Hawkish Dove or Dovish Hawk?

Optimal Monetary Policy with Reputational Concerns[†]

Tomás E. Caravello[‡] Alex Carrasco-Martínez[§]
MIT MIT

Pedro Martínez-Bruera[¶]
MIT

Job Market Paper

First Version: April 30, 2025
This Version: October 3, 2025

Abstract: We study the design of monetary policy when the public learns about the policymaker's preference for inflation stabilization. Relative to discretion, the optimal policy reacts more strongly: the central bank signals a stronger commitment to inflation stability, thereby anchoring short-run expectations and reducing the future cost of disinflation. Using cross-sectional variation of private forecasts about U.S. inflation and the output gap, we document that the data are consistent with the mechanisms highlighted in our model. A quantitative exercise shows that a simple delegation problem can provide a robust implementation of the optimal policy.

Keywords: Monetary policy, reputation, learning, cost-push shocks.
JEL codes: E43, E52, E71

[†]We thank our advisors, Ivan Werning, Christian K. Wolf, and Ricardo Caballero, for guidance throughout this project. We also thank Olivier Blanchard, Javier García-Cicco, Emilio Espino, Gonzalo Huertas, Vicente Jimenez G., Jackson Mejia, Andy Neumeyer, Jim Poterba, Francisco Roldán, and David Romer for helpful comments and suggestions, as well as seminar participants at the MIT Macro Lunch, Universidad de Buenos Aires, Universidad de San Andrés, and Universidad Torcuato di Tella.

[‡]tomasec@mit.edu

[§]alexcm@mit.edu

[¶]pmbruera@mit.edu

> “Not only do expectations about policy matter, but, at least under current conditions, little else matters.”
>
> –Woodford (2003a)

1 Introduction

Reputation is central to the effectiveness of monetary policy. In practice, the common wisdom is to track it in real time by measuring long-run inflation expectations: if they remain anchored, credibility is secure. Yet, long-run expectations have remained anchored and stable (Kiley 2025), even during episodes of sharp fluctuations. For instance, during the COVID-19 pandemic, U.S. long-run inflation expectations barely moved. At the same time, perceptions about the Fed shifted dramatically: it was initially seen as prioritizing the recovery over inflation, and later had to tighten aggressively once inflation surged—a costly restoration of credibility that traditional measures of anchored expectations failed to anticipate.

This gap between stable long-run expectations and shifting credibility echoes a growing body of work documenting that perceptions of monetary policy fluctuate over the business cycle (Hamilton et al. 2011, Bauer et al. 2024, Bocola et al. 2024). Furthermore, accusations that central banks fell behind the curve in 2021-2022—despite anchored long-run expectations— suggest that long-run expectations are not the full picture. Anchored long-run expectations are a necessary condition for credibility, but they are not sufficient. They miss a crucial dimension: the public's perception of how the central bank reacts to shocks, which shapes the formation of *short-run* inflation expectations.

This gap motivates our analysis. We develop a framework that treats reputation as the public's perception of how the central bank reacts to shocks. This reaction is fundamentally shaped by the relative priority the central bank assigns to inflation stability: the more weight placed on inflation stability, the stronger the reaction. A central bank perceived as highly committed—a hawkish reputation— reduces the pass-through of shocks to short-run expectations; it *anchors* them, reducing the real costs of stabilization. In this sense, we provide a short- and medium-run analog of anchored expectations in Bernanke (2007): rather than asking whether short-run inflation shifts long-run expectations, we focus on how a central bank's reputation shapes the propagation of shocks through short- and medium-run expectations, and thus the trade-offs it faces in real time.

Our benchmark is a central bank that stabilizes inflation period-by-period, without internalizing the effects of actions on beliefs (which we call *myopic*). Our main result is that the optimal policy reacts more strongly to shocks than the myopic benchmark; it *overreacts*. By acting tougher on inflation today, the central bank builds credibility, anchors expectations, and lowers the inflation-output trade-off in the future.

Our theory provides a novel way of measuring reputation in the data. Specifically, we exploit cross-sectional variation in forecasts of U.S. inflation and output to develop a model-based measure of reputation. This measure is a natural complement to the traditional use of long-run inflation expectations, giving central banks another way to track reputation in real time. We document that our proposed mechanism is present. Consistent with the model's predictions, we find that hawkish reputation reduces the pass-through of shocks to expectations, and an unexpected tightening shifts the private sector's beliefs towards perceiving the central bank as more hawkish.

While the data support the predictions of our model, implementing the optimal policy is not straightforward. It requires tracking reputation and anticipating how the private sector interprets announcements and actions— tasks that are fragile if the central bank's model is misspecified. The classic answer to the long-run credibility problem, going back to Rogoff (1985), was to delegate policy to a hawkish central banker. Our quantitative analysis shows that the very same solution applies to the short-run credibility problem as well: appointing a more hawkish myopic central banker closely approximates the optimal policy and remains robust under misspecification. In this sense, short-run reputation introduces a new dimension of credibility, but the remedy is familiar.

FRAMEWORK We build a tractable model in which the private sector is uncertain about the central bank's preference for inflation stability, relative to output gap stability. This is captured by λ in a dual-mandate loss function $y^2 + \lambda\pi^2$, where y is the output gap and π is inflation.

The private sector does not know the central bank's weight on inflation stability, λ. To form forecasts, it must "close the model" by assuming a policy rule. It assumes the central bank is myopic—that is, it does not internalize how its actions affect beliefs. Beliefs with more mass on high values of λ imply the central bank is perceived to have a higher commitment to low and stable inflation.

Reputation impacts inflation expectations and the cost of stabilization. In response to cost-push shocks, policymakers face a trade-off between stabilizing the output gap and inflation. A hawkish central bank (high λ) prefers to stabilize inflation, whereas a dovish central bank (low λ) prefers a stable output gap. Thus, the central bank's response is informative about how it will react to future shocks. When designing monetary policy, the central bank internalizes how its actions affect the future private sector's beliefs.

Optimal Policy: Overreaction Our main result is that the optimal policy reacts more aggressively than the myopic benchmark. The optimal policy can be expressed as the myopic policy plus an additional term that reflects the incentive to strengthen the perceived commitment to price stability. By definition, the myopic central bank ignores this channel, while the optimal policy internalizes it. Consequently, interest rates move more aggressively, which amplifies the response of the output gap and, through the NKPC, dampens the response of inflation.

The logic of overreaction extends in several directions. First, the incentive to overreact is most substantial when beliefs are more sensitive, which in our case is when reputation is intermediate—neither clearly hawkish nor dovish. Second, following a positive cost-push shock, the central bank generates an unexpected tightening when it is perceived as dovish, and raises rates by less than expected when perceived as hawkish. Third, policy actions are tilted toward improving reputation when the central bank is perceived as dovish, and toward letting it decline when perceived as hawkish. Finally, small, persistent shocks justify building reputation, whereas large, short-lived shocks justify spending it.

Empirical Preview Our theoretical model has direct empirical implications. It provides a novel way of measuring reputation in the data from cross-sectional variation in forecasts: it is the slope of a regression of forecasts of the output gap on forecasts of inflation—a smaller slope signals greater perceived hawkishness.

Using the Blue Chip Survey of Financial Forecasts (BCFF), which provides individual forecasts of output and inflation for the U.S. economy, we estimate the Fed's reputation. We document three robust facts. First, a more hawkish reputation anchors inflation expectations: the pass-through of cost-push shocks to inflation expectations is lower when the central bank is perceived as hawkish. Second, reputation rises after hawkish monetary policy surprises: in response to a monetary tightening, the private sector updates its beliefs toward a more substantial commitment to stable inflation. Finally, reputation is not affected by cost-push shocks.

Our theoretical model also sheds light on the structural interpretation of perceived Taylor-rule coefficients. We show that while an increase in the perceived Taylor-rule coefficient of the output gap corresponds to a rise in the central bank's perceived responsiveness to demand-driven fluctuations, an increase in the coefficient of inflation may not correspond to an improvement in reputation. A central bank with a very strong and a very weak perceived commitment to stable inflation may have the same reduced form coefficient. In the data, the time-series correlation between our reputation measure and the perceived Taylor-rule coefficient on inflation is negative.

New problems, same solutions? In the quantitative section, we numerically evaluate the relative performance of delegating policy to a hawkish but myopic central banker, as in Rogoff (1985). Calibrated to generate the same long-run reputation as the optimal policy, this central banker tracks the optimal policy closely, in both first and second moments. Moreover, it outperforms the optimal policy when the central bank underestimates its ability to change the private sector's beliefs.

Our findings align with models where the central bank's actions reveal preferences. In contrast with those models, reputation is not about the perceived inflation bias, but rather about perceived hawkishness. Instead of focusing on the steady state, we focus on stabilization.

Related Literature This perspective connects to a broad literature on credibility, commitment, and expectations formation, while shifting the focus to dynamic short- and medium-run interactions. Our work mainly contributes to three literatures.

First, we shift the perspective of the central bank's reputation from perceived inflation bias to perceived hawkishness. The classic literature emphasizes the strategic interaction that produces inflation bias (Barro and Gordon 1983; Kydland and Prescott 1977). There can be incentives to build reputation over time when there is uncertainty about the central bank's inflation bias (Backus and Driffill 1985, Barro 1986, Canzoneri 1985, Vickers 1986). However, private information about the state of the economy may complicate inference and create perverse incentives (Cukierman and Meltzer 1986), and central banks may mimic types to avoid detection (King et al. (2008); Lu et al. (2016); Kostadinov and Roldán (2020)). Our contribution reframes this logic: by overreacting, the central bank mimics greater hawkishness than its true preferences, not because of time inconsistency, but to anchor expectations.

Second, our empirical measure of reputation connects directly to recent work showing that perceptions of monetary policy vary over the business cycle (Hamilton et al. 2011, Bauer et al. 2024, Bocola et al. 2024). Those papers interpret shifts in the perceived Taylor-rule coefficients as changes in perceived hawkishness. Yet, given the skepticism about whether central banks actually follow such a rule (see, e.g., Svensson 2003; Nakamura et al. 2025), it is not clear how these coefficients map to the central bank's preferences. We take a complementary approach: rather than inferring reputation from reduced-form coefficients, we develop a model-based measure that directly captures perceived hawkishness. This distinction clarifies why there is no one-to-one mapping between Taylor-rule coefficients and reputation: a central bank with a stronger reputation may not need to react as aggressively because expectations are already well anchored.

Third, we contribute to the literature on robust solutions to credibility problems. Rogoff (1985) proposed delegating policy to a hawkish central banker as a way to reduce inflation bias. We show numerically that this prescription is not only a suitable approximation of the optimal policy, but also robust: appointing a hawkish myopic central banker provides a close approximation to the optimal policy, particularly when the central bank underestimates how strongly beliefs respond to its actions.

Finally, the work closest to ours is the contemporaneous work by Bocola et al. (2025). Both papers study optimal policy when the private sector is uncertain about the central bank's relative weight on inflation stability. We take a complementary approach and reach distinct implications along three dimensions. First, in our model, the central bank overreacts relative to a myopic benchmark, irrespective of its type; in Bocola et al. (2025), overreaction arises only for hawkish types. Second, our benchmark for overreaction is within-economy, so counterfactuals hold the environment fixed and change only policy behavior. By contrast, Bocola et al. (2025) defines overreaction relative to a perfect-information economy. Third, we model reputation on the extensive margin: changes in reputation do not reflect a change in the probability of being one of two types, but rather a shift in the whole distribution of beliefs over all types.[1]

[1] There are also modeling differences. Bocola et al. (2025) adopts a two-type environment and assumes the private sector internalizes that the central bank follows its optimal policy. We allow for a continuum of types and, in our baseline, assume agents believe the central bank is myopic. However, as shown in Section 4, this distinction is not essential for our results.

Roadmap The rest of the paper is organized as follows. Section 2 develops the main mechanism in a three-period economy. Section 3 describes the economy. Section 4 derives the optimal policy and compares it to the myopic benchmark. Section 5 presents empirical evidence from U.S. forecast data. Section 6 provides a quantitative exercise evaluating the robustness of delegation to a hawkish but myopic central banker.

2 The Reputation Channel in a Three-Period Model

To build intuition, this section introduces a simple model of reputation as perceived hawkishness. Private-sector learning makes it optimal to trade a deeper recession today in exchange for a better inflation-output trade-off tomorrow. To illustrate the mechanism, we use a three-period New Keynesian model. The Phillips curve is given by

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t^P \left[\pi_{t+1}\right] + \varepsilon_t \qquad t = 0, 1, 2,$$

where y_t is the output gap, π_t is inflation, and ε_t is a cost-push shock. $\mathbb{E}_2^P[\cdot]$ denotes private-sector expectations, and $\mathbb{E}_t^P[\pi_3] = 0$. The central bank has a dual mandate over the output gap and inflation,

$$\mathcal{W}_0 = -\frac{1}{2}\mathbb{E}_0^{CB}\left[\sum_{t=0}^{2}\beta^t\left(y_t^2 + \lambda\pi_t^2\right)\right]$$

where $\mathbb{E}_t^{CB}[\cdot]$ denotes the central bank's expectations, and λ its relative weight on inflation stability. The private sector knows the structure of the economy but not the value of λ. The true preference is fixed, but what evolves is the private sector's perception of it: its reputation. The central bank maximizes $\mathcal{W}_0$ subject to the Phillips curve. This toy model previews the trade-off at the heart of our results for optimal policy. While the full framework generalizes these insights, the simplified setting captures the main logic.

Solving by backward induction yields, at $t = 2$,

$$y_2 = -\underbrace{\frac{\kappa\lambda}{1+\kappa^2\lambda}}_{\psi^y}\varepsilon_2 \qquad \pi_2 = \underbrace{\frac{1}{1+\kappa^2\lambda}}_{\psi^\pi}\varepsilon_2.$$

At $t = 1$, given y_2 and π_2

$$y_1 = -\psi^y\left(\varepsilon_1 + \beta\mathbb{E}_1^P\left[\pi_2\right]\right) \qquad \pi_1 = \psi^\pi\left(\varepsilon_1 + \beta\mathbb{E}_1^P\left[\pi_2\right]\right).$$

At $t = 0$

$$y_0 = \overbrace{-\psi^y \left(\varepsilon_0 + \beta \mathbb{E}_1^P \left[\pi_0\right]\right)}^{\text{Myopic Stabilization}} \overbrace{+\psi^\pi \beta \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_1^P \left[\pi_2\right]}{\partial y_0} y_1\right]}^{\text{Intertemporal Smoothing}}$$

$$\pi_0 = \overbrace{\psi^\pi \left(\varepsilon_0 + \beta \mathbb{E}_1^P \left[\pi_0\right]\right)}^{\text{Myopic Stabilization}} \overbrace{+\kappa \psi^\pi \beta \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_1^P \left[\pi_2\right]}{\partial y_0} y_1\right]}^{\text{Intertemporal Smoothing}}.$$

Two terms characterize the optimal policy. The first, *myopic stabilization*, is the policy of a central bank without commitment, stabilizing the output gap and inflation each period. The second, *intertemporal smoothing*, captures the ability of current policy to influence future inflation expectations. Formally, the effect of the current output gap on future expected inflation is given by

$$\frac{\partial \mathbb{E}_1^P \left[\pi_2\right]}{\partial y_0} = \frac{\partial \mathbb{E}_1^P \left[\psi^\pi\right]}{\partial y_0} \mathbb{E}_1^P \left[\varepsilon_2\right].$$

Because ψ^π decreases with λ, the pass-through of shocks to expectations depends on the central bank's perceived hawkishness. Intuitively, if the central bank is perceived as more hawkish, agents expect shocks to pass through less strongly to expectations. This dependence of pass-through to beliefs is precisely where the reputation channel operates: a hawkish stance shifts beliefs toward higher values of λ, dampening the effect of shocks on expectations.

We refer to this as the *reputation channel*. When the private sector perceives the central bank as highly committed to stable inflation, it has a *hawkish* reputation; otherwise, it has a *dovish* one. Reputation shapes the sensitivity of inflation expectations to shocks. Having isolated the mechanism in a three-period model, we generalize it to an infinite-horizon New Keynesian economy.

3 The Economy

This section lays out the core equations that define equilibrium in our economy.[2]

[2] We provide a detailed description of the economy in Appendix A.

Private Sector Block The competitive equilibrium in this economy is summarized by a dynamic IS equation and a New Keynesian Phillips Curve (NKPC):

$$y_t = \mathbb{E}_t^P [y_{t+1}] - \frac{1}{\sigma} \left(i_t - \mathbb{E}_t^P [\pi_{t+1}] - \mathbb{E}_t^{CB} [r_t^n] \right) \tag{1}$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t^P [\pi_{t+1}] + \varepsilon_t \tag{2}$$

where y_t is the output gap, π_t is inflation, i_t is the risk-free nominal interest rate, r_t^n is the real natural rate, and ε_t is a markup (cost-push) shock. Following Woodford (2003a), demand and productivity shocks enter through the natural rate, r_t^n:

$$r_t^n \equiv \rho + \upsilon \mathbb{E}_t^P [\Delta a_{t+1}] - \mathbb{E}_t^P [\Delta z_{t+1}]$$

A positive demand shock increases the natural rate, while a positive productivity shock lowers it. The expectations operator $\mathbb{E}^P [\cdot]$ denotes private-sector beliefs, which satisfy the Law of Iterated Expectations. We assume both households and firms have the same information set. Equations (1) and (2) fix notation and make clear where private-sector expectations enter the equilibrium. Reputation will matter through its effect on expectations.

Central Bank and Monetary Policy Regime The central bank's dual mandate over the output gap and inflation is represented by the welfare loss function

$$\mathcal{W}_t = -\frac{1}{2} \mathbb{E}_t^{CB} \left[\sum_{k=0}^{\infty} \beta^k \left(y_{t+k}^2 + \lambda \pi_{t+k}^2 \right) \right] \tag{3}$$

where $\mathbb{E}_t^{CB} [\cdot]$ denotes the central bank's expectations. The parameter λ measures the weight placed on inflation stabilization relative to output-gap stabilization, and the objective will be optimized subject to the private-sector block, with expectations playing the key role in transmitting reputation.

This specification, while stylized, is standard: it arises from a second-order approximation to household welfare in the canonical New Keynesian model and is widely used in applied policy analysis (e.g., Federal Reserve Board 2016, Barnichon and Mesters 2023).

To avoid full revelation of the central bank's preferences from its actions, we include the following information friction: At the start of period t, the central bank announces the nominal interest rate, i_t, before the current demand shock is realized. Letting ν_t denote the central bank's forecast error for the demand shock, we write

$$\mathbb{E}_t^P \left[\Delta z_{t+1}\right] = \mathbb{E}_t^{CB} \left[\Delta z_{t+1}\right] + \mathbb{E}_t^P \left[\nu_t\right] \tag{4}$$

We assume $\mathbb{E}_t^P [\nu_t] = 0$, so the forecast is unbiased.

Private Sector Expectations' Formation Process The private sector knows the full structure of the economy, but not the central bank's relative weight on inflation, λ, nor their demand forecast $\mathbb{E}_t^{CB} [\Delta z_{t+1}]$. They believe the central bank is myopic, and share a common prior, μ_t, over possible values of $\lambda \in (0, \infty)$. In particular, they assume a central bank of type $\lambda = \tilde{\lambda}$ maximizes the dual mandate (3) subject to the ex-ante equilibrium conditions, (1) and (2)

$$\tilde{y}_t = \mathbb{E}_t^{CB} \left[\mathbb{E}_t^P \left[y_{t+1}\right]\right] - \frac{1}{\sigma} \left(i_t - \mathbb{E}_t^{CB} \left[\mathbb{E}_t^P \left[\pi_{t+1}\right]\right] - \mathbb{E}_t^{CB} \left[\tilde{r}_t^n\right]\right) \tag{5}$$

$$\tilde{\pi}_t = \kappa \tilde{y}_t + \beta \mathbb{E}_t^{CB} \left[\mathbb{E}_t^P \left[\pi_{t+1}\right]\right] + \varepsilon_t \tag{6}$$

where $\tilde{y}_t := \mathbb{E}_t^{CB} [y_t]$ and $\tilde{\pi}_t := \mathbb{E}_t^{CB} [\pi_t]$ denote the allocation the central bank seeks to implement.[3]

We impose no anticipated learning (Marcet and Sargent 1989; Eusepi and Preston 2018; Kreps 1998; Evans and Honkapohja 2001): agents forecast the future using today's beliefs, without accounting for the fact that they will update them in the future. This assumption avoids the infinite regress of beliefs about future beliefs and keeps the problem tractable.

The private sector's belief about λ determines how shocks today feed into expectations. Let $\psi^y := \frac{\kappa\lambda}{1+\kappa^2\lambda}$ and $\psi^\pi := \frac{1}{1+\lambda\kappa^2}$. From the private sector's perspective, the following result describes how a myopic central bank of type $\lambda = \tilde{\lambda}$ reacts to shocks.

Lemma 1. *Under no anticipated learning, a myopic central bank with $\lambda = \tilde{\lambda}$ implements:*

$$\tilde{y}_t \left(\tilde{\lambda}\right) = -\psi^y \left(\tilde{\lambda}\right) \sum_{s=0}^{\infty} \left(\beta \mathbb{E}_t^P \left[\psi^\pi\right]\right)^s \mathbb{E}_t \left[\varepsilon_{t+s}\right] = -\psi^y \left(\tilde{\lambda}\right) X_t \tag{7}$$

$$\tilde{\pi}_t \left(\tilde{\lambda}\right) = \psi^\pi \left(\tilde{\lambda}\right) \sum_{s=0}^{\infty} \left(\beta \mathbb{E}_t^P \left[\psi^\pi\right]\right)^s \mathbb{E}_t \left[\varepsilon_{t+s}\right] = \psi^\pi \left(\tilde{\lambda}\right) X_t \tag{8}$$

[3] We impose this assumption for tractability; in Appendix B we show that all conclusions continue to hold once it is relaxed.

where $\psi^y\left(\tilde{\lambda}\right) = \frac{\kappa\tilde{\lambda}}{1+\kappa^2\tilde{\lambda}}$, $\psi^\pi\left(\tilde{\lambda}\right) = \frac{1}{1+\kappa^2\tilde{\lambda}}$, *and* $X_t = \sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right]$.

