: The figure reports estimates of $b_{(k,s)}$(point estimates and 90% confidence intervals) in (33) for various horizons for the US and Brazil. To account for any potential central bank information effect, the chart uses only monetary events where interest rate swaps and stock prices move in opposite direction. The confidence intervals are constructed using robust standard errors. Inflation expectations are market-based and calculated from the difference between nominal and real bond yields. The sample covers 2010 through 2024.

unexpected interest rate hikes being associated with significant reductions in expectations about inflation in the long-run. The effect is statistically significant at the conventional level, and it is slightly more negative than in our baseline specification: a 100 basis points surprise in the three month interest rate swap rate causes, on average, a reduction of roughly 70 basis points (compared to 50 basis points in the baseline). The point estimate for the long-horizon regression for the US is slightly more negative than in our baseline specification, but still statistically insignificant. Consistent with our baseline results, we continue to find that a tightening surprise in Brazil is associated with a very persistent and sizable drop in inflation expectations, with little evidence of mean reversion. In the US, we see a hump-shaped response of inflation expectations, similar to the results in our baseline specification.