Proof. See Appendix A ■

Given these perceived allocations, the object of interest is how agents form expectations. The k-period-ahead forecasts of output gap and inflation are:

$$\mathbb{E}_t^P\left[y_{t+k}\right] = -\mathbb{E}_t^P\left[\psi^y\right]\mathbb{E}_t^P\left[X_{t+k}\right] + \mathbb{E}_t^P\left[\eta_{t+k}\right] \tag{9}$$

$$\mathbb{E}_t^P\left[\pi_{t+k}\right] = \mathbb{E}_t^P\left[\psi^\pi\right]\mathbb{E}_t^P\left[X_{t+k}\right] + \kappa\mathbb{E}_t^P\left[\eta_{t+k}\right] \tag{10}$$

where $\eta_t := -\frac{1}{\sigma}\nu_t$ denotes the effect of the central bank's forecast error on the output gap. Under our modeling assumptions, $\mathbb{E}_t^P\left[\eta_{t+k}\right] = 0$, but it need not be the case in the data.[4] These forecasts depend on the expected values of ψ^y and ψ^π, the parameters through which reputation shapes the transmission of shocks.

The Central Bank's reputation Private-sector expectations for inflation and the output gap depend on beliefs about λ only through their effect on $\mathbb{E}_t^P\left[\psi^y\right]$ and $\mathbb{E}_t^P\left[\psi^\pi\right]$. Since $\psi^y = \kappa^{-1}\left(1-\psi^\pi\right)$, the two move in opposite directions: a higher $\mathbb{E}_t^P\left[\psi^\pi\right]$ means a lower $\mathbb{E}_t^P\left[\psi^y\right]$, and vice versa: stabilizing inflation comes at the cost of volatility of the output gap. We refer to $\mathbb{E}_t^P\left[\psi^\pi\right]$ as the *reputation* of the central bank, and say reputation *improves* when $\mathbb{E}_t^P\left[\psi^\pi\right]$ falls, that is, when the central bank is perceived as more hawkish. Shifts in reputation correspond to movements in the full distribution of beliefs over λ, rather than just a change in the odds of any single type.

From (9)-(10), a lower $\mathbb{E}_t^P\left[\psi^\pi\right]$ shapes expectations through two channels:

1. **Direct:** holding $\mathbb{E}_t^P\left[X_{t+k}\right]$ fixed, the private sector expects less inflation and a deeper recession in equilibrium.

2. **Indirect:** a stronger reputation makes the private sector discount future shocks more heavily: if they are confident that the central bank will stabilize inflation in the future, those shocks matter less for today's expectations.

From the central bank's perspective, a lower value of $\mathbb{E}_t^P\left[\psi^\pi\right]$ is always desirable: If the private sector is confident that the central bank will stabilize inflation, expectations barely

[4] We demonstrate in Appendix C that the case with evolving beliefs over η_{t+k} is isomorphic to a setting where the private sector learns about the central bank's inflation bias.

respond to shocks.[5] Achieving and maintaining such a reputation is costly, and the optimal policy trades off these costs against the benefits.

To formalize how beliefs about λ translate into reputation, we establish two results:

Lemma 2. *Suppose the private sector has beliefs μ_t over λ. Consider a different set of beliefs μ'_t that first-order stochastically dominates μ_t. Then, $\mathbb{E}^P_t[\psi^\pi]$ is lower under μ'_t than under μ_t.*

Proof. See Appendix A ■

Proposition 2 establishes that if beliefs place more weight on larger values of λ, the central bank is perceived as more hawkish. Therefore, expected inflation is lower for any given sequence of shocks. Figure 1 plots the case where first-order stochastic dominance leads to an improvement in reputation.

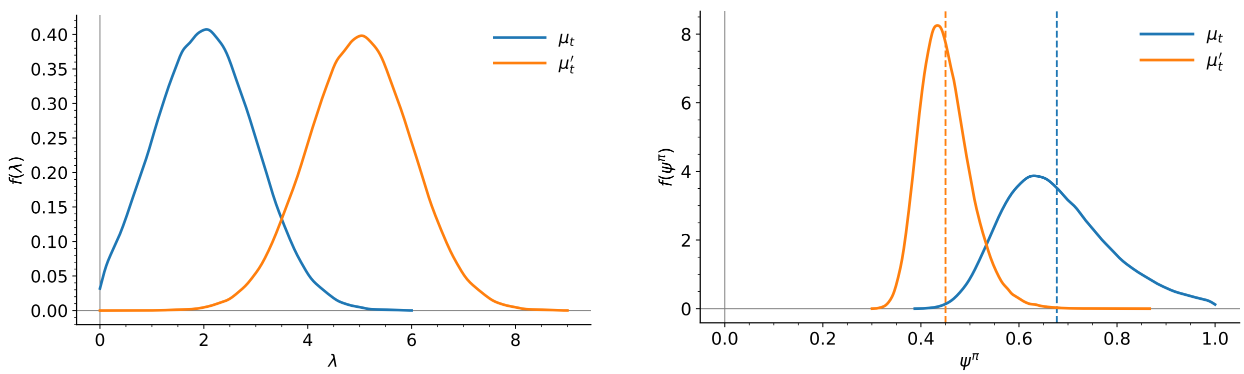

(a) Beliefs over λ **(b)** Beliefe over ψ^π

Figure 1: First-order stochastic dominance: $\mu'_t \overset{FOSD}{\succsim} \mu_t$

Lemma 3. *Suppose the private sector has beliefs μ_t over λ. Consider a different set of beliefs μ'_t that is a mean-preserving spread of μ_t. Then, $\mathbb{E}^P_t[\psi^\pi]$ is lower under μ_t that under μ'_t*

Proof. See Appendix A ■

Lemma 3 says that greater uncertainty about the central bank's preferences raises expected inflation. Because ψ^π is convex in λ, the private sector prices in the possibility of a dovish central bank. When in doubt, they expect softer responses to shocks.

[5]This may no longer be true at the Zero Lower Bound. A more dovish reputation can serve as a way to get out of the ZLB. A larger pass-through of shocks on inflation expectations can offset the effect of a negative natural interest rate. We explore this situation in Appendix C.

We have characterized how current beliefs shape reputation and expectations; next, we examine how these beliefs are updated over time.

Timing and Belief Updating At the start of each period, the central bank privately forecasts future shocks. Based on these forecasts, it sets the nominal interest rate for that period. The private sector observes the policy choice but not the bank's internal forecast, takes its consumption and production decisions, and only then updates its beliefs once outcomes for period t are realized. In other words, consumption and production decisions come first, and belief updating about the bank's preferences comes afterward.[6] Figure 2 summarizes the sequence.

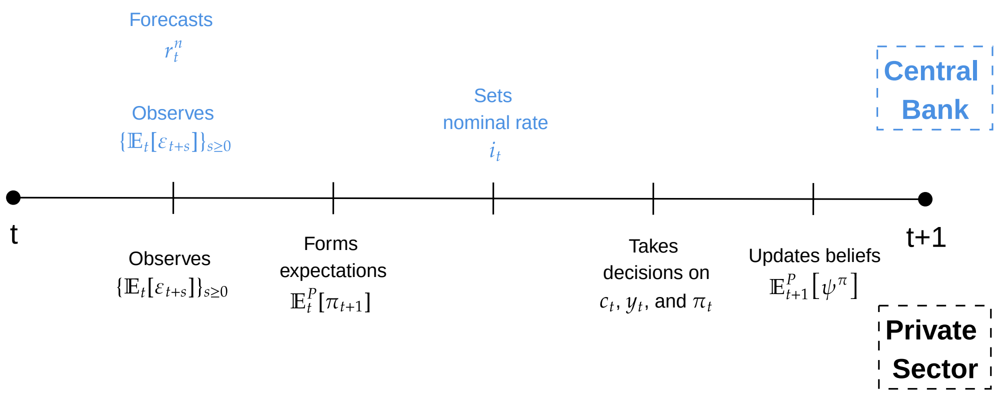

Figure 2: Timing within period

Given beliefs μ_t, households and firms update beliefs each period as follows:

1. The cost-push shocks are realized $\{\mathbb{E}_t\left[\varepsilon_{t+s}\right]\}_{s\geqslant 0}$.

   - Under μ_t, the public believes that a central bank of type $\lambda = \tilde{\lambda}$ chooses the allocations $\tilde{y}_t^M\left(\tilde{\lambda}\right)$ and $\tilde{\pi}_t^M\left(\tilde{\lambda}\right)$ given by (7) and (8), leading to the ex-post realizations
$$\tilde{y}_t^M\left(\tilde{\lambda}\right) + \eta_t \quad \text{and} \quad \tilde{\pi}_t^M\left(\tilde{\lambda}\right) + \kappa\eta_t$$
   - The central bank sets the policy rate i_t to achieve allocations $\tilde{y}_t$ and $\tilde{\pi}_t$

[6] We maintain this assumption for tractability in the main text. In Appendix B we relax it and show our results continue to hold.

2. The true realizations are

$$y_t = \tilde{y}_t + \eta_t \quad \text{and} \quad \pi_t = \tilde{\pi}_t + \kappa\eta_t$$

3. Upon observing π_t and y_t, the public cannot tell whether the deviations from their expectations are due to the forecast error η_t or a (myopic) central bank's type λ. A myopic central bank of type $\lambda = \tilde{\lambda}$ would choose $\tilde{\pi}^M\left(\tilde{\lambda}\right) = \psi^\pi\left(\tilde{\lambda}\right)X_t$, where $X_t = \sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t\left[\psi^\pi\right]\right)^s\mathbb{E}_t^P\left[\varepsilon_{t+s}\right]$. The private sector infers that such type's forecast error must have been

$$\eta_t\left(\tilde{\lambda}\right) = \kappa^{-1}\left(\pi_t - \psi^\pi\left(\tilde{\lambda}\right)X_t\right) = \kappa^{-1}\left(\tilde{\pi}_t - \psi^\pi\left(\tilde{\lambda}\right)X_t + \kappa\eta_t\right) \tag{11}$$

Beliefs are updated via Bayes' rule:

$$Pr\left(\tilde{\lambda}\middle|\pi_t, y_t\right) = \frac{Pr\left(\pi_t, y_t|\,\tilde{\lambda}\right)}{Pr\left(\pi_t, y_t\right)}Pr\left(\tilde{\lambda}\right) \Longleftrightarrow \mu_{t+1}\left(\tilde{\lambda}\right) = \frac{f_\eta\left(\eta_t\left(\tilde{\lambda}\right)\right)}{\int f_\eta\left(\eta_t\left(a\right)\right)\mu_t(a)da}\mu_t\left(\tilde{\lambda}\right) \tag{12}$$

where $f_\eta\left(\cdot\right)$ is the time-invariant probability distribution of η_t.[7]

Because ψ^π is a sufficient statistic for how λ affects inflation expectations (Proposition 1), it is equivalent, and more transparent, to describe beliefs over ψ^π rather than λ.

We assume $\eta_t \sim \mathcal{N}\left(0, \tau_\eta^{-1}\right)$ and that the prior over ψ^π is a truncated normal on $[0,1]$:

$$\tilde{\psi}^\pi \sim \Psi_t\Big|_{\Psi_t\in[0,1]}, \quad \Psi_t \sim \mathcal{N}\left(\overline{\psi}_t, \tau_t^{-1}\right)$$

Lemma 4. *Under these assumptions, the prior is conjugate, and*

$$\mathbb{E}_t^P\left[\psi^\pi\right] = \overline{\psi}_t - \frac{1}{\sqrt{\tau_t}}\frac{\phi\left(\sqrt{\tau_t}\left(1-\overline{\psi}_t\right)\right) - \phi\left(-\sqrt{\tau_t}\overline{\psi}_t\right)}{\Phi\left(\sqrt{\tau_t}\left(1-\overline{\psi}_t\right)\right) - \Phi\left(-\sqrt{\tau_t}\overline{\psi}_t\right)} \tag{13}$$

where ϕ and Φ are the standard normal PDF and CDF, and

$$\frac{\mathbb{E}_t^P\left[\psi^\pi\right]}{\partial\overline{\psi}_t} = \tau_t\mathbb{V}_t^P\left[\psi^\pi\right]$$

[7] Out conclusions are unchanged if forecast errors are expressed in deviations from expected inflation, $\tilde{\pi}_t - \tilde{\pi}_t^M$, expected output gap $\tilde{y}_t - \tilde{y}_t^M$, or the first-order condition $\tilde{y}_t^M + \lambda\kappa\tilde{\pi}_t^M$. We use inflation deviations from simplicity

moreover, $\mathbb{V}_t^P[\psi^\pi]$ is hump-shaped, and peaks when $\overline{\psi}_t = \frac{1}{2}$.

Proof. This result is a direct consequence of the truncated normal functional form assumption. ■

Lemma 4 establishes that the private sector's Bayesian learning has a conjugate prior, whose properties are illustrated in Figure 3.

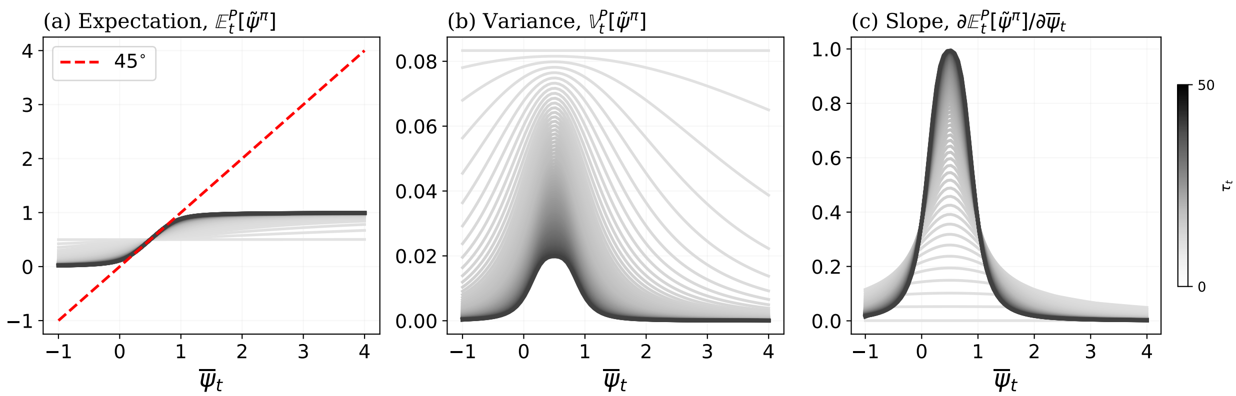

Figure 3: Prior beliefs as a function of $\overline{\psi}_t$

A dovish central bank accommodates most cost-push shocks, so ψ^π is near one; a hawkish central bank prioritizes inflation stability, so ψ^π is close to zero, unlike two-type models (e.g. Backus and Driffill 1985, Bocola et al. 2025), we allow a continuum of types. Hence, reputation varies along an extensive margin: beliefs shift the entire distribution over ψ^π, not just the probability of a given type. Reputation is, then, the perception of how hawkish the central bank is, rather than the probability that the central bank is the hawkish one.

Under the above structure, $\Psi_{t+1} \sim \mathcal{N}\left(\overline{\psi}_{t+1}, \tau_{t+1}^{-1}\right)$ with

$$\overline{\psi}_{t+1} = \omega_t X_t^{-1}\left(\tilde{\pi}_t + \kappa\eta_t\right) + \left(1-\omega_t\right)\overline{\psi}_t \tag{14}$$

where

$$\omega_t = \frac{\left(\kappa^{-1}X_t\right)^2 \tau_\eta}{\tau_t + \left(\kappa^{-1}X_t\right)^2 \tau_\eta} \quad \text{and} \quad \tau_{t+1} = \tau_t + \left(\kappa^{-1}X_t\right)^2 \tau_\eta$$

ω_t is the Kalman gain: it rises when shocks are large relative to noise, but the *effective* updating weight $\omega_t X_t^{-1}$ is hump-shaped in X_t because large shocks dilute the signal about its type.

Policy surprises affect beliefs through $\tilde{\pi} - \mathbb{E}_t^P[\psi^\pi] X_t$: in the presence of inflationary shocks, implementing lower inflation than expected shifts beliefs towards more hawkish types. Using, (14),

$$\overline{\psi}_{t+1} = \overline{\psi}_t + \omega_t \left[\left(\mathbb{E}_t^P[\psi^\pi] - \overline{\psi}_t \right) + \kappa \frac{\eta_t}{X_t} + \left(\frac{\tilde{\pi}_t}{X_t} - \mathbb{E}_t^P[\psi^\pi] \right) \right] \tag{15}$$

Reputation only improves if

$$\tilde{\pi}_t + \kappa \eta_t < \overline{\psi}_t X_t$$

When reputation is already high, $\mathbb{E}_t^P[\psi^\pi] < \frac{1}{2}$, mean reversion ($\mathbb{E}_t^P[\psi^\pi] > \overline{\psi}_t$) implies that only a considerable hawkish surprise can shift beliefs further. These dynamics describe how policy choices today affect reputation, which will be central to the optimal policy problem we study next.

4 Reputation and Optimal Policy

We now turn to the central bank's optimal policy when its reputation evolves endogenously. At each date t, the central bank chooses the (ex-ante) output gap, $\tilde{y}_t$, and inflation, $\tilde{\pi}_t$, balancing the conventional stabilization trade-off with the benefits from influencing how it is perceived.

Formally, given the private sectors beliefs, μ_t, the central bank maximizes

$$\mathcal{W}_t = -\frac{1}{2} \mathbb{E}_t^{CB} \left[\sum_{s=0}^{\infty} \beta^s \left(\tilde{y}_t^2 + \lambda \tilde{\pi}_t^2 \right) \right] \tag{PP}$$

subject to

(i) The New Keynesian Phillips Curve (NKPC)

$$\tilde{\pi}_t = \kappa \tilde{y}_t + \beta \mathbb{E}_t^P[\tilde{\pi}_{t+1}] + \varepsilon_t \qquad \forall t$$

which links today's inflation to the output gap, expected inflation in $t+1$, and cost-push shocks.

(ii) The private sector's learning structure,

$$\mathbb{E}_t^P[\tilde{\pi}_{t+1}] = \mathbb{E}_t^P[\psi^\pi] \mathbb{E}_t^P[X_{t+1}]$$

$$\mathbb{E}_t^P[\psi^\pi] = \overline{\psi}_t - \frac{1}{\sqrt{\tau_t}} \frac{\phi\left(\tau_t\left(1-\overline{\psi}_t\right)\right) - \phi\left(-\tau_t\overline{\psi}_t\right)}{\Phi\left(\tau_t\left(1-\overline{\psi}_t\right)\right) - \Phi\left(-\tau_t\overline{\psi}_t\right)}$$

$$\overline{\psi}_{t+1} = \overline{\psi}_t + \omega_t \left[\left(\mathbb{E}_t^P[\psi^\pi] - \overline{\psi}_t\right) + \kappa\frac{\eta_t}{X_t} + \left(\frac{\tilde{\pi}_t}{X_t} - \mathbb{E}_t^P[\psi^\pi]\right)\right]$$

where $\omega_t = \frac{\left(\kappa^{-1}X_t\right)^2\tau_\eta}{\tau_t + \left(\kappa^{-1}X_t\right)^2\tau_\eta}$ and $\tau_{t+1} = \tau_t + \left(\kappa^{-1}X_t\right)^2\tau_\eta$

We do not need to write the dynamic IS equation (5) explicitly. Given any desired path for (ex-ante) inflation and the output gap, the central bank can choose the interest rate to implement it.[8] More importantly, Forward Guidance is irrelevant. At time t, the private sector's inflation expectations depend only on the central bank's current reputation and cost-push shocks. Promises about the future path of variables have no effect, making the policy time-consistent.

The formal derivation is in Appendix B. Here, we focus on the economics. Policy is no longer just about the current inflation-output trade-off; it is also about shaping the trade-off it will face in the future.

Overreaction Our main result is that the optimal policy reacts more aggressively than the myopic benchmark. The reason is simple: reputation acts as a technology that lets the central bank trade a deeper recession today for a better inflation-output trade-off tomorrow. With positive cost-push shocks, the myopic problem is the familiar trade-off between inflation and recession. But here, the central bank's actions also reveal information about its preferences: conditional on a forecast error η_t, higher inflation signals a lower taste for inflation stabilization, that is, lower λ. The optimal output gap satisfies

$$\tilde{y}_t\left(\tilde{\psi}^\pi\right) = \overbrace{\tilde{y}_t^M\left(\tilde{\psi}^\pi\right)}^{\text{Myopic Stabilization}} + \overbrace{\tilde{\psi}^\pi \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB}\left[\beta \frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+1+s}\right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s}\left(\tilde{\psi}^\pi\right)\right]}^{\text{Intertemporal Smoothing}} \tag{16}$$

This equation has two components:

1. **Myopic Stabilization:** responds to shocks without accounting for reputation, reflecting the textbook inflation-output trade-off.

[8] We assume the Zero Lower Bound never binds.

2. **Intertemporal smoothing:** By acting hawkish, the central bank can improve the inflation-output trade-off in the future.

If the central bank strengthens its reputation, future shocks will have a smaller impact on inflation expectations. This channel creates incentives to overreact. If shocks are persistent, a current shock is informative about the future, so acting aggressively reduces their damage in the future. Moreover, even if the current shock is *iid*, the central bank wants to insure against shocks in the future.

The policy trade-off is therefore between a deeper *current* recession and a milder *future* one. Relative to the myopic benchmark, the optimal policy calls for a larger recession. For inflation, the logic is reversed: a larger recession today means, by the NKPC, lower inflation than the myopic benchmark:

$$\tilde{\pi}_t\left(\tilde{\psi}^\pi\right) = \tilde{\pi}_t^M\left(\tilde{\psi}^\pi\right) + \kappa\tilde{\psi}^\pi \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB}\left[\beta \frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+1+s}\right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s}\left(\tilde{\psi}^\pi\right)\right] \tag{17}$$

Proposition 1. *When there is uncertainty about cost-push shocks, then* $|\tilde{y}_t| > \left|\tilde{y}_t^M\right|$ *and* $|\tilde{\pi}_t| < \left|\tilde{\pi}_t^M\right|$

Proof. See Appendix B ■

More formally, we can express $\frac{\partial \mathbb{E}_{t+s}^P[\tilde{\pi}_{t+s+1}]}{\partial \tilde{\pi}_t}$ as

$$\frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_t} = \frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial \mathbb{E}_{t+s}^P\left[\psi^\pi\right]} \frac{\partial \mathbb{E}_{t+s}^P\left[\psi^\pi\right]}{\partial \overline{\psi}_{t+s}} \frac{\partial \overline{\psi}_{t+s}}{\partial \tilde{\pi}_t} \tag{18}$$

The first two terms capture how reputation affects inflation expectations: a higher (lower) reputation lowers (raises) expected inflation in response to positive shocks. Multiplying by $\tilde{y}_{t+s}$ yields a negative product: the optimal response to a positive shock is a recession, and a worse reputation amplifies both inflation expectations and therefore the required contraction. The third term reflects how today's inflation affects reputation: with $X_t > 0$, higher inflation damages reputation; with $X_t < 0$, it improves it.

The policy implication of Proposition 1 is simple: the central bank must raise interest rates more than the myopic benchmark. In this model, a deeper recession can only be caused by higher rates.[9] The NKPC then converts the output gap into lower inflation. Figure 4

[9]Under commitment, this would correspond to committing to a path of higher rates.

plots the interest rate and inflation.

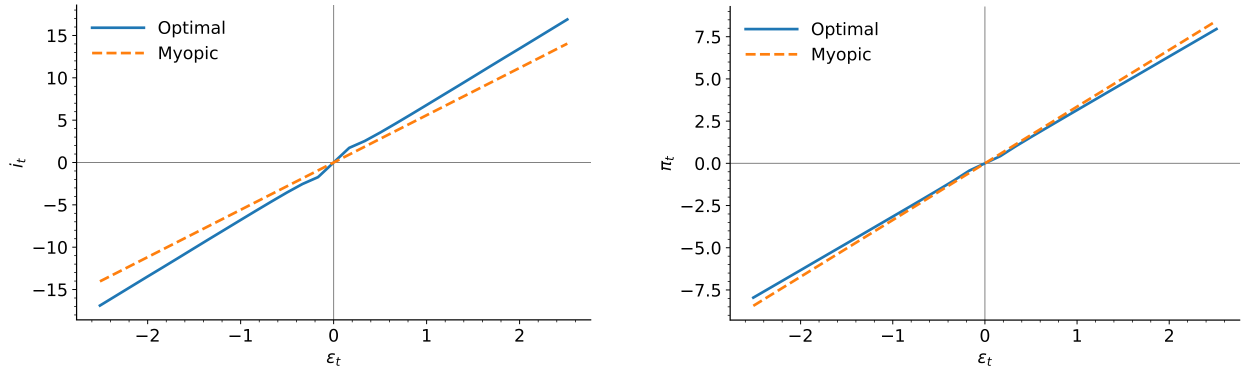

(a) Policy function for the interest rate i_t

(b) Policy function for inflation π_t

Figure 4: Policy functions. We assume that the shock ε_t follows an AR(1) process.

When shocks are persistent, reputation works in the opposite direction from commitment. Under commitment, the central bank smooths the recession over time, trading a smaller current recession for a larger future one. With reputational concerns, this logic flips: a larger current contraction signals a high λ, resulting in a softer one in the future. This result stands in stark contrast to the traditional reputation literature, where the optimal output-gap policy shifts from discretion toward commitment. In our framework, reputation pushes policy even further in that direction.

The logic of overreaction extends in several directions, which we explore below.

State-dependence Overreaction is strongest when the central bank's reputation sits in the middle ground: neither hawkish nor dovish, but uncertain. From (16), overreaction grows when inflation expectations are more sensitive to the actions of the central bank. Under our functional form assumptions, (18) becomes

$$\frac{\partial \mathbb{E}^P_{t+s}\left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_t} = \frac{\partial \mathbb{E}^P_{t+s}\left[\tilde{\pi}_{t+s+1}\right]}{\partial \mathbb{E}^P_{t+s}\left[\psi^\pi\right]} \tau_{t+s} \mathbb{V}^P_{t+s}\left[\psi^\pi\right] \frac{\partial \overline{\psi}_{t+s}}{\partial \tilde{\pi}_t}$$

where $\mathbb{V}^P_{t+s}\left[\psi^\pi\right]$ is the dispersion of beliefs about ψ^π. Taking current uncertainty as a proxy for future uncertainty, this expression shows that expectations are most sensitive when belief dispersion is high.

Belief dispersion peaks when reputation is in doubt: the private sector is unsure whether the central bank is hawkish or dovish. At the extremes, actions are uninformative. Hawkish types are certainly not dovish, and vice versa. But in the middle ground, actions become informative, and the incentive to overreact is most substantial.

Proposition 2. *The output gap overreaction* $\left|\tilde{y}_t - \tilde{y}_t^M\right|$ *is increasing in* $\mathbb{E}_t^P\left[\psi^\pi\right]$ *for small values of* $\mathbb{E}_t^P\left[\psi^\pi\right]$ *and decreases for large values of* $\mathbb{E}_t^P\left[\psi^\pi\right]$*.*

Proof. See Appendix B ■

Figure 5 plots the policy rate overreaction as a function of the central bank's reputation in response to a positive cost-push shock. When the public is genuinely convinced that the bank is dovish, does the incentive to overreact fade simply because it becomes tough to convince them otherwise?

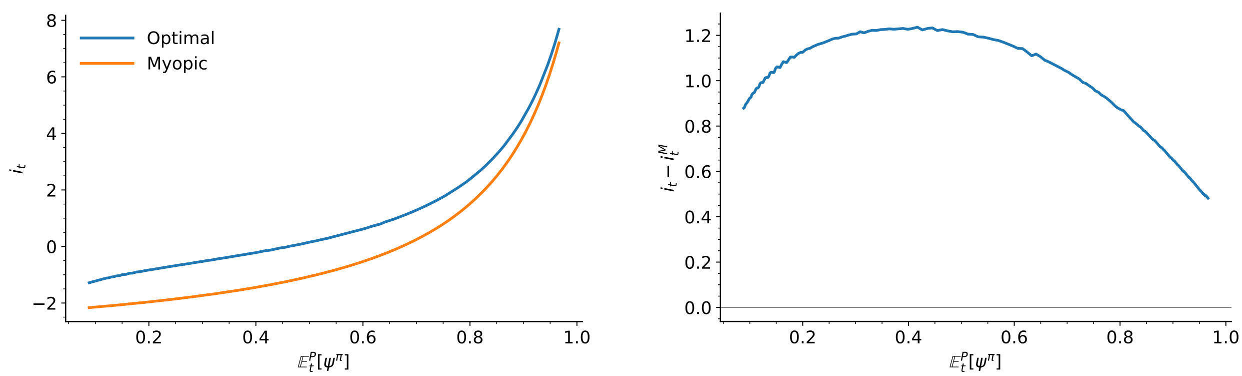

(a) Policy function for the interest rate i_t **(b)** Interest rate overreaction $i_t - i_t^M$

Figure 5: Policy rate overreaction and the central bank's reputation. Positive shock $\varepsilon_t > 0$

So far, we have compared the optimal policy to the myopic benchmark. We now turn to a different comparison: how it departs from what the private sector expects.

Hawkish if (Perceived as) Dovish, Dovish if Hawkish Overreaction means responding more aggressively than the myopic benchmark, but this need not produce a hawkish surprise relative to market expectations. The direction of the surprise depends entirely on how hawkish or dovish the bank is perceived to be. From the central bank's perspective, there will be a positive monetary policy surprise when

$$\tilde{\varepsilon}_t^m = i_t - \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P\left[i_t\right]\right] > 0$$

or equivalently when

$$\pi_t < \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P\left[\pi_t\right]\right] = \mathbb{E}_t^P\left[\psi^\pi\right] X_t$$

Under the myopic policy, this occurs if $\psi^\pi < \mathbb{E}_t^P[\psi^\pi]$. Under the optimal policy, overreaction raises the bar for being 'perceived as hawkish': if the central bank's ψ^π matches the private sector's perception-$\psi^\pi = \mathbb{E}_t^P[\psi^\pi]$, it would still like to appear more hawkish because of the intertemporal smoothing.

Proposition 3. *Suppose $\varepsilon_t > 0$. There exist a $\widehat{\psi} < \psi^\pi$ such that there is a positive pure monetary policy surprise ($\tilde{\varepsilon}_t^m > 0$) if $\mathbb{E}_t^P[\psi^\pi] < \widehat{\psi}$ and negative otherwise.*

Proof. See Appendix B ■

Proposition 3 formalizes this idea: a central bank perceived as dovish will deliver a hawkish surprise, while one perceived as hawkish will deliver a dovish surprise. Figure 6 illustrates this relationship.

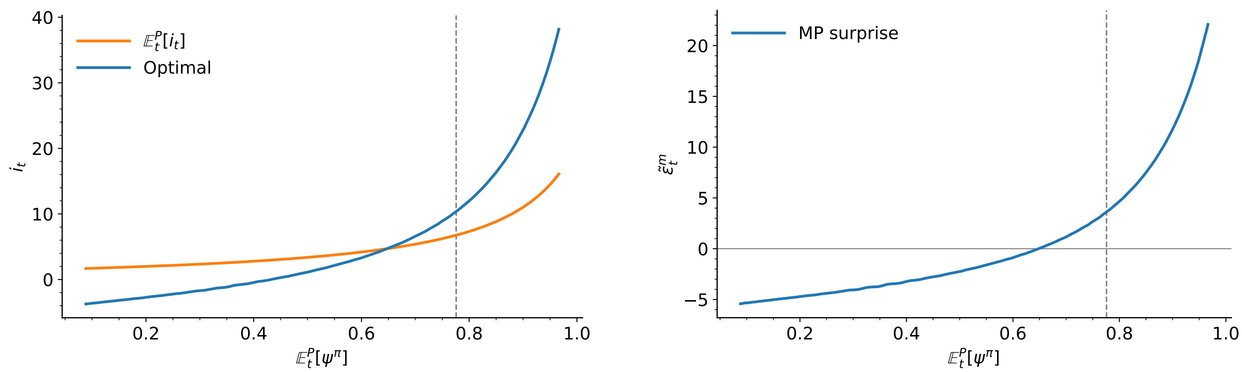

(a) Policy function for the interest rate i_t

(b) Monetary policy surprise

Figure 6: Monetary Policy Surprises as a function of reputation. Positive shock $\varepsilon_t > 0$

From the private sector's perspective, a monetary policy surprise, ε_t^m, is composed of the pure monetary policy surprise, $\tilde{\varepsilon}_t^m$, and the central bank's forecast error, ν_t. More formally,

$$\varepsilon_t = \tilde{\varepsilon}_t^m + \nu_t$$

This composition reinforces the idea that forecast errors not only have the same real effect as monetary policy shocks, but also have the same effect on the private sector's learning (Bauer and Swanson 2023)

How hawkish or dovish the bank is perceived depends on how its actions influence its reputation. Our next focus is on how the central bank manages its reputation.

Reputation Targeting From our previous discussion, we know the central bank will implement a hawkish monetary policy surprise whenever its reputation is below a certain threshold. But is it enough to improve reputation? The answer is no, because the private sector's beliefs embed a form of skepticism: if the central bank is thought to be hawkish ($\mathbb{E}_t^P[\psi^\pi] < \frac{1}{2}$), the private sector is not fully convinced it is so, and will tend to revise towards dovishness; and if it is thought to be dovish ($\mathbb{E}_t^P[\psi^\pi] > \frac{1}{2}$), the private sector will tend to revise towards hawkishness.[10]

Equation (15) makes this clear:

$$\overline{\psi}_{t+1} = \overline{\psi}_t + \omega_t \left[\left(\mathbb{E}_t^P[\psi^\pi] - \overline{\psi}_t \right) - \frac{\kappa}{\sigma} X_t^{-1} \varepsilon_t^m \right]$$

When the central bank is perceived as hawkish ($\mathbb{E}_t^P[\psi^\pi] > \frac{1}{2}$), $\overline{\psi}_t > \mathbb{E}_t^P[\psi^\pi]$ and skepticism works against reputation-building: independent of the monetary policy surprise, the private sector is skeptical about the central bank's reputation and will update towards a more dovish central bank unless proven otherwise. The central bank must deliver a monetary policy surprise large enough to offset this effect. The more hawkish the reputation, the larger the surprise required.

When the central bank is perceived as dovish ($\mathbb{E}_t^P[\psi^\pi] > \frac{1}{2}$), the opposite holds: skepticism works with the central bank's goals. It naturally pushes beliefs upwards, and a positive monetary policy surprise only accelerates this process.

The interaction between these two forces, monetary policy surprises and skepticism, determines a reputation target:

Proposition 4. *There exists a $\widehat{\psi}$ such that $\mathbb{E}_t^P[\psi^\pi] > \mathbb{E}_t^{CB}\left[\mathbb{E}_{t+1}^P[\psi^\pi]\right]$ when $\mathbb{E}_t^P[\psi^\pi] > \widehat{\psi}$ and $\mathbb{E}_t^P[\psi^\pi] < \mathbb{E}_t^{CB}\left[\mathbb{E}_{t+1}^P[\psi^\pi]\right]$ otherwise. Furthermore, when $\psi^\pi > \frac{1}{2}$ then $\widehat{\psi} < \psi^\pi$*

Proof. See Appendix B ■

Under reasonable calibrations, $\psi^\pi > \frac{1}{2}$, so the reputation target is more hawkish than the central bank's true preferences. The central bank improves its reputation when perceived as dovish, and 'uses it' when perceived as hawkish. A direct consequence of this result is that there are no multiple equilibria in the long run.

[10]This skepticism is a consequence of assuming the private sector beliefs have a truncated normal distribution.

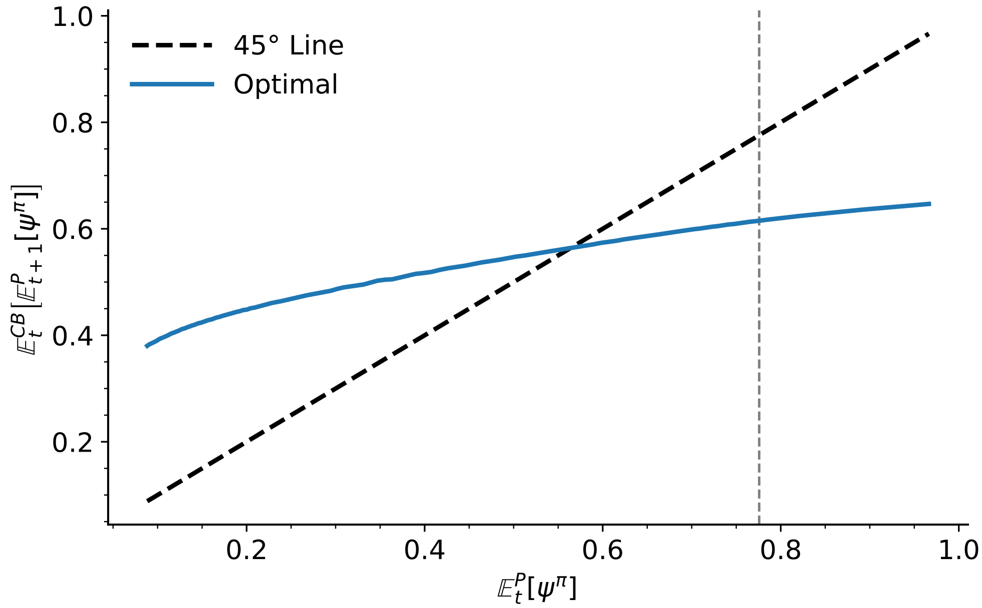

Figure 7: Reputation Dynamics.

Figure 7 shows that the central bank manages its reputation toward a hawkish target, building it when perceived as dovish, letting it decline when perceived as very hawkish. This management of reputation raises the question of when the central bank should build or spend its reputation. Should it take advantage of calm periods to strengthen its credibility, or wait for turbulent moments to 'show what it is made of'? Next, we answer this question.

Good Times and Bad Times The central bank's incentives to overreact depend on the state of the economy. When shocks are small and persistent ('good times'), both the anticipation of future recessions and the increased informativeness of policy push in the same direction: overreaction rises with the size of the shock. When shocks are large and short-lived ('bad times'), these forces point in opposite directions. Large shocks lower the informativeness of policy because high inflation is expected regardless of the central bank's policy. In this case, overreaction increases for small shocks, but decreases once they are sufficiently large.

Formally, the size of a shock affects overreaction through two channels:

1. **Recession smoothing:** larger current shocks predict larger future shocks ($\tilde{y}_{t+s} \downarrow$). Overreaction helps to smooth the future recession through reputation-building.

2. **Informativeness of Policy:** the sensitivity of expectations to policy, $\omega_t X_t^{-1}$, rises

with shock size when shocks are small, but falls when they are large.

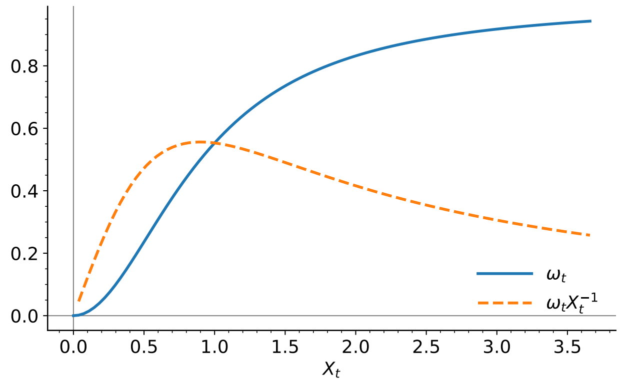

Figure 8: Effective Weight.

In good times, both channels point in the same direction. Overreaction is increasing in the size of the shock. In bad times, the informativeness channel eventually dominates, and overreaction decreases in the size of the shock.

Proposition 5. *Suppose ε_t is independent from $\{\mathbb{E}_t[\varepsilon_{t+1+s}]\}_{s\geqslant 0}$. The output gap overreaction $\left|\tilde{y}_t - \tilde{y}_t^M\right|$ is increasing in $|\varepsilon_t|$ for small values of $|\varepsilon_t|$ and decreases for large values of $|\varepsilon_t|$.*

Proof. See Appendix B ■

Taken together, these results describe a strategy for managing credibility:

1. The central bank overreacts relative to a myopic benchmark.
2. Overreaction should be larger when the reputation is neither hawkish nor dovish.
3. Monetary policy surprises should be hawkish (dovish) if the central bank is perceived as dovish (hawkish).
4. Target a more hawkish level of reputation.
5. Build reputation in good times, spend it in bad times.

While these results were developed in the context of cost-push shocks, the same logic extends to demand shocks. In this case, the private sector's perception about λ plays a more minor role, and what matters is how strongly the private sector believes the central bank is willing to move interest rates to stabilize the economy. As we will see next, all our results hold.

Markov Perfect Equilibrium So far, we assumed for tractability that the private sector held misspecified beliefs and assumed the central bank to be myopic. In reality, the private sector may be aware that the central bank internalizes the impact of its actions on their beliefs. To capture this, we consider a Markov Perfect Equilibrium, in which private beliefs are correctly specified. In Appendix C we prove that, as long as higher inflation signals a higher λ when there are positive shocks (and the opposite when shocks are negative), the central bank continues to overreact relative to the myopic benchmark. The magnitude of overreaction may be smaller than in the misspecified-beliefs case, since private agents anticipate the central bank's incentives, but the qualitative result remains. This result confirms that overreaction is a robust feature of optimal policy and sets the stage for our empirical analysis in the next section. Furthermore, in a three-period economy, all of our results hold.

Beyond Cost-Push Shocks So far, our discussion has focused on cost-push shocks, which are relatively rare compared to demand fluctuations. Do the same theoretical results carry over to this (more general) structure? To address this question, we assume the central bank cannot fully stabilize demand shocks. In practice, complete stabilization of demand shocks would require extreme volatility in interest rates, and central banks are usually gradualist.

To capture situations where the Divine Coincidence does not hold, consider the modified welfare function:

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s}^2 + \lambda\tilde{\pi}_{t+s}^2 + \varphi i_{t+s}^2\right)\right] \tag{19}$$

The third term, inspired by Woodford (2003b), penalizes large movements in the policy rate. When $\varphi \to 0$, we recover the benchmark where the central bank fully stabilizes demand shocks. A positive φ justifies gradualism and allows for sizeable deviations from the Divine Coincidence. This formulation captures the idea, emphasized by Bauer et al. (2024), of a central bank that remains responsive to evolving economic conditions and adjusts policy in a data-driven manner.

In Appendix C, we prove that if prices are fixed and the private sector is uncertain about the central bank's interest rate smoothing parameter, φ, then the optimal policy problem is isomorphic to the case where the private sector learns about λ. The optimal policy prescribes overreaction in the interest rate and underreaction in the output gap. Under functional form assumptions similar to the cost-push case, all our conclusions still hold. The logic of optimal policy is not tied to cost-push shocks, but also holds for demand-driven fluctuations. Relative to a myopic benchmark, the policy rate overreacts to manage reputation.

Reputation about the Long Run Following the traditional literature on reputation and monetary policy, suppose the private sector knows the central bank's relative weight on inflation, λ, but is uncertain about its inflation bias. Uncertainty about the inflation bias is analogous to a situation where the private sector learns about the central bank's inflation target. In Appendix C, we analyze this case and derive the implications for optimal policy. In an effort to stabilize future outcomes, the central bank overreacts to news shocks. This policy can be interpreted as akin to Average Inflation Targeting (see Powell 2020; Eggertsson and Kohn 2023).

5 Reputation in the Data

Our theory provides a new way of measuring reputation in the data, which allows us to test the main predictions of our theory.

Suppose that in each period t there is a continuum of independent forecasters. They have a common prior, μ_t, over the central bank's preference parameter λ and its assessment of demand, so they agree on $\mathbb{E}_t^P[\psi^\pi]$, $\mathbb{E}_t^P[\psi^y]$, and $\mathbb{E}_t^P[\eta_{t+k}]$. They differ only in their expectation about the sequence of cost-push shocks that hit the economy. Formally, forecaster i's projections for the output gap and inflation are

$$\mathbb{E}_t^i[y_{t+k}] = -\mathbb{E}_t^P[\psi^y]\,\mathbb{E}_t^i[X_{t+k}] + \mathbb{E}_t^P[\eta_{t+k}] \tag{20}$$

$$\mathbb{E}_t^i[\pi_{t+k}] = \mathbb{E}_t^P[\psi^\pi]\,\mathbb{E}_t^i[X_{t+k}] + \kappa\mathbb{E}_t^P[\eta_{t+k}] \tag{21}$$

where

$$\mathbb{E}_t^i[X_{t+k}] = \sum_{s-0}^{\infty}\left(\beta\mathbb{E}_t^P[\psi^\pi]\right)^s\mathbb{E}_t^i[\varepsilon_{t+k+s}]$$

Under this assumption, all the variation in forecasts comes from differences in forecasts of the shocks, not from disagreement about the central bank's preferences.

Suppose we have access to survey data on expectations at period t. Consider the regression

$$\mathbb{E}_t^i[y_{t+k}] = \gamma_{1t} + \gamma_{2t}\mathbb{E}_t^i[\pi_{t+k}] + u_{t,k}^i \tag{22}$$

Proposition 6. *If forecasters form their expectations according to* (20) *and* (21) *then*

$$\gamma_{2t} = -\frac{\mathbb{E}_t^P\left[\psi^y\right]}{\mathbb{E}_t^P\left[\psi^\pi\right]} = -\kappa^{-1}\left(\frac{1}{\mathbb{E}_t^P\left[\psi^\pi\right]} - 1\right) \tag{23}$$

Moreover, a more negative value of γ_{2t} *corresponds to an improvement in the central bank's reputation.*

Proof. See Appendix D ■

Proposition 6 summarizes our main empirical result: we can recover the central bank's reputation from a simple regression of output-gap forecasts on inflation forecasts. The intuition is straightforward. Consider a myopic central bank. Its first-order condition is

$$y_t = -\kappa\lambda\pi_t$$

A larger λ implies a steeper slope in the regression of the output gap on inflation. Proposition 6 extends this idea to the case where the private sector is uncertain about λ.

It is important to stress why the survey data on forecasts is essential. Cross-sectional variation in survey forecasts at a given t provides variation in expected outcomes while holding reputation fixed, which allows us to retrieve reputation from $\gamma_{2,t}$. In contrast, time variation in realized outcomes is not sufficient to identify reputation, since it fluctuates.

An improvement in reputation, captured by a decline in $\mathbb{E}_t^P\left[\psi^\pi\right]$, corresponds to a more negative value of $\gamma_{2,t}$. As shown in Lemma 2, this can arise when the private sector shifts its beliefs towards a more hawkish central bank. Alternatively, as shown in Lemma 3, it can also occur if beliefs become more tightly concentrated around their mean.

POSITIVE ESTIMATE OF $\gamma_{2,t}$ Although our framework predicts a negative correlation between inflation and the output gap, two mechanisms can generate a positive estimate of $\gamma_{2,t}$: disagreement about the central bank's forecast errors, and perceived gradualism in monetary policy. We discuss each, and then show that the comparative statics with respect to reputation still hold.

First, if there is disagreement about the central bank's forecast error η_t, then the estimands of $\gamma_{2,t}$ become a weighted average:

$$\gamma_{2,t} = \omega_t\left(-\frac{\mathbb{E}_t^P\left[\psi^y\right]}{\mathbb{E}_t^P\left[\psi^\pi\right]}\right) + (1-\omega_t)\left(\frac{1}{\kappa}\right)$$

where

$$\omega_t = \frac{\mathbb{E}_t^P\left[\psi^\pi\right]^2 Var\left(\mathbb{E}_t^i\left[X_{t+k}\right]\right)}{\mathbb{E}_t^P\left[\psi^\pi\right]^2 Var\left(\mathbb{E}_t^i\left[X_{t+k}\right]\right)+\kappa^2 Var\left(\mathbb{E}_t^i\left[\eta_{t+k}\right]\right)}$$

$\gamma_{2,t}$ can turn positive when disagreement about the cost-push shocks is small relative to the disagreement about the demand shocks not fully stabilized by the central bank. More generally, $\gamma_{2,t}$ also depends on the cross-sectional variance of the forecasts of these shocks. When disagreement about the forecast error falls, $\gamma_{2,t}$ becomes more negative. It is essential to notice that what matters is not disagreement about demand shocks per se, but rather disagreement about demand shocks that are not fully stabilized by policy. Outside major crises such as the Global Financial Crisis or the COVID-19 pandemic, the period covered was characterized by stability and low volatility. We therefore view this channel as unlikely to be the primary driver of the time variation in reputation. Importantly, as long as disagreement about unstabilized demand shocks is not too large relative to disagreement about cost-push shocks, the comparative static with respect to reputation still holds: improvements in reputation are reflected in a more negative $\gamma_{2,t}$.

A second reason why $\gamma_{2,t}$ can have positive values is that the central bank does not fully stabilize demand shocks. To allow for situations where the Divine Coincidence does not hold, suppose the central bank follows the triple mandate in (19)

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s}^2+\lambda\tilde{\pi}_{t+s}^2+\varphi i_{t+s}^2\right)\right]$$

Assume the private sector does not know λ or φ, but continues to assume the central bank is myopic. For simplicity, assume there are only demand shocks.[11] In this environment, the correlation between inflation and the output gap can be either positive or negative.

Lemma 5. *Suppose the central bank's preferences are the dual mandate from* (19) *and there are only demand shocks. Define* $\widehat{\psi} := \frac{\varphi}{1+\lambda\kappa^2+\sigma^2\varphi}$. *Then, the estimand of* $\gamma_{2,t}$ *in equation* (22) *can be positive or negative, and increasing in* $\mathbb{E}_t^P\left[\widehat{\psi}\right]$

Proof. See Appendix D ■

Lemma 5 extends the logic of our results to demand shocks. As in the main model, improvements in reputation are also captured by a smaller value of $\gamma_{2,t}$. Unlike the cost-

[11]In this setting, the assumption $\mathbb{E}_t^i\left[\eta_{t+k}\right] = \mathbb{E}_t^P\left[\eta_{t+k}\right]$ can also be read as no disagreement about future monetary policy shocks, as in Bauer et al. (2024).

push shock case, the sign of the correlation can flip. Contemporaneous demand shocks push inflation and the output gap in the same direction, yielding a positive correlation. In addition to this effect, anticipated demand shocks also act like cost-push shocks, raising inflation expectations and generating a negative comovement. When the central bank is perceived as dovish, the first effect dominates; when it is perceived as hawkish, the second one dominates. As in the cost-push shock case, a stronger reputation dampens the impact of news shocks.

This extension also modifies the reputation measure by incorporating φ, which captures gradualism. When the private sector believes the central bank is less willing to adjust rates (higher perceived φ), $\mathbb{E}_t^P\left[\widehat{\psi}\right]$ increases. For the perceived inflation-output trade-off, the logic remains the same: a higher perceived weight on inflation stabilization implies a more negative slope. If reputation improves, whether through a higher perceived weight on inflation or lower perceived gradualism, the estimand of $\gamma_{2,t}$ decreases.[12]

In sum, there are two reasons why the estimand of $\gamma_{2,t}$ can be positive: disagreement about demand shocks, or gradualism in monetary policy. In the first case, disagreement shifts the level of $\gamma_{2,t}$. Still, we do not believe changes in disagreement are a significant source of time variation in reputation during the Great Moderation. In the second case, the measure of reputation is modified to incorporate perceived gradualism. A higher perceived φ lowers reputation, while a lower perceived φ improves reputation. The comparative statics, however, remain unchanged: reputation improves when the central bank is perceived as either more hawkish or less gradualist, and in all such cases the estimand of $\gamma_{2,t}$ decreases.

5.1 Data and Estimation

Our primary data source is the Blue Chip Financial Forecasts (BCFF) survey, which provides individual forecasts of interest rates and macroeconomic variables such as GDP growth and CPI inflation for the U.S. economy. This survey has been used in the literature to measure the private sector's expectations, and more recently, their perceptions about monetary policy (Bauer et al. (2024)). Each month, forecasters report projections for their current quarter and four to five quarters ahead. We use the sample from January 1992 to June 2023.

The forecasts of output growth and CPI inflation are reported in quarter-over-quarter

[12]In Appendix C we show that this measure of reputation is negatively correlated with the perceived Taylor rule coefficient on the output gap. Consistent with this, in Appendix D we document the same negative correlation in the data using the estimates of Bauer et al. (2024).

annualized terms. Following Bauer et al. (2024), we transform them into year-over-year inflation and output gap forecasts. For inflation, we combine realized CPI with the survey responses to build forecasts of year-over-year inflation. For output, we construct GDP forecasts by cumulating quarterly growth projections, using the contemporaneous Archival Federal Reserve Economic Data (ALFRED) vintage for the level of real GDP.[13] Potential output is taken from Congressional Budget Office (CBO) projections, also retrieved in real time from ALFRED. Finally, we calculate the output gap as

$$\mathbb{E}_t^i\left[y_{t+k}\right] = 100 \times \frac{\mathbb{E}_t^i\left[Y_{t+k}\right] - \mathbb{E}_t\left[Y_{t+k}^n\right]}{\mathbb{E}_t\left[Y_{t+k}^n\right]}$$

where Y_{t+k} is real output, and Y_{t+k}^npotential output at horizon $t+k$.

We estimate (22) using monthly data from January 1992 to June 2023. Figure 9a plots the time series of $\gamma_{2,t}$. Figure 9b illustrates the underlying cross-sectional variation, comparing two moments in time: January 2010, when reputation was relatively high ($\gamma_{2,t} = -0.254$), and January 2015, when it was relatively low ($\gamma_{2,t} = 0.386$). When reputation is high, the comovement of forecasts of inflation and the output gap becomes more negative.

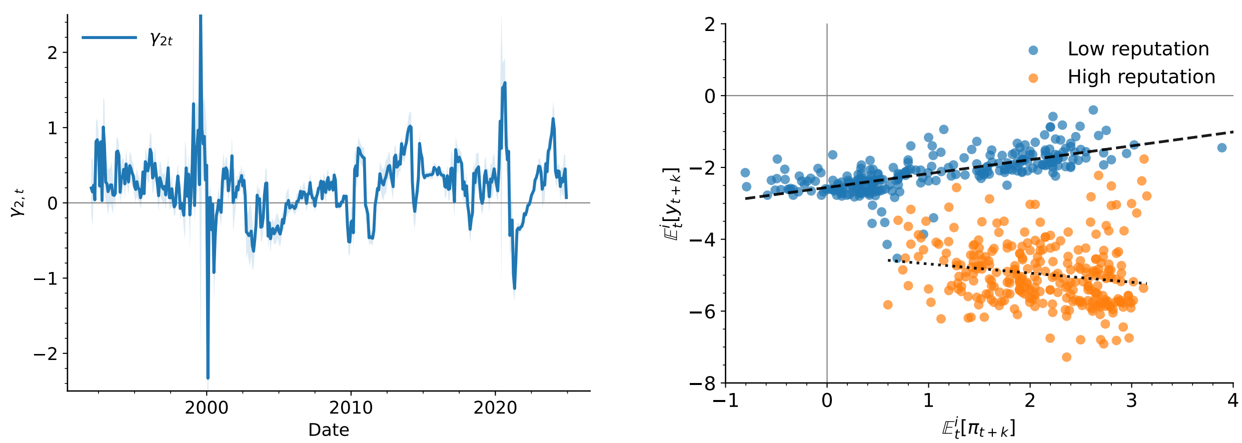

(a) Monthly estimates for $\gamma_{2,t}$ **(b)** Reputation: High vs Low

Figure 9: Dynamics of $\gamma_{2,t}$

An Augmented Dickey-Fuller test on γ_{2t} rejects the null of a unit root at the 1% level. The absence of a unit root is consistent with our learning model. Reputation fluctuates around a long-run target, deteriorating when it is above (more hawkish) and improving when it is

[13]ALFRED provides the vintages available to forecasters at each survey date. If the exact date is missing, we use the closest vintage. We assume the survey took place on the first of each month.

below (more dovish). This property holds whether the central bank is myopic or follows the optimal policy.

5.2 Three Facts about Reputation

With our measure of reputation in hand, we now study its empirical properties. Specifically, we document three facts that summarize how reputation responds to shocks, each of which is consistent with our model. We cut our sample in December 2019 to exclude the impact of the COVID-19 Pandemic.

Fact #1: Reputation Anchors Inflation Expectations Our first finding shows that our measure of reputation behaves as the theory predicts: a more hawkish reputation reduces the pass-through of shocks to inflation expectations. Better reputation anchors inflation expectations.

To show this, we follow Ramey and Zubairy (2018) and Kolesár and Plagborg-Møller (2024) and split the sample into high and low reputation regimes, defined relative to the mean of γ_{2t}, which we denote by $\bar{\gamma}_2$. We estimate the effect of the oil price news shocks from Känzig (2021) on inflation expectations under each regime. Oil price news shocks are a natural test for our theory: they are cost-push shocks that raise inflation, and because they are anticipated, they affect inflation expectations. We estimate

$$\mathbb{E}_t^i\left[\pi_{t+k}\right] = I_{\gamma_{2t-1}>\bar{\gamma}_2}\left(\beta_{1,L}^h + \beta_{2,L}^h \varepsilon_t\right) + \left(1 - I_{\gamma_{2t-1}>\bar{\gamma}_2}\right)\left(\beta_{1,H}^h + \beta_{2,H}^h \varepsilon_t\right) + u_{i,t,k}$$

Figure 10 shows the result. When reputation is dovish, the pass-through from oil price news shocks to inflation expectations is large; when reputation is hawkish, it is smaller. The difference vanishes at longer horizons, reflecting that these shocks are short-lived and reputation matters the most for the short run.

Figure 10 shows that the pass-through of cost-push shocks on reputation depends heavily on the central bank's reputation. The pass-through is larger when the reputation is dovish, which supports the idea that a hawkish reputation anchors inflation expectations. The difference vanishes with longer horizons, which has two potential explanations. First, oil price news shocks are short-lived, so their impact vanishes with long horizons. Second, the central bank's reputation is more important in the short run than in the long run.

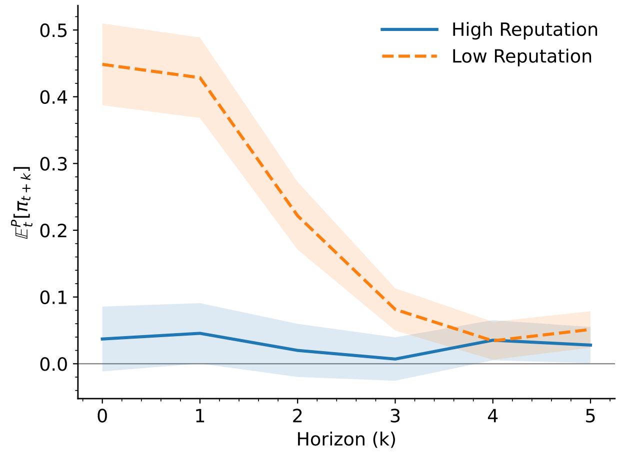

Figure 10: Impulse response of inflation expectations to an oil price news shock for high and low reputation.
Shaded areas denote 95% confidence intervals.

Fact #2: Reputation Responds to Monetary Policy Shocks Our second fact is that monetary policy surprises move reputation. Using high-frequency shocks from Gertler and Karadi (2015), we estimate

$$\gamma_{2,t+h} = \alpha_h + \beta_h \varepsilon_t^m + \delta_{1,h}\gamma_{1,t-1} + \delta_{2,h}\gamma_{2,t-1} + u_{i,t,h} \qquad i = 1,2 \tag{24}$$

where ε_t^m is the FF4 monetary policy surprise from Gertler and Karadi (2015). We find that an unexpected tightening makes the central bank appear more hawkish. Figure 11 shows that γ_{2t} declines, consistent with an improvement in reputation.

To understand how a positive monetary policy surprise can generate a positive average effect, define $mps_t = i_t - \mathbb{E}_t^P[i_t]$ as the monetary policy surprise. Rewriting (15), we have

$$\mathbb{E}\left[\overline{\psi}_{t+1}\middle| mps_t = \varepsilon^m\right] - \mathbb{E}\left[\overline{\psi}_{t+1}\middle| mps_t = 0\right] = -\mathbb{E}\left[\omega_t X_t^{-1}\right]\frac{\kappa}{\sigma}\varepsilon^m$$

Our findings are consistent with $\mathbb{E}\left[\omega_t X_t^{-1}\right] > 0$. One possibility is that, on average, shocks are positive. Another possibility is that, even if on average shocks are zero ($\mathbb{E}[X_t] = 0$), the distribution is not symmetric. In particular, if there are more positive cost-push shocks than negative ones, then the term will be positive (up to fifth order). Using (14) and $\mathbb{E}[X_t] = 0$, we have

$$\mathbb{E}\left[\omega_t X_t^{-1}\right] = -\left(\frac{\kappa^{-2}\tau_\eta}{\tau_t}\right)^3 \mathbb{E}\left[X_t^3\right] + \mathbb{E}\left[\mathcal{O}\left(X_t^5\right)\right] > 0$$

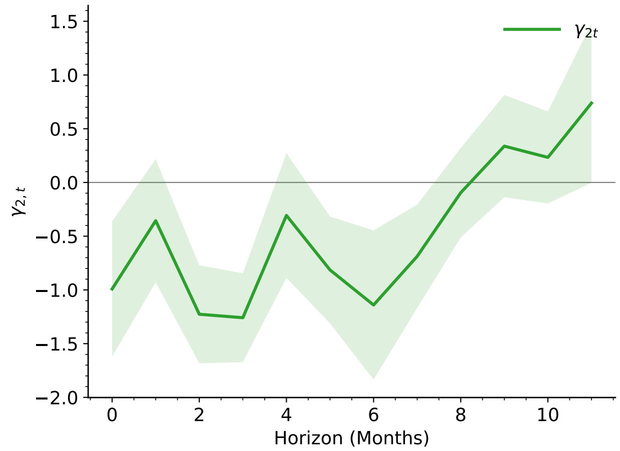

Figure 11: Impulse response to a monetary tightening.
Shaded areas denote 68% confidence intervals.

We show in Appendix D that this finding is robust to controlling for the "Fed Information Effect" (Nakamura and Steinsson 2018, Bauer and Swanson 2022, Bauer and Swanson 2023). In particular, controlling for information in between FOMC announcements does not change our conclusions.

Finally, our model is consistent with the evidence that monetary policy shocks are partly predictable (Ramey 2016, Bauer and Swanson 2023). When reputation is below target, an unexpected tightening is optimal when the economy is booming or there are positive cost-push shocks. As a result, policy surprises may be *ex post* predictable from information available between FOMC meetings.[14] The model, therefore, not only explains why monetary policy shocks may appear predictable, but also shows that such predictability can be itself optimal.

Fact #3: Reputation Does Not Respond to Non-Monetary Shocks Does reputation change with shocks other than surprises from the central bank? To answer this question, we estimate (24) using oil price news shocks from Känzig (2021). This shock embodies a typical cost-push shock.

Figure 12a plots the impulse response to an oil price news shock. The impact on reputation is negligible. To further evaluate the effect of cost-push shocks on reputation, we estimate the medium-scale VAR for the U.S. economy of Angeletos et al. (2020) and recover

[14]Even at the long-run reputation target, private-sector skepticism may induce the central bank to deliver a hawkish monetary policy surprise.

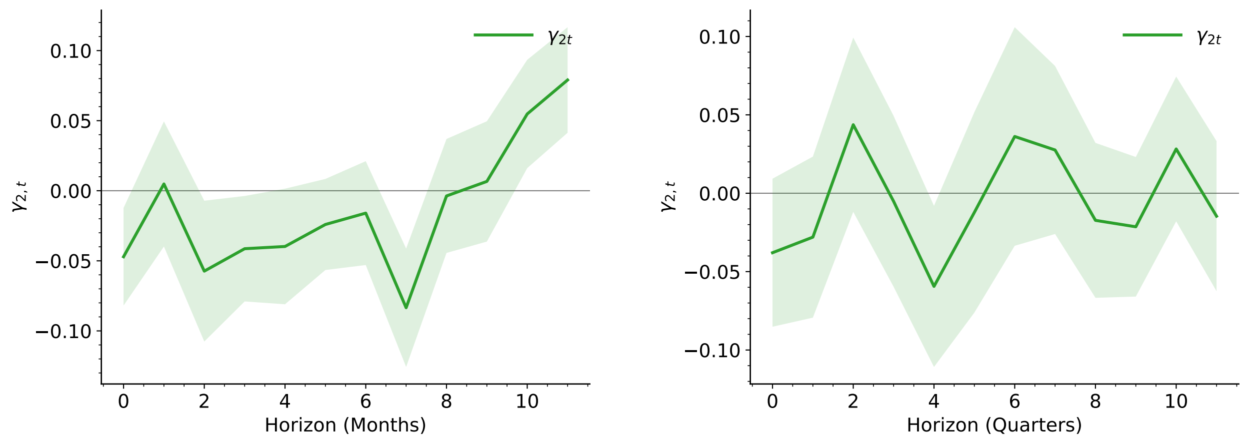

(a) Oil price news shock Känzig (2021) **(b)** Wold innovation in inflation

Figure 12: Impulse responses to cost-push shocks.
Shaded areas denote 68% confidence intervals.

the Wold innovations for inflation. A Wold innovation for inflation represents a linear combination of structural shocks that increase inflation unexpectedly. These innovations embody a more general combination of cost-push shocks.

Figure 12b plots the impulse response of reputation to a one standard deviation Wold innovation. Since the VAR was estimated at the quarterly frequency, we use the quarterly average of reputation in each period. As with the case of oil price news shocks, the impact is not significant. This evidence is suggestive that the central bank is at its long-run target for reputation. To see this, let $\boldsymbol{\varepsilon}_t := \{\mathbb{E}_t[\varepsilon_{t+s}]\}_{s\geqslant 0}$. From (14), the impact of a cost-push shock on the central bank's reputation is

$$\mathbb{E}\left[\overline{\psi}_{t+1}\middle|\boldsymbol{\varepsilon}_t = \boldsymbol{\varepsilon}\right] - \mathbb{E}\left[\overline{\psi}_{t+1}\middle|\boldsymbol{\varepsilon}_t = \mathbf{0}\right] = \mathbb{E}\left[\omega_t X_t^{-1}\left(\pi_t - \overline{\psi}_t X_t\right)\middle|\boldsymbol{\varepsilon}_t = \boldsymbol{\varepsilon}\right] \tag{25}$$

The term in brackets is equal to zero whenever the reputation is constant. Then, this evidence suggests that on average the central bank's reputation is at its long-run target: it increases when it is below and decreases when it is above.[15]

[15] This by itself cannot distinguish whether the central bank is acting myopically or following the optimal policy.

5.3 Extensions

The empirical analysis has so far focused on our baseline model. In this subsection, we look at a number of extensions to check the robustness of the results. We consider central bank commitment, inflation inertia, Taylor rules, belief disagreements, and structural breaks. In each case, the evidence points in the same direction: the data remain consistent with the mechanisms highlighted in our model.

COMMITMENT Suppose that the private sector believes the central bank operates under commitment, designing optimal policy from the timeless perspective. The first-order condition is, then

$$\tilde{y}_t - y_{t-1} + \kappa\lambda\tilde{\pi}_t = 0 \tag{26}$$

Lemma 6. *Under the timeless perspective, the k-periods ahead forecast is*

$$\begin{aligned}
\mathbb{E}_t^i\left[y_{t+k}\right] =& \mathbb{E}_t^i\left[\varphi^y y_{t+k-1}\right] - \alpha\mathbb{E}_t^P\left[\psi^y\right]\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right] \\
\mathbb{E}_t^i\left[\pi_{t+k}\right] =& \mathbb{E}_t^i\left[\varphi^\pi y_{t+k-1}\right] + \alpha\mathbb{E}_t^P\left[\psi^\pi\right]\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]
\end{aligned}$$

where $\alpha < 1$, and φ^y and φ^π are positive constants.

Proof. See Appendix D ■

Lemma 6 highlights two differences relative to the myopic benchmark. First, shocks have a smaller direct impact ($\alpha < 1$): under commitment, the central bank smooths policy to reduce volatility. Second, the lagged output gap becomes a new state variable, reflecting the central bank honoring its past promises.

Empirically, this implies we can recover reputation under commitment by estimating

$$\mathbb{E}_t^i\left[y_{t+k}\right] = \gamma_{1,t} + \gamma_{2,t}\mathbb{E}_t^i\left[\pi_{t+k}\right] + \gamma_{3,t}\mathbb{E}_t^i\left[y_{t+k-i}\right] + u_{i,t,k}$$

By the Frisch-Waugh-Lovell Theorem, $\gamma_{2,t}$ coincides with the estimand in (22), so Proposition 6 continues to hold. In practice, the correlation coefficient between the commitment-based and baseline estimates is 0.84. Figure 13a shows the monthly series under commitment, confirming the robustness of our measure.

Inertia Suppose that past inflation feeds into current inflation, so that the NKPC exhibits inertia. In this case, the New Keynesian Phillips Curve (2) becomes

$$\tilde{\pi}_t = \kappa \tilde{y}_t + \delta \pi_{t-1} + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1}\right] + \varepsilon_t \quad \delta + \beta < 1.$$ [16]

Lemma 7. *Under inertia, the k-periods ahead forecast is*

$$\begin{aligned} \mathbb{E}_t^i \left[y_{t+k}\right] &= - \mathbb{E}_t^i \left[\varphi^y \pi_{t+k-1}\right] - \alpha \mathbb{E}_t^P \left[\psi^y\right] \sum_{s=0}^{\infty} \left(\alpha \beta \mathbb{E}_t^P \left[\psi^\pi\right]\right)^s \mathbb{E}_t^i \left[\varepsilon_{t+k+s}\right] \\ \mathbb{E}_t^i \left[\pi_{t+k}\right] &= \mathbb{E}_t^i \left[\varphi^\pi \pi_{t+k-1}\right] + \alpha \mathbb{E}_t^P \left[\psi^\pi\right] \sum_{s=0}^{\infty} \left(\alpha \beta \mathbb{E}_t^P \left[\psi^\pi\right]\right)^s \mathbb{E}_t^i \left[\varepsilon_{t+k+s}\right] \end{aligned}$$

where $\alpha > 1$, and φ^y and φ^π are positive constants.

Proof. See Appendix D ■

Proposition Lemma 7 highlights two differences relative to the myopic benchmark. First, shocks have a larger direct impact ($\alpha > 1$): inertia amplifies the impact of shocks on inflation. Second, the lagged inflation becomes a new state variable, reflecting that past inflation acts as a cost-push shock in the presence of inertia.

By the Frisch-Waugh-Lovell theorem, the coefficient γ_{2t} coincides with the estimand in (22), so Proposition 6 continues to hold. In practice, the correlation coefficient between the inertia and baseline estimates is 0.82. Figure 13b plots the monthly series under inertia.

Zero Lower Bound A further concern is that the private sector's model abstracts from the Zero Lower Bound (ZLB). When the ZLB binds, Lemma 1 no longer applies, and forecasts are not given by (9) and (10). To address this issue, we re-estimate our baseline specification (22) after excluding forecasts that assume the policy rate will remain at the ZLB—that is, those for which $\mathbb{E}_t^i \left[i_{t+k}\right] = 0$— during the 2009-2015 period, when the constraint was binding.[17] Figure 14 compares the two series, which display a correlation coefficient of 0.78

[16] In conventional macro models with inertia, the sum of the coefficients of the NKPC is weakly smaller than one. See Werning (2022).

[17] In practice, we exclude observations with$\mathbb{E}_t^i \left[i_{t+k}\right] < 0.20$. Since policy rates typically move in increments of 0.25, values below 0.20 are best interpreted as ZLB.

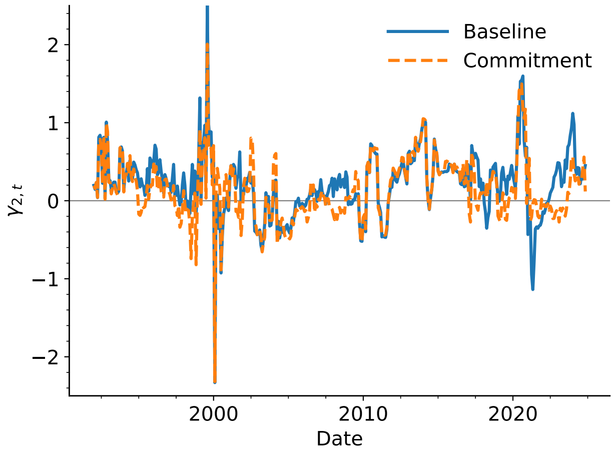

(a) Commitment

(b) Inertia

Figure 13: Monthly estimates for $\gamma_{2,t}$ under Commitment and Inertia

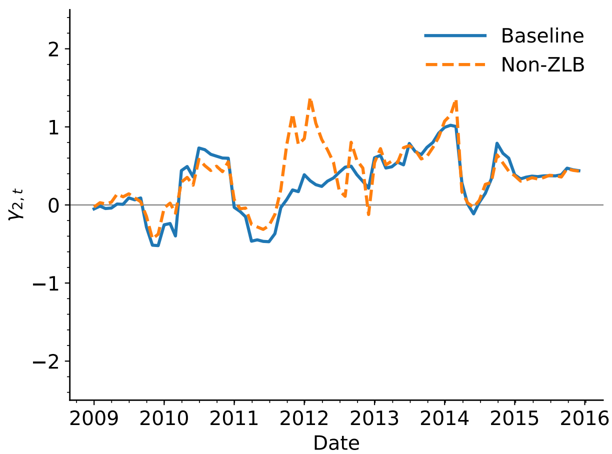

Figure 14: Monthly estimates for $\gamma_{2,t}$ at the ZLB

Identification under a Taylor Rule Recent work estimates the private sector's perception of monetary policy by estimating the Taylor rule from forecasts (Bauer et al. 2024, Bocola et al. 2024). While there is skepticism about whether the central bank actually follows such a rule (see, e.g. Svensson 2003 or Nakamura et al. 2025), we ask: is there a one-to-one relationship between our measure of reputation and the perceived Taylor rule coefficient?

Lemma 8. *Suppose we estimate the following Taylor rule*

$$\mathbb{E}_t^i \left[i_{t+k}\right] = \alpha_{1t} + \alpha_{2t}\mathbb{E}_t^i \left[\pi_{t+k}\right] + u_{i,t+k}$$

If $\beta > \frac{\kappa^{-1}\sigma}{1+\kappa^{-1}\sigma}$ *then* $\alpha_{2,t}$ *is decreasing in* $\mathbb{E}_t^P \left[\psi^\pi\right]$ *when reputation is low, and increasing in*

$\mathbb{E}_t^P [\psi^\pi]$ *when reputation is high.*

Proof. See Appendix D ■

Proposition 8 highlights a tension between two forces. On one hand, when the central bank is perceived as more hawkish, the private sector expects it to react more strongly to shocks, raising α_{2t}. On the other hand, a stronger reputation means inflation expectations are better anchored, so the central bank can react less to the same shocks, lowering α_{2t}. When reputation is low (high $\mathbb{E}_t^P [\psi^\pi]$), the first effect dominates; when reputation is high, the second effect dominates.

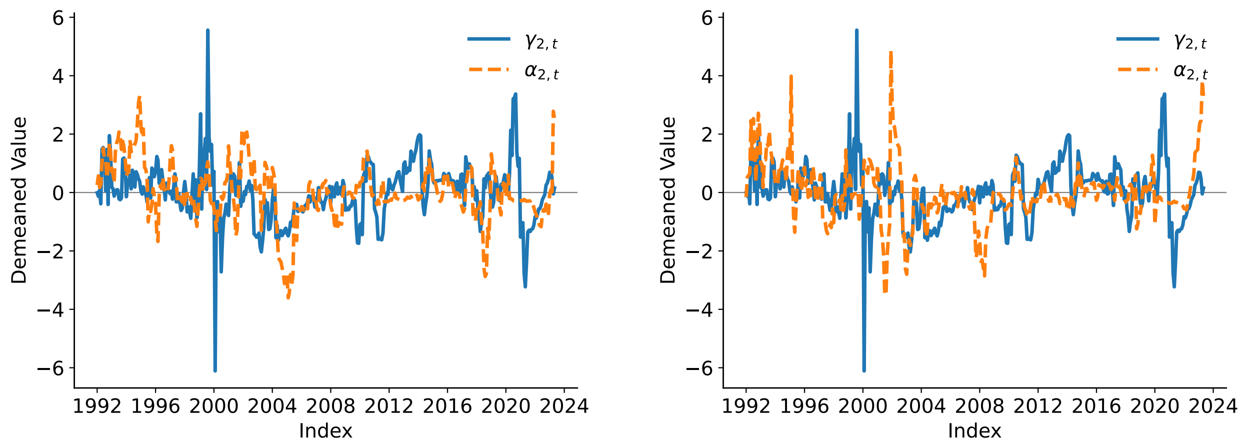

(a) Baseline **(b)** Inertial

Figure 15: Perceived Taylor-rule inflation coefficient for inflation ($\alpha_{2,t}$) from Bauer et al. (2024) and our reputation measure ($\gamma_{2,t}$).

Figure 15 plots the perceived inflation coefficient of the Taylor rule, $\alpha_{2,t}$ estimated by Bauer et al. (2024) alongside our model-consistent measure of reputation, $\gamma_{2,t}$. We display both the estimates for the baseline and inertial Taylor rules. For comparison, we standardize both measures. The correlation coefficient between the two series is 0.27 in the baseline rule and 0.17 in the inertial rule. This positive correlation is at odds with the conventional interpretation: if a larger inflation coefficient is taken to signal greater hawkishness, then the series should move in opposite directions. Through the lens of our model, however, an increase in the perceived inflation coefficient coincides with a deterioration in reputation, indicating that the common practice of equating the two is misleading.[18] Increases in

[18]In Appendix C, we show that an increase in the perceived Taylor-rule coefficient on the output gap corresponds to a higher perceived willingness of the central bank to adjust interest rates to stabilize the

the perceived Taylor rule coefficient on inflation do not necessarily reflect a more hawkish reputation through the lens of our model.

Disagreements about the Central Bank's Reputation So far, we assumed the private sector agreed on the central bank's reputation. Suppose instead that beliefs over λ, μ_t^i, differ across agents, while still assuming that forecasts of shocks are independent of beliefs about the central bank's preferences.

Lemma 9. *When beliefs are heterogeneous and* $Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k}\right]\right) = \phi^k Var\left(\mathbb{E}_t^i\left[\varepsilon_t\right]\right)$ *with* $\phi \leqslant 1$ *, then*

$$\gamma_{2,t} = -\overline{\mathbb{E}}\left[\omega^i \frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right] \qquad \omega^i = \frac{\sum_{s=0}^{\infty}\left(\beta^s \mathbb{E}_t^i\left[\psi^\pi\right]^{s+1}\right)^2 \phi^s}{\overline{\mathbb{E}}\left[\sum_{s=0}^{\infty}\left(\beta^s \mathbb{E}_t^i\left[\psi^\pi\right]^{s+1}\right)^2 \phi^s\right]}$$

where $\overline{\mathbb{E}}\left[\cdot\right]$ *denotes the cross-sectional mean. Then,* $\gamma_{2,t}$ *is increasing in* $Var\left(-\frac{\mathbb{E}_t^i[\psi^y]}{\mathbb{E}_t^i[\psi^\pi]}\right)$.

Proof. See Appendix D ∎

Lemma 9 shows that with heterogeneous beliefs, the estimand γ_{2t} is a weighted average of the individual estimands $\gamma_{2t}^i = -\frac{\mathbb{E}_t^i[\psi^y]}{\mathbb{E}_t^i[\psi^\pi]}$. Formally,

$$\gamma_{2,t} = -\overline{\mathbb{E}}\left[\frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right] + Cov\left(\omega^i, -\frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right)$$

The second term is positive, so belief heterogeneity biases γ_{2t} upward relative to the cross-sectional mean. Disagreement creates an additional channel through which γ_{2t} can change: a reduction in the disagreement about the central bank's preferences lowers γ_{2t}. If we interpret cross-sectional disagreement as different draws from the same prior beliefs, μ_t, then this result is consistent with Lemma 3.

Model Heterogeneity We initially assumed that the private sector shared a common model when forming their forecasts. Suppose instead that each forecaster i relies on a different model. In particular, each forecaster uses a different value of the slope of the NKPC, κ_i. While beliefs about the central bank's preference λ remain homogeneous, the derived beliefs over $\psi^\pi = \frac{1}{1+\kappa^2\lambda}$ are heterogeneous.

economy. This interpretation is in line with Bauer et al. (2024), which interprets an increase in the coefficient as an increase in perceived responsiveness to economic conditions.

Under this assumption, Lemma 9 holds: the estimand of $\gamma_{2,t}$ is a weighted average of the individual estimands, $\gamma^i_{2,t} = -\frac{\mathbb{E}^i_t[\psi^y]}{\mathbb{E}^i_t[\psi^\pi]}$, placing more weight on forecaster with a flatter NKPC. Since beliefs over λ are still homogeneous, the comparative statics from Lemma 2 and Lemma 3 continue to hold. A shift in beliefs toward a more hawkish central bank makes each $\gamma^i_{2,t}$ more negative, and therefore their weighted average, $\gamma_{2,t}$, shifts in the same direction.

Structural Breaks All our comparative statics assume no structural breaks. In particular, there are no changes in the slope of the NKPC, κ. How does the estimand γ_{2t} depend on κ? Taking the derivative

$$\frac{\partial \gamma_{2,t}}{\partial \kappa} = \kappa^{-2}\left(\frac{1}{\mathbb{E}^P_t\left[\psi^\pi\right]} - 1\right) + \kappa^{-1}\frac{\mathbb{E}^P_t\left[\frac{\partial \psi^\pi}{\partial \kappa}\right]}{\mathbb{E}^P_t\left[\psi^\pi\right]^2} = \kappa^{-2}\left(\frac{1}{\mathbb{E}^P_t\left[\psi^\pi\right]} - 1\right) - 2\kappa^{-1}\frac{\mathbb{E}^P_t\left[\psi^\pi\psi^y\right]}{\mathbb{E}^P_t\left[\psi^\pi\right]^2}$$
$$= \kappa^{-2}\frac{1}{\mathbb{E}^P_t\left[\psi^\pi\right]}\left(1 - \mathbb{E}^P_t\left[\psi^\pi\right] - 2\frac{\mathbb{E}^P_t\left[\psi^\pi\left(1-\psi^\pi\right)\right]}{\mathbb{E}^P_t\left[\psi^\pi\right]}\right)$$

This expression is not unambiguously positive or negative. Under certainty, we know $\gamma_{2t} = -\kappa\lambda$, so higher κ makes γ_{2t} more negative. Under uncertainty, however, the mapping is ambiguous. In practice, because our analysis excludes the COVID-19 pandemic, we do not view structural breaks as a significant driver of variation in γ_{2t}. Thus, while structural breaks could matter in principle, they are not an important concern in our sample.

To summarize, the evidence confirms that the reputation channel shapes expectations. With this channel established in the data, we now ask how the optimal policy can be implemented in practice. The following section turns to the quantitative analysis, where we establish that a simple institutional design can replicate the optimal policy.

6 New Problems, Same Old Solutions?

Implementing the optimal policy is not straightforward. It requires the central bank to track reputation in real time and anticipate how its actions influence the private sector's beliefs, which raises a natural question: can a simple rule approximate the optimal policy, and can it do so with greater robustness under misspecification?

In his seminal contribution, Rogoff (1985) proposed appointing a more hawkish central banker to offset inflation bias. In this section, we ask whether a similar delegation can work in our setting. Here, there is no inflation bias, but the optimal policy depends on reputation and how the private sector learns. The natural analog to Rogoff's proposal is to appoint a central banker who is both myopic and more hawkish. Does this rule approximate the optimal policy? Is it more robust if the central bank has the incorrect model of private sector learning? The goal of this exercise is not to provide a full quantitative model, but to clarify when the hawkish-myopic approximation succeeds and when it fails.

We benchmark against the optimal policy of a utilitarian planner with $\lambda = 10$, a standard value in the literature (see Galí 2003). The model is calibrated at a quarterly frequency. The slope of the New Keynesian Phillips Curve is $\kappa = 0.17$; the discount factor is $\beta = 0.99$; and the intertemporal elasticity of substitution is $\sigma = 1$. The cost-push shock follows an AR(1) process with persistence $\rho = 0.90$ and standard deviation $\sigma_u = 0.2$. Finally, we arbitrarily set $\sigma_\eta = 3$ and fix the precision of private sector beliefs at $\tau = 15$.[19] Table 6.1 summarizes the full set of parameter values.

Table 6.1: Calibrated Parameter Values

Category	Parameter	Value	Description / Target
Preferences	β	0.99	Subjective discount factor
Phillips Curve	κ	0.17	Slope of New-Keynesian Phillips Curve
Policy Preferences	λ	10	Weight on inflation in loss function
Shocks	ρ	0.9	Persistence of cost-push shock
	σ_u	0.2	Std. dev. of cost-push shock
	σ_η	3	Std. dev. of forecast error
Learning	τ	10	Precision of private sector beliefs

Because $\lim_{t\to\infty} \tau_t \to \infty$, beliefs eventually collapse to a single value and the reputation channel loses its force. To prevent this, we consider a modified economy in which precision remains constant,

$$\tau_{t+1} = \tau_t = \tau$$

This assumption preserves the central bank's ability to influence its reputation: if precision kept rising, that influence would vanish in the limit. By holding τ fixed, reputation does not converge to a single value but instead continues to fluctuate over time.

From (16), a key ingredient in optimal policy design is how inflation today shifts beliefs about the central bank's type. In Section 6 we prove that, for any prior μ_t and distribution

[19] We vary the value of τ later in this section.

of the forecast error η_t

$$\frac{\partial}{\partial \tilde{\pi}_t} \mathbb{E}^P_{t+s}\left[\tilde{\pi}_{t+s+1}\right] \approx -\kappa^{-1} \mathbb{COV}^P_{t+s}\left[\psi^\pi, s_\eta\left(\psi^\pi; \tilde{\pi}_t, \eta_t\right)\right]$$

The covariance term captures how informative inflation is about preferences: with $X_t > 0$ (a positive cost-push shock), higher inflation shifts beliefs towards dovish types; with $X_t < 0$ towards hawkish types. This insight is robust: it does not rely on functional form assumptions about beliefs or shocks. Because inflation is informative in this way, the optimal policy systematically overreacts relative to the myopic benchmark.

This logic implies that appointing a more hawkish but myopic central banker can mimic the optimal policy's overreaction and deliver welfare gains. For any myopic central banker of type $\tilde{\lambda}$, the steady-state welfare is

$$\begin{aligned} \mathcal{W}\left(\tilde{\lambda}\right) &= -\frac{1}{2} \mathbb{E}^{CB}\left[\tilde{y}\left(\tilde{\lambda}\right)^2 + \lambda \tilde{\pi}\left(\tilde{\lambda}\right)^2\right] \\ &= -\frac{1}{2} \overbrace{\left(\psi^y\left(\tilde{\lambda}\right)^2 + \lambda \psi^\pi\left(\tilde{\lambda}\right)^2\right)}^{\mathcal{W}_1(\tilde{\lambda})} \times \overbrace{\frac{1}{1 - \left(\beta \mathbb{E}^P\left[\psi^\pi \middle| \tilde{\lambda}\right] \rho\right)^2} \times \mathbb{E}^{CB}\left[\varepsilon^2\right]}^{\mathcal{W}_2(\tilde{\lambda})} \end{aligned} \tag{27}$$

Welfare combines two terms: a static trade-off between inflation and output gap stabilization ($\mathcal{W}_1\left(\tilde{\lambda}\right)$), and a dynamic term capturing how policy shapes reputation in the long run ($\mathcal{W}_2\left(\tilde{\lambda}\right)$). A myopic central bank optimizes only the first, whereas the optimal policy internalizes both. Since reputation improves with hawkishness,

$$\mathcal{W}'\left(\tilde{\lambda}\right)\Big|_{\tilde{\lambda}=\lambda} > 0$$

so appointing a more hawkish myopic central banker raises welfare. This result is robust and holds regardless of the learning model. Much like Rogoff (1985), hawkishness is substituted for the optimal policy. There, it substitutes commitment; here, it substitutes anticipating how the private sector interprets the central bank's actions.

How hawkish should the appointed central banker be? Under our calibration, a myopic central bank must set $\lambda \approx 37$ to generate the same long-run reputation as the optimal policy with $\lambda = 10$. Put differently, Arthur Burns implementing the optimal policy is equivalent to appointing a myopic Paul Volcker. Implementing the optimum is equivalent to appointing a central banker more than three times as hawkish. Figure 16a illustrates: the hawkish-myopic and optimal policies align at the steady state, but diverge from it at large shocks. Figure 16b shows that, when reputation is not at its steady state, the hawkish-myopic and optimal policies diverge.

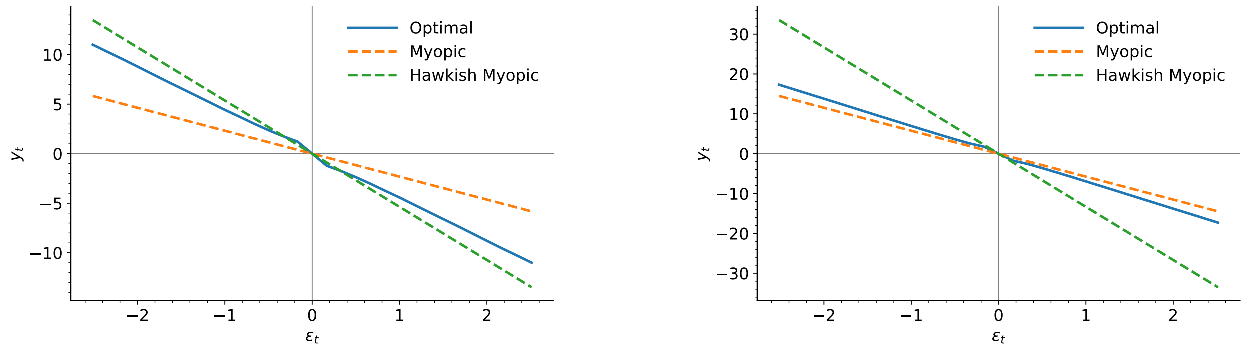

(a) Steady-state Reputation

(b) Reputation away from Steady State

Figure 16: Policy Functions: Optimal Policy vs Hawkish Myopic

The welfare expression (27) gives insights for a steady-state, but provides little information about how different policies perform at the business cycle frequency. We assess the relative performance of the hawkish-myopic central bank by simulating the economy for 5,000 periods. To eliminate dependence on initial conditions, we discard the first 1,000 periods as burn-in. Table 6.2 summarizes the outcomes.

Table 6.2: Performance of alternative policies

Policy	Welfare Loss	$\text{Cov}(y_t, \pi_t)$	$\sigma_y/\sigma_y^{\text{opt}}$	$\sigma_\pi/\sigma_\pi^{\text{opt}}$
Myopic	0.241	-1.203	0.645	1.684
Optimal	0.000	-1.076	1.000	1.000
Robust	0.055	-0.983	1.123	0.790

Notes: Welfare Loss is defined as $\dfrac{W^{\text{opt}} - W}{|W^{\text{opt}}|}$ with W^{opt} the baseline Optimal welfare.

Consistent with the theory, and relative to the myopic policy, the optimal policy trades off larger volatility of the output gap for lower volatility of inflation. Consequently, the optimal policy yields a more hawkish reputation and reduces the negative relationship between inflation and the output gap. The hawkish-myopic policy delivers similar welfare, reflected by a relatively small welfare loss. However, the policy mix is different: it prioritizes inflation stability relatively more, at the cost of larger volatility of the output gap.

The final, and perhaps most important test, is robustness. Optimal policy is sensitive to assumptions about how the private sector learns. If the central bank does not have the correct specification, welfare losses can be significant. Delegating policy to a hawkish myopic central banker can mitigate this risk. Suppose the central bank underestimates the precision of beliefs τ. We show in B that this implies underestimating the sensitivity of expectations to actions. We repeat the simulation exercise, assuming $\tau^{True} = 18 > \tau = 10$. Table 6.3 summarizes the outcomes.

Table 6.3: Performance under model misspecification. $\tau^{True} > \tau$

Policy	Welfare Loss	$\text{Cov}(y_t, \pi_t)$	$\sigma_y/\sigma_y^{\text{opt}}$	$\sigma_\pi/\sigma_\pi^{\text{opt}}$
Optimal (Higher)	0.000	−0.788	1.000	1.000
Optimal – Misspecified (Higher)	0.190	−1.312	0.932	1.796
Robust – Misspecified (Higher)	0.092	−0.983	1.009	1.196

Notes: Welfare Loss is defined as $\dfrac{W^{\text{opt}} - W}{|W^{\text{opt}}|}$ with W^{opt} the baseline Optimal welfare.

Because the central bank underestimates how sensitive beliefs are, overreaction is suboptimal. Reputation is therefore more dovish, leading to higher inflation volatility and a less favorable inflation-output trade-off. Delegating policy to a hawkish myopic central banker reduces welfare losses. By overreacting relative to the misspecified optimal policy, the hawkish delegate moves closer to the actual optimal response.

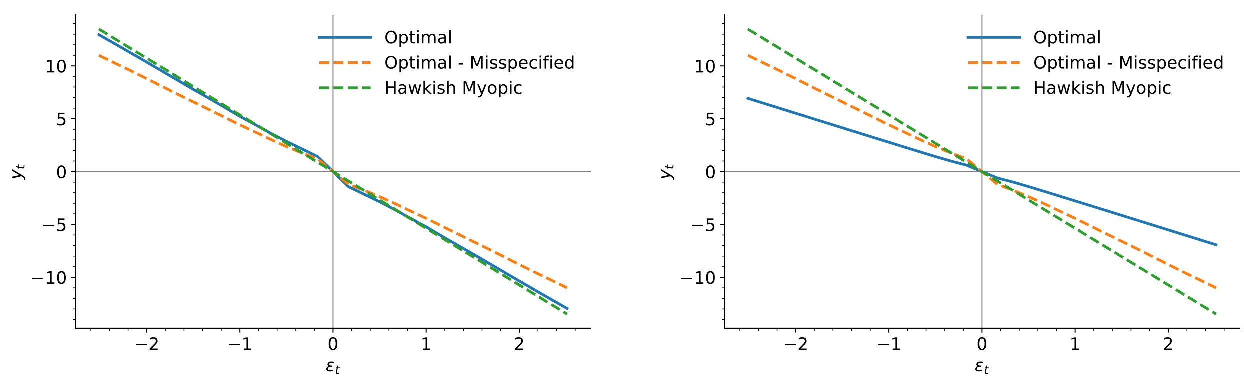

(a) $\tau^{True} > \tau$ **(b)** $\tau^{True} < \tau$

Figure 17: Comparison of policy functions under model misspecification.

Figure 16a compares the different policy functions. When the central bank underestimates its impact on the private sector's beliefs, the policies followed by a hawkish myopic central banker resemble the optimal policy more. Table 6.3 confirms this insight for the business cycle as well.

When the central bank overestimates the impact of its actions on the private sector's beliefs, the opposite occurs. Since there is too much overreaction, appointing a hawkish myopic central banker only makes things worse. Figure 16b plots the policy functions, and Table 6.4 displays the simulation exercise for $\tau^{True} = 2$.

Table 6.4: Performance under model misspecification. $\tau^{True} < \tau$

Policy	Welfare Loss	Cov(y_t, π_t)	$\sigma_y/\sigma_y^{\text{opt}}$	$\sigma_\pi/\sigma_\pi^{\text{opt}}$
Optimal (Lower)	0.000	-0.856	1.000	1.000
Optimal – Misspecified (Lower)	0.219	-1.002	1.614	0.742
Robust – Misspecified (Lower)	0.384	-1.014	1.851	0.638

Notes: Welfare Loss is defined as $\dfrac{W^{\text{opt}} - W}{|W^{\text{opt}}|}$ with W^{opt} the baseline Optimal welfare.

The hawkish myopic central banker overreacts even more, leading to larger output volatility and lower inflation volatility, when the opposite is needed. While appointing a hawkish myopic central banker can reduce welfare losses when the central bank underestimates the sensitivity of beliefs, it amplifies them when it is overconfident.

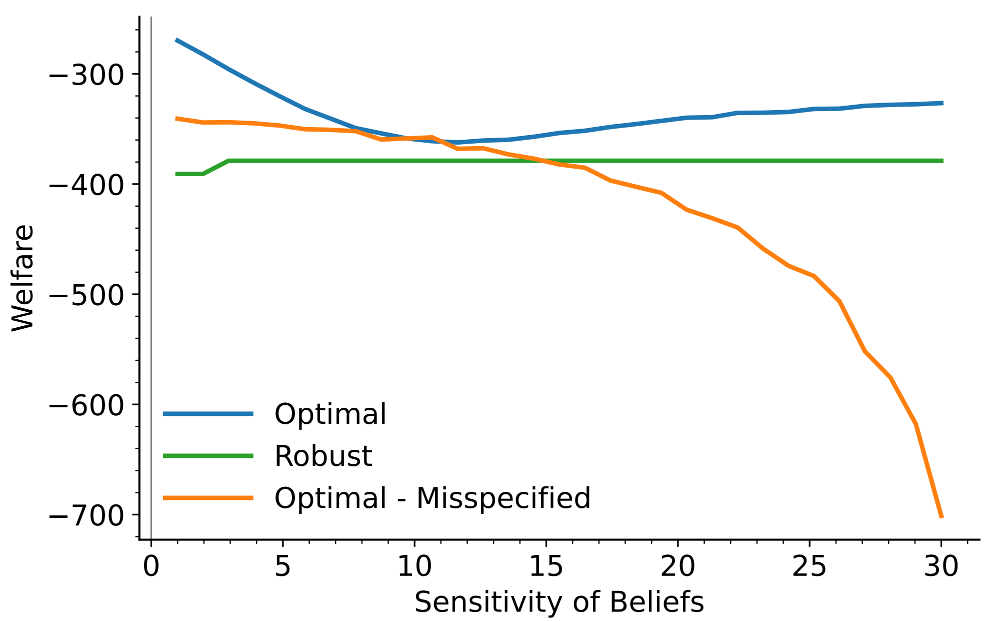

Figure 18: Welfare under alternative policies as a function of τ

Figure 18 plots welfare as a function of τ^{True}. Consistent with our previous findings, when the true value of τ is large enough, it becomes convenient to appoint a more hawkish myopic central banker. On the other hand, if the true value of τ is lower than expected, then it is no longer convenient.

Taken together, these exercises deliver a simple message. First, delegation to a hawkish-myopic central banker can replicate the outcomes of the optimal policy, in both a steady state and the business cycle. Second, the delegate must be substantially more hawkish. Third, and most importantly, delegation is robust to misspecification when the central bank underestimates its ability to change the private sector's beliefs. A simple institutional fix thus captures most of the benefits of dynamic optimization while remaining transparent and robust.

7 Concluding Remarks

This paper studies how a central bank with a dual mandate over output and inflation should design its monetary policy when reputation/credibility is an issue.

Theoretically, we compare the optimal responses under two monetary policy frameworks: A myopic central bank, which takes the private sector's expectations as given and optimizes period by period, and a central bank with reputational concerns, which includes the public's learning process in its optimized policy design. In response to a cost-push shock, the central bank optimally overreacts relative to the reactions of the myopic central bank. We discuss how this overreaction depends on the central bank's reputation, the size of the shock, and its implications for monetary policy surprises and reputation dynamics.

Empirically, we use U.S. survey data on forecasts to show that the private sector's perception of the central bank's reputation closely tracks our model. Reputation anchors the private sector's expectations, and the only way to improve it is by hawkish monetary surprises.

Finally, our analysis has practical implications. Delegating policy to a more hawkish but myopic central banker provides a simple way to approximate the optimal policy while avoiding significant welfare losses under misspecification. Such delegation can be implemented through appointment decisions, mandate design, or institutional arrangements that tilt policy in a hawkish direction. In this sense, an old solution to credibility problems provides a robust answer to a new one.

References

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-Cycle Anatomy," *American Economic Review*, 110, 3030–3070.

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

Barnichon, R. and G. Mesters (2023): "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, 113, 2809–2845.

Barro, R. J. (1986): "Reputation in a Model of Monetary Policy with Incomplete Information," *Journal of Monetary Economics*, 17, 3–20.

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

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024): "Perceptions About Monetary Policy," *The Quarterly Journal of Economics*, 139, 2227–2278.

Bauer, M. D. and E. T. Swanson (2022): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," NBER Working Paper 29939, National Bureau of Economic Research.

——— (2023): "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, 113, 664–700.

Bernanke, B. S. (2007): "Inflation Expectations and Inflation Forecasting," in *Speech at the Monetary Economics Workshop of the NBER Summer Institute*, "Anchored" means relatively insensitive to incoming data.

Bocola, L., A. Dovis, K. Jørgensen, and R. Kirpalani (2024): "Bond Market Views of the Fed," Working Paper 32620, National Bureau of Economic Research.

——— (2025): "Monetary Policy without an Anchor," Working paper.

Caballero, R. J. and A. Simsek (2022): "Monetary Policy with Opinionated Markets," *American Economic Review*, 112, 2353–92.

Canzoneri, M. B. (1985): "Monetary Policy Games and the Role of Private Information," *The American Economic Review*, 75, 1056–1070, publisher: American Economic Association.

Cukierman, A. and A. H. Meltzer (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54, 1099–1128.

Eggertsson, G. B. and D. P. Kohn (2023): "The Inflation Surge of the 2020s: The Role of Monetary Policy," *Brookings Papers on Economic Activity*, Summer, 45–131.

Eusepi, S. and B. Preston (2018): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56, 3–59.

Evans, G. W. and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*, Princeton University Press.

Federal Reserve Board (2016): "Tealbook A: Economic and Financial Conditions: Current Situation and Outlook," Tech. rep., Board of Governors of the Federal Reserve System, report to the FOMC on Economic Conditions and Monetary Policy.

Galí, J. (2003): "New Perspectives on Monetary Policy, Inflation, and the Business Cycle," in *Advances in Economic Theory and Econometrics, Eighth World Congress*, ed. by M. Dewatripont, L. Hansen, and S. Turnovsky, Cambridge, UK: Cambridge University Press, vol. III, 151–197.

García-Schmidt, M. and M. Woodford (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109, 86–120.

Gertler, M. and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Hamilton, J. D., S. Pruitt, and S. Borger (2011): "Estimating the Market-Perceived Monetary Policy Rule," *American Economic Journal: Macroeconomics*, 3, 1–28.

Hicks, J. R. (1946): *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*, Oxford: Clarendon Press, 2nd ed., originally published in 1939.

Känzig, D. R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 111, 1092–1125.

Kiley, M. T. (2025): "Monetary Policy Strategy and the Anchoring of Long-Run Inflation Expectations," Finance and Economics Discussion Series 2025-027, Board of Governors of the Federal Reserve System (U.S.).

King, R. G., Y. K. Lu, and E. S. Pastén (2008): "Managing Expectations," *Journal of Money, Credit and Banking*, 40, 1625–1666.

Kolesár, M. and M. Plagborg-Møller (2024): "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly," .

Kostadinov, R. and F. Roldán (2020): "Credibility Dynamics and Disinflation Plans," IMF Working Paper WP/20/85, International Monetary Fund.

Kreps, D. M. (1998): "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*, ed. by D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Schwartz, Cambridge University Press, 242–274.

Krugman, P. R. (1998): "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 29, 137–206, economic Studies Program.

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

Känzig, D. R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 111, 1092–1125.

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.

Marcet, A. and T. J. Sargent (1989): "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48, 337–368.

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

Nakamura, E. and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

Powell, J. H. (2020): "New Economic Challenges and the Fed's Monetary Policy Review," in *Jackson Hole Economic Policy Symposium: Navigating the Decade Ahead: Implications for Monetary Policy*, Jackson Hole, Wyoming: Federal Reserve Bank of Kansas City.

Ramey, V. A. (2016): "Macroeconomic Shocks and Their Propagation," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 71–162.

Ramey, V. A. and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *The Quarterly Journal of Economics*, 100, 1169–1189.

Sastry, K. (2022): "Disagreement About Monetary Policy," *SSRN Electronic Journal*, available at SSRN: https://ssrn.com/abstract=3421723 or http://dx.doi.org/10.2139/ssrn.3421723.

Svensson, L. E. O. (2003): "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules," *Journal of Economic Literature*, 41, 426–477.

Vickers, J. (1986): "Signalling in a Model of Monetary Policy with Incomplete Information," *Oxford Economic Papers*, 38, 443–455, publisher: Oxford University Press.

Werning, I. (2022): "Expectations and the Rate of Inflation," Working Paper 30260, National Bureau of Economic Research.

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

——— (2003b): "Optimal Interest-Rate Smoothing," *The Review of Economic Studies*, 70, 861–886.

A The Economy

A.1 Private Sector and Government

Households There is a continuum of identical households. In each period, the representative household derives utility from consuming a continuum of differentiated final goods, $c_t(j)$ for $j \in [0,1]$, and working N_t units. Assuming separability between consumption and labor and isoelastic functions, the lifetime utility is given by

$$\mathcal{U}_0 \equiv \mathbb{E}_0^P \left[\sum_{t=0}^{\infty} \beta^t Z_t \left(\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \frac{N_t^{1+\varphi}}{1+\varphi} \right) \right]$$

where $C_t = \left[\int_0^1 c_t(j)^{1-\frac{1}{\epsilon_t}} dj \right]^{\frac{1}{1-\frac{1}{\epsilon_t}}}$ is the Dixit-Sitglitz consumption aggregator, ϵ_t is the elasticity of substitution across final goods varieties, and $\mathbb{E}_t^P [\cdot]$ is the private sector's expectation operator, which we assume satisfies the Law of Iterated Expectations.[20] In our analysis below, we allow the elasticity of substitution between varieties to be time-varying to allow for markup shocks that will simulate an inflationary episode.

Labor markets are competitive, and the representative household takes nominal wages W_t as given. Moreover, households can trade one-period nominal risk-free bonds, B_t, which the government issues. The representative household's budget constraint is given by

$$P_t C_t + B_t = W_t N_t + \Pi_t + (1 + i_{t-1}) B_{t-1} + T_t$$

where $P_t = \left[\int_0^1 P_t(j)^{1-\epsilon_t} dj \right]^{\frac{1}{1-\epsilon_t}}$ is the aggregate price index, i_{t-1} is the short-term nominal interest rate from period $t-1$ to t, Π_t denotes the nominal firms' profits, and T_t are lump sum transfers. Then, given B_{-1}, aggregate prices, government policy, and firms' profits, households choose a path for consumption, labor, and asset portfolio that maximizes their utility $\mathcal{U}_0$ subject to the budget constraint at every period. As a result of this optimization process, households' optimal behavior is captured by an aggregate Euler equation, consumption-labor optimal allocation and the budget constraint.

[20] We fully describe the household's information set and belief updating process later in this appendix.

Firms Firms and households share the same information and understanding of the functioning of the economy, so their expectation operator is also $\mathbb{E}_t^P[\cdot]$. There is monopolistic competition in the final goods market. Each producer of variety $j \in [0,1]$ has access to the same production technology

$$Y_t(j) = A_t N_t(j)^{1-\alpha}$$

where A_t is stochastic aggregate productivity and $\alpha \in (0,1)$ controls the degree of decreasing returns to scale in this economy. Firms set prices à la Calvo, with a probability $1-\theta \in [0,1]$ of changing prices every period. There is a production subsidy to offset the firms' market power so that the non-stochastic steady state is efficient.

Government Each period, the government issues short-term nominal bonds to finance lump-sum transfers and past government debt. We abstract from government expenditure. Then, the government's budget constraint is $(1+i_{t-1})B_{t-1} + T_t = B_t$. In addition, there is a monetary authority (Central Bank) that sets the path of nominal interest, $\{i_t\}_{t\geqslant 0}$.

Log-linear model Throughout the paper, we study the optimal monetary policy using a linear-quadratic approach. Thus, we log-linearize the model and obtain a version around its deterministic, efficient steady state.[21] Hence, up to first-order approximation, the aggregate Euler Equation and goods market clearing condition (i.e., $Y_t = C_t$) leads to the dynamic IS equation:

$$y_t = \mathbb{E}_t^P[y_{t+1}] - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t^P[\pi_{t+1}] - r_t^n\right) \tag{A.1}$$

where y_t is the log deviation of the output concerning the efficient level of output (output gap), $\pi_t = P_t/P_{t-1}$ is the price inflation rate between $t-1$ and t, and

$$r_t^n \equiv \rho + \vartheta \mathbb{E}_t^P[\Delta a_{t+1}] - \mathbb{E}_t^P[\Delta z_{t+1}]$$

with $z_t \equiv \log Z_t$, $a_t \equiv \log A_t$, and ϑ being a function of the model's parameters. Similarly, up to first order, the solution to the firms' problem, together with households' labor supply,

[21] When prices are flexible, reputation has no bite in the economy, so the log-linear approximation does not rely on assuming a particular value for reputation in the long run.

implies a New-Keynesian Phillips Curve[22]:

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t^P \left[\pi_{t+1}\right] + \varepsilon_t \tag{A.2}$$

where ε_t denotes the log deviation of the efficient output level (*frictionless* concept) with respect to the natural output level (*flexible-prices* concept), and κ is a function of the model's parameters. In our model, these disturbances arise from markup shocks, $\frac{\epsilon_t}{\epsilon_t - 1} \geqslant 1$, but, in general, our analysis does not change if we consider different sources of these cost-push shocks.

A.2 Central Bank and Monetary Policy Regime

The monetary authority has a dual mandate over output gaps and inflation, given by the following welfare loss function:

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{k=0}^{\infty} \beta^k \left(y_{t+k}^2 + \lambda \pi_{t+k}^2\right)\right] \tag{A.3}$$

where $\mathbb{E}_t^{CB}[\cdot]$ is the Central Bank's rational expectations operator. This loss function captures the idea that the Central Bank wishes to stabilize inflation around zero and output around its efficient level. The importance of inflation stabilization, relative to output stabilization, is controlled by the fixed parameter $\lambda \geqslant 0$.

Central Bank's Information Set and Policy Regime Each period t, the Central Bank announces the nominal interest rate i_t before observing the realization of the demand shocks hitting the economy that period. Thus, the information set at t of the Central Bank is given by the history of output gaps and inflation, the cost-push shocks, and the private sector's behavioral equations: $\mathcal{I}_t^{CB} \equiv \{\mathcal{M}^{CB}, h_{t-1}, \mu_t, \boldsymbol{\varepsilon}_t\}$ where $\mathcal{M}^{CB}$ contains the structure of the economy, i.e., equations (A.1) and (2) plus the data generator process of the exogenous variables, $h_{t-1} \equiv \{y_s, \pi_s, i_s, \varepsilon_s, z_s, a_s\}_{s \leqslant t-1}$, μ_t encodes the private sector's beliefs at t, and $\boldsymbol{\varepsilon}_t = \{\mathbb{E}_t[\varepsilon_{t+s}]\}_{s \geqslant 0}$ the cost-push shocks.

Even though the Central Bank's model for the economic structure is correctly specified, it sets monetary policy at t with imperfect information on the realization of z_t, and a_t.

[22] This is a direct consequence of the expectation process satisfying the Law of Iterated Expectations. See Evans and Honkapohja (2001).

Formally, given $\mathcal{I}_t^{CB}$, the Central Bank privately forecasts exogenous variable realizations, $\tilde{r}_t^n \equiv \mathbb{E}_t^{CB}[r_t^n]$, and chooses i_t to maximize (A.3) subject to the private sector's ex-ante equilibrium conditions, (A.1) and (A.2):

$$\tilde{y}_t = \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[y_{t+1}]\right] - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[\pi_{t+1}]\right] - \tilde{r}_t^n\right)$$

$$\tilde{\pi}_t = \kappa\tilde{y}_t + \beta\mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[\pi_{t+1}]\right] + \varepsilon_t$$

where $\tilde{y}_t \equiv \mathbb{E}_t^{CB}[y_t]$ and $\tilde{\pi}_t \equiv \mathbb{E}_t^{CB}[\pi_t]$ are the Central Bank's ex-ante expected values of output and inflation, respectively.[23] Under the learning structure of our model (described below), it turns out that the terms with double expectations are simply the private sector's expectation of the Central Bank's choices:

$$\mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[y_{t+1}]\right] = \mathbb{E}_t^P[\tilde{y}_{t+1}] \qquad \text{and} \qquad \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[\pi_{t+1}]\right] = \mathbb{E}_t^P[\tilde{\pi}_{t+1}].$$

Under this monetary policy regime, differences between the output gap and its Central Bank's ex-ante expected value, $\eta_t \equiv y_t - \tilde{y}_t$, may arise from forecasting errors about the realization of aggregate demand, productivity, or markup shocks. Note that after the realization of error η_t, the resulting inflation rate is $\pi_t = \tilde{\pi}_t + \kappa\eta_t$.

Re-writing the Central Bank's problem Given the quadratic objective function and the i.i.d. property of η_t with respect to the Central Bank's expectation operator, we can rewrite the Central Bank's objective function in terms of the ex-ante output and inflation (i.e., $\tilde{y}_{t+k}$ and $\tilde{\pi}_{t+k}$),

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{k=0}^{\infty}\beta^k\left(\tilde{y}_{t+k}^2 + \lambda\tilde{\pi}_{t+k}^2\right)\right] + t.i.p. \tag{A.4}$$

where $t.i.p.$ are terms independentt of policy. Finally, under the assumption that disturbances are properly bounded, zero lower bound is never binding so that we can re-express the Central Bank's problem as directly choosing the (ex-ante) allocation path $\{\tilde{y}_{t+k}, \tilde{\pi}_{t+k}\}_{k=0}^{\infty}$ to maximize (3) subject to

$$\tilde{\pi}_t = \kappa\tilde{y}_t + \beta\mathbb{E}_t^P[\tilde{\pi}_{t+1}] + \varepsilon_t \tag{A.5}$$

[23] The Central Bank's ex-ante expectations of the private sector's forward looking terms is present in the literature of monetary policy with disagreements. See, e.g. Caballero and Simsek (2022) and Sastry (2022).

and the below-described public's learning process. It is worth emphasizing that if the Central Bank chooses to implement output $\tilde{y}_t$ and inflation $\tilde{\pi}_t$, the final equilibrium outcome at t are given by $\tilde{y}_t + \eta_t$ and $\tilde{\pi}_t + \kappa\eta_t$.

Although this Central Bank's problem, given by (3) and (A.5), seems to be a linear-quadratic one, the below-described expectation formation process introduces a nonlinearity into the Phillips Curve. In particular, $\mathbb{E}_t^P[\tilde{\pi}_{t+1}]$ depends on $\{\tilde{\pi}_s\}_{s\leqslant t}$ in a nonlinear way. The main objective of this paper is to characterize how this nonlinear dependence shapes the optimal monetary policy.

A.3 Proof of Lemma 1

Proof. The first-order condition of a Central Bank with parameter $\lambda = \tilde{\lambda}$ is

$$\tilde{y}_t + \kappa\tilde{\lambda}\tilde{\pi}_t = 0$$

Taking expectations as given, this yields

$$\begin{aligned}\tilde{y}_t\left(\tilde{\lambda}\right) =& - \psi^y \left(\varepsilon_t + \beta\mathbb{E}_t^{CB}\left[\mathbb{E}_t^P\left[\pi_t\right]\right]\right)\\ \tilde{\pi}_t\left(\tilde{\lambda}\right) =& \psi^\pi \left(\varepsilon_t + \beta\mathbb{E}_t^{CB}\left[\mathbb{E}_t^P\left[\pi_t\right]\right]\right)\end{aligned}$$

We conjecture that the Central Bank follows a linear policy rule, and they are not aware of the private sector's perceived bias:

$$\begin{aligned}\tilde{y}_t\left(\tilde{\lambda}\right) =& - \psi^y\left(\tilde{\lambda}\right)\sum_{s=0}^{\infty}\Theta_s\mathbb{E}_t^P\left[\varepsilon_{t+s}\right]\\ \tilde{\pi}_t\left(\tilde{\lambda}\right) =& \psi^\pi\left(\tilde{\lambda}\right)\sum_{s=0}^{\infty}\Theta_s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

Matching coefficients for $s = 0$

$$\Theta_0 = 1$$

For $s = 1$, using that there is no anticipated learning

$$\Theta_1 = \beta\mathbb{E}_t^P\left[\psi^\pi\right]$$

For any arbitrary s, we have

$$\Theta_s = \beta \mathbb{E}_t^P \left[\psi^\pi\right] \Theta_{s-1} = \left(\beta \mathbb{E}_t^P \left[\psi^\pi\right]\right)^s$$

Putting everything together,

$$\begin{aligned}\tilde{y}_t\left(\tilde{\lambda}\right) =& -\psi^y \sum_{s=0}^{\infty} \left(\beta \mathbb{E}_t^P \left[\psi^\pi\right]\right) \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\ \tilde{\pi}_t\left(\tilde{\lambda}\right) =& \psi^\pi \sum_{s=0}^{\infty} \left(\beta \mathbb{E}_t^P \left[\psi^\pi\right]\right) \mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

■

A.4 Proof of Lemma 2

Proof. Rewrite the expected value as

$$\begin{aligned}\mathbb{E}_t^P\left[\psi^\pi\right] = \int_0^1 \psi^\pi f_\mu^{\psi^\pi}\left(\psi^\pi\right) d\psi^\pi &= \int_0^1 \left(\int_0^{\psi^\pi} ds\right) f_\mu^{\psi^\pi}\left(\psi^\pi\right) d\psi^\pi \\ &= \int_0^1 \left(\int_s^{\psi} f_\mu^{\psi^\pi}\left(\psi^\pi\right) d\psi^\pi\right) ds = \int_0^1 \left(1 - F_\mu^{\psi^\pi}\left(\psi^\pi\right)\right) d\psi^\pi\end{aligned}$$

If μ_t' first-order stochastically dominates μ_t then $F_{\mu_t}^{\lambda}(\lambda) \geqslant F_{\mu_t'}^{\lambda}(\lambda)$. Since $\psi^\pi(\lambda)$ is a decreasing function of λ then $F_{\mu_t'}^{\psi^\pi}\left(\psi^\pi\right) \geqslant F_{\mu_t}^{\psi^\pi}\left(\psi^\pi\right)$ and therefore $\mathbb{E}_t^P\left[\psi^\pi\right]$ is lower under μ_t' than under μ_t ■

A.5 Proof of Lemma 3

Proof. This is a direct consequence of Jensen's inequality and the convexity of $\psi^\pi(\lambda)$. ■

B Solution to the Optimal Policy Problem

B.1 Learning Structure

In this subsection, we analytically characterize the learning structure of our model. Using our functional form assumptions in (11) yields

$$\begin{aligned}\mu_{t+1}\left(\tilde{\psi}^{\pi}\right) \propto& \sqrt{\tau_\eta} exp\left(-\frac{\tau_\eta}{2}\left(\kappa^{-1}\left(\tilde{\pi}_t - \tilde{\psi}^{\pi}X_t + \kappa\eta_t\right)\right)^2\right) \times \sqrt{\tau_t} exp\left(-\frac{\tau_t}{2}\left(\tilde{\psi}^{\pi} - \overline{\psi}_t\right)^2\right)\\ =& exp\left(-\frac{\tau_\eta\left(\kappa^{-1}X_t\right)^2 + \tau_t}{2}\left(\tilde{\psi}^{\pi} - \overline{\psi}_{t+1}\right)^2\right)\end{aligned}$$

where

$$\overline{\psi}_{t+1} = \omega_t X_t^{-1}\left(\tilde{\pi}_t + \kappa\eta_t\right) + \left(1-\omega_t\right)\overline{\psi}_t \quad \text{where} \quad \omega_t = \frac{\left(\kappa^{-1}X_t\right)^2\tau_\eta}{\tau_t + \left(\kappa^{-1}X_t\right)^2\tau_\eta}$$

therefore the model exhibits a conjugate prior and $\mu_{t+1}\left(\tilde{\psi}^{\pi}\right) \sim \Psi_{t+1}|\,\Psi_{t+1} \in [0,1]$ where $\Psi_{t+1} \sim \mathcal{N}\left(\overline{\psi}_{t+1}, \tau_{t+1}^{-1}\right)$, with $\tau_{t+1} = \tau_t + \left(\kappa^{-1}X_t\right)^2\tau_\eta$

B.2 The Central Bank's Problem

After observing the cost push shocks $\{\mathbb{E}_t\left[\varepsilon_{t+s}\right]\}_{s\geqslant 0}$ the central bank maximizes

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t\left[\sum_{s=0}^{\infty}\left(\tilde{y}_{t+s}^2 + \lambda\tilde{\pi}_t^2\right]\right]$$

subject to

(i) The New-Keynesiean Phillips Curve,

$$\tilde{\pi}_t = \kappa\tilde{y}_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] + \varepsilon_t$$

(ii) The private sector's expectation formation process and learning structure,

$$\begin{aligned}\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] =& \mathbb{E}_t^P\left[\psi^{\pi}\right]\mathbb{E}_t^P\left[X_{t+1}\right]\\ \mathbb{E}_t^P\left[\psi^{\pi}\right] =& \overline{\psi}_t - \frac{1}{\sqrt{\tau}}\frac{\phi\left(\tau\left(1-\overline{\psi}_t\right)\right) - \phi\left(-\tau\overline{\psi}_t\right)}{\Phi\left(\tau\left(1-\overline{\psi}_t\right)\right) - \Phi\left(-\tau\overline{\psi}_t\right)}\end{aligned}$$

$$\overline{\psi}_{t+1} = \omega_t X_t^{-1} \left(\tilde{\pi}_t + \kappa \eta_t\right) + \left(1 - \omega_t\right) \overline{\psi}_t$$

where $\omega_t = \frac{\left(\kappa^{-1} X_t\right)^2 \tau_\eta}{\tau + \left(\kappa^{-1} X_t\right)^2 \tau_\eta}$ and $\tau_{t+1} = \tau_t + \left(\kappa^{-1} X_t\right)^2 \tau_\eta$.

Let μ_t denote the Lagrange multipliers. The first-order conditions are

$$\begin{aligned}
\tilde{y}_t =& \kappa \mu_t \\
\lambda \tilde{\pi}_t =& - \mu_t + \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_t} \mu_{t+s} \right] \\
\tilde{\pi}_t =& \kappa \tilde{y}_t + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1}\right] + \varepsilon_t
\end{aligned}$$

Combining the three equations, we obtain

$$\tilde{y}_t = \tilde{y}_t^M + \psi^\pi \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s} \right] \tag{C.1}$$

$$\tilde{\pi}_t = \tilde{\pi}_t^M + \kappa \psi^\pi \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s} \right] \tag{C.2}$$

B.3 Solution Method

We assume the cost-push shock follows an AR(1) process, with persistence ρ. Given $\left(\overline{\psi}_t, \varepsilon_t\right)$, together with $\eta_t \sim F$ being white noise, the central bank solves

$$V\left(\overline{\psi}_t, \varepsilon_t\right) = \min_{\{y_{t+s}, \pi_{t+s}\}_{s \geqslant 0}} \frac{1}{2} \mathbb{E}_t^{CB} \left[\sum_{s=0}^{\infty} \beta^s \left(y_{t+s}^2 + \lambda \pi_{t+s}^2\right) \right]$$

subject to

$$\begin{aligned}
\varepsilon_{t+1} =& \rho \varepsilon_t + \nu_{t+1} \\
\pi_t =& \kappa y_t + \beta \mathbb{E}_t^P \left[\pi_{t+1}\right] + \varepsilon_t \\
\mathbb{E}_t^P \left[\pi_{t+s}\right] =& \mathbb{E}_t^P \left[\psi^\pi\right] \mathbb{E}_t^P \left[X_{t+s}\right] \\
\mathbb{E}^P \left[\psi^\pi\right] =& \overline{\psi}_t - \frac{1}{\sqrt{\tau}} \frac{\phi\left(\sqrt{\tau_t \left(1 - \overline{\psi}_t\right)}\right) - \phi\left(-\sqrt{\tau_t \overline{\psi}_t}\right)}{\Phi\left(\sqrt{\tau_t \left(1 - \overline{\psi}_t\right)}\right) - \Phi\left(-\sqrt{\tau_t \overline{\psi}_t}\right)}
\end{aligned}$$

$$
\begin{aligned}
X_t =& \frac{1}{1-\beta\mathbb{E}_t^P\left[\psi^\pi\right]\rho}\varepsilon_t \\
\overline{\psi}_{t+1} =& \overline{\psi}_t + \omega_t X_t^{-1}\left(\pi_t + \kappa\eta_t - \overline{\psi}_t X_t\right) \\
\omega_t =& \frac{\left(\kappa^{-1}X_t\right)^2\tau_\eta}{\left(\kappa^{-1}X_t\right)^2\tau_\eta + \tau}
\end{aligned}
$$

We can rewrite the problem recursively as

$$
V\left(\overline{\psi},\varepsilon\right) = \min_y \quad \frac{1}{2}\left(y^2 + \lambda\left\{\kappa y + G\left(\overline{\psi},\epsilon\right)\right\}^2\right) + \beta\mathbb{E}\left[V\left(\overline{\psi}',\epsilon'\right)\middle|\varepsilon\right]
$$

subject to

$$
\begin{aligned}
\overline{\psi}' =& \frac{\omega\left(\overline{\psi},\varepsilon\right)}{G\left(\overline{\psi},\varepsilon\right)}\left(\kappa\left(y+\eta\right) + G\left(\overline{\psi},\varepsilon\right)\right) + \left(1-\omega\left(\overline{\psi},\varepsilon\right)\right)\overline{\psi} \\
\varepsilon' =& \rho\varepsilon + \nu
\end{aligned}
$$

where

$$
\begin{aligned}
g\left(\overline{\psi}\right) =& \overline{\psi} - \frac{1}{\sqrt{\tau}}\frac{\phi\left(\sqrt{\tau}\left(1-\overline{\psi}\right)\right) - \phi\left(-\sqrt{\tau}\overline{\psi}\right)}{\Phi\left(\sqrt{\tau}\left(1-\overline{\psi}\right)\right) - \Phi\left(-\sqrt{\tau}\overline{\psi}\right)} \\
\omega\left(\overline{\psi},\varepsilon\right) =& \frac{\left(\kappa^{-1}G\left(\overline{\psi},\varepsilon\right)\right)^2\tau_\eta}{\left(\kappa^{-1}G\left(\overline{\psi},\varepsilon\right)\right)^2\tau_\eta + \tau} \\
G\left(\overline{\psi},\varepsilon\right) =& \left(1-\beta\rho g\left(\overline{\psi},\varepsilon\right)\right)^{-1}\varepsilon
\end{aligned}
$$

We discretize ε using the Rowenhorst method: $\boldsymbol{\varepsilon}$ (S-states vector) and $\boldsymbol{P}$ (transition matrix). We define a grid for the forecast error $\boldsymbol{\eta}$. We build a grid for the parameters of the initial prior $\overline{\boldsymbol{\psi}}$ (I elements). We propose an $I \times S$ matrix $\boldsymbol{V}^{old}$ where their elements (i,s) specify the value for $V\left(\overline{\psi}_i,\varepsilon_s\right)$. Given $\boldsymbol{V}^{old}$:

1. Compute
$$
\tilde{\boldsymbol{V}} = \boldsymbol{V}^{old} \times \boldsymbol{P}^T
$$

2. For each s:

   i. Take $\tilde{\boldsymbol{V}}_{\bullet,s}$

   ii. Compute $G\left(\overline{\psi}_i,\varepsilon_s\right)$, $\omega\left(\overline{\psi}_i,\varepsilon_s\right)$ and $g\left(\overline{\psi}_i\right)$.

iii. Compute $\overline{\psi}^{\prime}(y,\eta)$.

iv. Solve

$$V\left(\overline{\psi}_i,\varepsilon_i\right)=\max_y\left\{\frac{1}{2}\left(y^2+\lambda\left(\kappa y+G\right)^2\right)+\beta\sum_\eta interp\left(\tilde{\boldsymbol{V}},\overline{\psi}^{\prime}\left(y,\eta\right)\right)\boldsymbol{F}_\eta\right\}$$

v. Store the solution in $\boldsymbol{V}^{new}$.

3. Check $\left\|\boldsymbol{V}^{new}-\boldsymbol{V}^{old}\right\|_\infty$ and update/stop.

B.4 Proof of Proposition 1

Proof. From (C.1) and (C.2), it suffices to show that

$$\mathcal{I}_t=\sum_{s=0}^{\infty}\beta^s\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial\tilde{\pi}_t}\tilde{y}_{t+s}\right]\tag{C.3}$$

is negative when $X_t>0$, and positive when $X_t<0$. Recall that

$$\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]=\mathbb{E}_{t+s}^P\left[\psi^\pi\right]\mathbb{E}_{t+s}^P\left[X_{t+s+1}\right]$$

Then, we have

$$\begin{aligned}\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial\tilde{\pi}_t}&=\mathbb{E}_{t+s}^P\left[\psi^\pi\right]\mathbb{E}_{t+s}^P\left[X_{t+s+1}\right]\left(\frac{\frac{\partial\mathbb{E}_{t+s}^P\left[\psi^\pi\right]}{\partial\tilde{\pi}_t}}{\mathbb{E}_{t+s}^P\left[\psi^\pi\right]}+\frac{\frac{\partial\mathbb{E}_{t+s}^P\left[X_{t+s+1}\right]}{\partial\tilde{\pi}_t}}{\partial\tilde{\pi}_t}\right)\\&=\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]\left(\frac{\frac{\partial\mathbb{E}_{t+s}^P\left[\psi^\pi\right]}{\partial\tilde{\pi}_t}}{\mathbb{E}_{t+s}^P\left[\psi^\pi\right]}+\frac{\frac{\partial\mathbb{E}_{t+s}^P\left[X_{t+s+1}\right]}{\partial\tilde{\pi}_t}}{\partial\tilde{\pi}_t}\right)\end{aligned}$$

Since $\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]\tilde{y}_{t+s}<0$, we need to show that the term in parenthesis is positive when $X_t>0$ and negative when $X_t<0$. From the learning structure (14) and the definition of X_t: with a positive cost-push shock, higher inflation signals a lower commitment with inflation stability, whereas the opposite occurs with a negative shocks.

■

B.5 Proof of Proposition 2

Proof. Recall that

$$\frac{\partial \mathbb{E}_t^P\left[\psi^\pi\right]}{\partial \overline{\psi}_t} = \tau_t Var_t\left(\psi^\pi\right)$$

and that $Var_t\left(\psi^\pi\right)$ is increasing in $\overline{\psi}_t$ when $\mathbb{E}_t^P\left[\psi^\pi\right] < \frac{1}{2}$ and decreasing when $\mathbb{E}_t^P\left[\psi^\pi\right] > \frac{1}{2}$. Then, $\frac{\partial \mathbb{E}_{t+s}^P[\tilde{\pi}_{t+s+1}]}{\partial \tilde{\pi}_t}$ increases with $\overline{\psi}_t$ is small, and decreases when $\overline{\psi}_t$ is large. Everything else equal, this implies that $\mathcal{I}_t$ increases with $\overline{\psi}_t$ at first, and then decreases. ■

B.6 Sensitivity of expectations to belief precision

Recall the sensitivity of beliefs is given by (18)

$$\frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial \overline{\psi}_t} = \frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial \mathbb{E}_{t+s}^P\left[\psi^\pi\right]} \frac{\partial \mathbb{E}_{t+s}^P\left[\psi^\pi\right]}{\partial \overline{\psi}_{t+s}} \frac{\partial \overline{\psi}_{t+s}}{\partial \tilde{\pi}_t}$$

Conditional on a certain reputation, $\mathbb{E}_t^P\left[\psi^\pi\right]$, the sensitivity of expectations depends on two elements:

1. The sensitivity of reputation to the prior's location. Under our functional form assumptions,

$$\frac{\partial \mathbb{E}_{t+s}^P\left[\psi^\pi\right]}{\partial \overline{\psi}_{t+s}} = \tau_{t+s} \mathbb{V}_{t+s}^P\left[\psi^\pi\right]$$

   which is increasing in τ_{t+s}. The tighter the beliefs are around the prior's location, the more sensitive they become to changes in it.

2. The sensitivity of the central bank's choices on the prior's location

$$\frac{\partial \overline{\psi}_{t+s}}{\partial \tilde{\pi}_t} = \left(\prod_{j=1}^{s}\left(1-\omega_{t+s-j}\right)\right)\omega_t$$

   which, for $s \geqslant 2$ increases in τ when τ is small, and decreases when τ is large.[24] If $\overline{\psi}_{t+1}$ depends almost entirely on actions at period t, then those actions will have little impact on $\overline{\psi}_{t+s}$ for $s \geqslant 2$.

Taken together, these facts imply that when τ is relatively small, an increase in τ increases

[24] For $s = 1$ is it always decreasing in τ.

the sensitivity of expectations to the central bank's actions. And when τ is large, an increase in τ may decrease the sensitivity of expectations to the central bank's actions.

B.7 Proof of Proposition 3

Proof. From the central bank's perspecive, the monetary policy surprise is

$$\tilde{\varepsilon}_t^m = \mathbb{E}_t^{CB}\left[i_t - \mathbb{E}_t^P\left[i_t\right]\right] = -\frac{\sigma}{\kappa}\left[\tilde{\pi}_t - \mathbb{E}_t^P\left[\psi^\pi\right]X_t\right]$$

Then, we have

$$\tilde{\varepsilon}_t^m = -\frac{\sigma}{\kappa}\left[\left(\tilde{\pi}_t - \tilde{\pi}_t^M\right) + \left(\tilde{\pi}_t^M - \mathbb{E}_t^P\left[\psi^\pi\right]X_t\right)\right] = -\left[\left(\tilde{\pi} - \tilde{\pi}^M\right) + \left(\psi^\pi - \mathbb{E}_t^P\left[\psi^\pi\right]\right)X_t\right]$$

Suppose for simplicity that $X_t > 0$. The first term is negative from Propositon 1. The second term is negative when $\psi^\pi < \mathbb{E}_t^P\left[\psi^\pi\right]$. When $\psi^\pi < \mathbb{E}_t^P\left[\psi^\pi\right]$ then both effects point in the same direction, and thus $\tilde{\varepsilon}_t^m$. When $\psi^\pi < \mathbb{E}_t^P\left[\psi^\pi\right]$ the effects point in opposite directions. From Proposition 2, the first term is zero then $\mathbb{E}_t^P\left[\psi^\pi\right] = 0$, peaks at an intermediate value and then decreases to zero when $\mathbb{E}_t^P\left[\psi^\pi\right] = 1$. The second term is strictly decreasing in $\mathbb{E}^P\left[\psi^\pi\right]$, and is positive at $\mathbb{E}_t^P\left[\psi^\pi\right] = 0$. Then, there exist a value $\widehat{\psi} < \psi^\pi$ such that $\tilde{\varepsilon}_t^m = 0$.

■

B.8 Proof of Proposition 4

Proof. Suppose for simplicity that $X_t > 0$. Rewrite (15) as

$$\overline{\psi}_{t+1} = \overline{\psi}_t + \omega_t\left[\left(\mathbb{E}_t^P\left[\psi^\pi\right] - \overline{\psi}_t - \frac{\kappa}{\sigma}X_t^{-1}\varepsilon_t^m\right)\right]$$

The first term is negative when $\mathbb{E}_t^P\left[\psi^\pi\right] > \frac{1}{2}$ and positive when $\mathbb{E}_t^P\left[\psi^\pi\right] < \frac{1}{2}$. From Proposition 3, there exist a value $\widehat{\psi} < \psi^\pi$ such that $\varepsilon_t^m > 0$ when $\mathbb{E}_t^P\left[\psi^\pi\right] > \widehat{\psi}$. Then, there exist a threshold such that reputation improves, from the ex-ante perspective of the central bank.

■

B.9 Proof of Proposition 5

Proof. When the current shock is *iid*, it has no direct impact on $\tilde{y}_{t+s}$ beyond its effect on reputation. Since the effective weight $\omega_t X_t^{-1}$ is decreasing in the size of the shock for large values, then the size of overreaction decreases with the size of the shock. ■

C Extensions of the Baseline Model

C.1 Uncertainty about Interest Rate Smoothing

Suppose for simplicity that prices are fixed. We slightly modify the information structure. The central bank and the private sector have the same information about the natural rate, but there are random monetary policy shocks η_t. Let $\tilde{i}_t$ denote the ex-ante interest rate. We have $i_t = \tilde{i}_t - \frac{1}{\sigma}\eta_t$. Then, the (ex-ante) economy private sector is only characterized by the Euler Equation

$$\tilde{y}_t = \mathbb{E}_t^P [\tilde{y}_{t+1}] - \frac{1}{\sigma} \left(\tilde{i}_t - \mathbb{E}_t [r_t^n]\right) \tag{D.1}$$

Under the baseline model, the Divine Coincidence holds, and the central bank can always implement the first best. Instead, suppose the central bank has the triple mandate (19):

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB} \left[\sum_{s=0}^{\infty} \left(\tilde{y}_{t+s}^2 + \lambda \tilde{\pi}_{t+s}^2 + \varphi \tilde{i}_{t+s}^2\right)\right]$$

If $\varphi > 0$, then the divine coincidence does not hold. In particular, the larger φ the less the central bank is willing to move interest rates to stabilize the economy. Since prices are fixed, then the private sector's perception of λ does not matter. Instead, suppose the private sector does not know how much is the central bank is willing to move rates to stabilize the economy. That is, they have beliefs μ_t over φ. To form their forecasts, they also assume the central bank optimizes under discretion.

A myopic central bank of type $\varphi = \tilde{\varphi}$ has the following first-order condition

$$\tilde{y}_t = \sigma \tilde{\varphi} \tilde{i}_t$$

Define $\psi^y (\tilde{\varphi}) := \frac{\sigma \tilde{\varphi}}{1+\sigma^2 \tilde{\varphi}}$, taking expectations as exogenous, then the central bank implements

$$\tilde{y}_t (\tilde{\varphi}) = \psi^y (\tilde{\varphi}) \left(\mathbb{E}_t [r_t^n] + \sigma \mathbb{E}_t^P [\tilde{y}_{t+1}]\right)$$

We conjecture that the central bank follows linear policy rule, and they are not aware of the private sector's perceived bias:

$$\tilde{y}_t(\tilde{\varphi}) = \psi(\tilde{\varphi}) \sum_{s=0}^{\infty} \Theta_s \mathbb{E}_t \left[r^n_{t+s} \right]$$

Matching coefficients for $s = 0$

$$\Theta_0 = 1$$

For $s = 1$, using that there is no anticipated learning

$$\Theta_1 = \sigma \mathbb{E}^P_t \left[\psi^y \right]$$

For any arbitrary s, we have

$$\Theta_s = \sigma \mathbb{E}^P_t \left[\psi \right] \Theta_{s-1} = \left(\sigma \mathbb{E}^P_t \left[\psi^y \right] \right)^s$$

Putting everything together,

$$\tilde{y}_t(\tilde{\varphi}) = \psi^y(\tilde{\varphi}) \sum_{s=0}^{\infty} \left(\sigma \mathbb{E}^P_t \left[\psi^y \right] \right)^s \mathbb{E}_t \left[r^n_{t+s} \right] = \psi^y(\tilde{\varphi}) X_t \tag{D.2}$$

$$\tilde{i}_t(\tilde{\varphi}) = (1 - \sigma \psi^y(\tilde{\varphi})) \sum_{s=0}^{\infty} \left(\sigma \mathbb{E}^P_t \left[\psi \right] \right)^s \mathbb{E}_t \left[r^n_{t+s} \right] = \psi^i(\tilde{\varphi}) X_t \tag{D.3}$$

In this model, we define $\mathbb{E}^P_t [\psi^y]$ as the central bank's reputation. When the private sector's perception about φ switches to a smaller value, $\mathbb{E}^P_t [\psi^y]$ decreases and the solution moves closer to the one from the Divine Coincidence. When the perception of φ goes to zero, then we recover $\mathbb{E}^P_t [\psi^y] = 0$ and we recover the Divine coincidence. When $\varphi \to \infty$, then $\mathbb{E}^P_t [\psi^y] \to \frac{1}{\sigma}$

Taylor Rule Under the assumption of Section 5, the $k-$periods ahead forecast of output gap and interest rates is

$$\mathbb{E}^i_t \left[y_{t+k} \right] = \mathbb{E}^P_t \left[\psi^y \right] \mathbb{E}^i_t \left[X_{t+k} \right] \tag{D.4}$$

$$\mathbb{E}^i_t \left[i_{t+k} \right] = \mathbb{E}^P_t \left[\psi^i \right] \mathbb{E}^i_t \left[X_{t+k} \right] \tag{D.5}$$

In this model, the assumption is that there is no disagreement about the future monetary policy surprises. Suppose we follow Bauer et al. (2024) and estimate the perceived Taylor-rule coefficient

$$\mathbb{E}_t^i \left[i_{t+k} \right] = \gamma_{1,t} + \gamma_{2,t} \mathbb{E}_t^i \left[y_{t+k} \right]$$

Then the coefficient becomes

$$\gamma_{1,t} = \frac{1}{\mathbb{E}_t^P \left[\psi^y \right]} - \sigma \geqslant 0$$

which is increasing in the central bank's reputation. Our model provides a structural interpretation of the perceived Taylor rule from Bauer et al. (2024). It is not about how much the central bank prioritizes inflation stability relative to output gap stability, but about how much it is willing to move rates to stabilize the economy. A larger perceived Taylor rule coefficient coincides with an increase in the perceived responsiveness of the central bank to demand-driven fluctuations.

Optimal Policy We keep the private sector's learning structure. The central bank observes a forecast of the natural rate and optimizes under discretion. The final allocation will be $y_t = \tilde{y}_t + \eta_t$ and i_t from D.3. The private sector is unsure whether the final allocation was due to the central bank's preference φ, the monetary policy shock, η_t.

Each period, given μ_t, the private sector updates their prior distribution to their posterior as follows:

1. Given μ_t, the private sector believes that a central bank of type $\varphi = \tilde{\varphi}$ chooses the allocations $\tilde{y}^M(\tilde{\varphi})$ and $i_t(\tilde{\varphi})$ given by (D.2) and (D.3), leading to the ex-post realizations

$$y_t = \tilde{y}^M(\tilde{\varphi}) + \eta_t(\tilde{\varphi})$$

2. The true realization of output gap is given by the allocation chosen by the central bank plus the forecast error

$$y_t = \tilde{y}_t + \eta_t$$

3. Upon observing y_t and i_t, the private sector does not know whether the current realizations are due to the monetary surprise η_t or a (myopic) central bank's preferences $\varphi = \tilde{\varphi}$. Then, from the private sector's point of view, the monetary policy shock of a central bank of type $\varphi = \tilde{\varphi}$ must be

$$\eta_t(\tilde{\varphi}) = \tilde{y}_t - \psi^y(\tilde{\varphi}) X_t + \eta_t$$

Under the same functional assumptions for η_t and beliefs over φ^y, beliefs over φ^y also have a conjugate prior, and the updating of the location parameter is

$$\overline{\psi}_{t+1} = \omega_t X_t^{-1} \left(\tilde{y}_t + \eta_t\right) + \left(1 - \omega_t\right) \overline{\psi}_t \quad \text{where} \quad \omega_t = \frac{X_t^2 \tau_\eta}{\tau_t + X_t^2 \tau_\eta}$$

Then, the same properties will hold. After observing the demand shocks $\left\{\mathbb{E}_t \left[r_{t+s}^n\right]\right\}_{s \geqslant 0}$ the central bank maximizes

$$\mathcal{W}_t = -\frac{1}{2} \mathbb{E}_t \left[\sum_{s=0}^{\infty} \beta^s \left(\tilde{y}_{t+s} + \varphi \tilde{i}_{t+s}^2\right)\right]$$

subject to

(i) The Euler Equation

$$\tilde{y}_t = \mathbb{E}_t^P \left[\tilde{y}_{t+1}\right] - \frac{1}{\sigma} \left(\tilde{i}_t - \mathbb{E}_t \left[r_t^n\right]\right)$$

(ii) The private sector's expectation formation process and learning structure,

$$\begin{aligned} \mathbb{E}_t^P \left[\tilde{y}_{t+1}\right] &= \mathbb{E}_t^P \left[\psi^y\right] \mathbb{E}_t^P \left[X_{t+1}\right] \\ \mathbb{E}_t^P \left[\psi^y\right] &= \overline{\psi}_t - \frac{1}{\sqrt{\tau}} \frac{\phi\left(\tau\left(1 - \overline{\psi}_t\right)\right) - \phi\left(-\tau \overline{\psi}_t\right)}{\Phi\left(\tau\left(1 - \overline{\psi}_t\right)\right) - \Phi\left(-\tau \overline{\psi}_t\right)} \\ \overline{\psi}_{t+1} &= \omega_t X_t^{-1} \left(\tilde{y}_t + \eta_t\right) + \left(1 - \omega_t\right) \overline{\psi}_t \end{aligned}$$

where $\omega_t = \frac{\left(\kappa^{-1} X_t\right)^2 \tau_\eta}{\tau + \left(\kappa^{-1} X_t\right)^2 \tau_\eta}$.

Let μ_t denote the Lagrange multipliers, the first-order conditions are

$$\begin{aligned} \varphi \tilde{i}_t &= \frac{1}{\sigma} \mu_t \\ \tilde{y}_t &= \mu_t - \sum_{s=1}^{\infty} \mathbb{E}_t^{CB} \left[\frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{y}_{t+s+1}\right]}{\partial \tilde{y}_t} \mu_{t+s}\right] \\ \tilde{y}_t &= \mathbb{E}_t^P \left[\tilde{y}_{t+1}\right] - \frac{1}{\sigma} \left(\tilde{i}_t - \mathbb{E}_t \left[r_t^n\right]\right) \end{aligned}$$

Combining the three equations we obtain

$$\tilde{y}_t = \tilde{y}_t^M - \psi^y \sum_{s=1}^{\infty} \mathbb{E}_t^{CB} \left[\frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{y}_{t+s+1}\right]}{\partial \tilde{y}_t} \tilde{i}_{t+s}\right] \tag{D.6}$$

$$\tilde{i}_t = \tilde{i}_t^M + \sigma\psi^y \sum_{s=1}^{\infty} \mathbb{E}_t^{CB}\left[\frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{y}_{t+s+1}\right]}{\partial \tilde{y}_t}\tilde{i}_{t+s}\right] \tag{D.7}$$

The insurance term in this case is given by

$$\mathcal{I}_t = \sum_{s=1}^{\infty} \mathbb{E}_t^{CB}\left[\frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{y}_{t+s+1}\right]}{\partial \tilde{y}_t}\tilde{i}_{t+s}\right]$$

Notice that this model is isomorphic to a model where the central bank tries to learn about the relative weight given to inflation. Thus, the same principles will hold. The main difference is that the central bank will now prioritize more output gap stability relative to the myopic benchmark. However, the policy prescriptions for the interest rate remains unchanged.

C.2 Uncertainty about Inflation Target

Suppose that the central bank's objective function is given by

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{s=0}^{\infty} \beta^s \left(\tilde{y}_{t+s}^2 + \lambda\left(\pi_{t+s} - \varphi\right)^2\right)\right] \tag{D.8}$$

In contrast with the main model, the private sector is uncertain about the central bank's inflation target φ. We also assume the private sector is aware of the value of λ. Like our main model, the price sector believes the central bank is myopic, and we maintain the information structure. A myopic central bank with inflation target $\varphi = \tilde{\varphi}$ has the following first-order condition

$$\tilde{y}_t + \lambda\kappa\left(\tilde{\pi}_t - \varphi\right) = 0$$

Define $\psi^\pi = \frac{1}{1+\lambda\kappa^2}$ and $\psi^y = \frac{\lambda\kappa}{1+\lambda\kappa^2}$. Taking expectations as exogenous, the central bank implements

$$\begin{aligned}\tilde{y}_t\left(\tilde{\varphi}\right) =& -\psi^y\left(\varepsilon_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] - \frac{\tilde{\varphi}}{\lambda\kappa}\right)\\ \tilde{\pi}_t =& \psi^\pi\left(\varepsilon_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+s}\right] + \kappa\tilde{\varphi}\right)\end{aligned}$$

We conjecture that the central bank follows a linear policy rule, and they are not aware of the private sector's perceived bias:

$$\begin{aligned}\tilde{y}_t(\tilde{\varphi}) =& \Psi_1^y \tilde{\varphi} - \Psi_2^y \mathbb{E}_t^P[\varphi] - \Psi_3^y \sum_{s=0}^{\infty} \Theta_s^y \mathbb{E}_t[\varepsilon_{t+s}] \\ \tilde{\pi}_t(\tilde{\varphi}) =& \Psi_1^\pi \tilde{\varphi} + \Psi_2^\pi \mathbb{E}_t^P[\varphi] + \Psi_3^\pi \sum_{s=0}^{\infty} \Theta_s^\pi \mathbb{E}_t[\varepsilon_{t+s}]\end{aligned}$$

Then

$$\Psi_1^y = \psi^\pi \qquad \Psi_2^y = \frac{\kappa\beta\psi^\pi}{1-\beta\psi^\pi}\psi^y \qquad \Psi_3^y = \psi^y$$

$$\Psi_1^\pi = \kappa\psi^\pi \qquad \Psi_2^\pi = \frac{\kappa\beta\psi^\pi}{1-\beta\psi^\pi}\psi^\pi \qquad \Psi_3^\pi = \psi^\pi$$

Finally, we match the coefficients for $\{\Theta_s^y, \Theta_s^\pi\}_{s\geqslant 0}$. Start with $s = 0$

$$\Theta_0^y = \Theta_0^\pi = 1$$

For $s = 1$

$$\Theta_1^y = \beta\psi^\pi \qquad \Theta_1^\pi = \beta\psi^\pi$$

For an arbitrary s, we have

$$\Theta_s = \beta\psi^\pi\Theta_{s-1} = (\beta\psi^\pi)^s$$

Putting everything together,

$$\tilde{y}_t(\tilde{\varphi}) = \Psi_1^y \tilde{\varphi} - \Psi_2^y \mathbb{E}_t^P[\varphi] - \Psi_3^y \sum_{s=0}^{\infty} (\beta\psi^\pi)\, \mathbb{E}_t[\varepsilon_{t+s}] = \Psi_1^y \tilde{\varphi} - \Psi_2^y \mathbb{E}_t^P[\varphi] - \Psi_3^y X_t \tag{D.9}$$

$$\tilde{\pi}_t(\tilde{\varphi}) = \Psi_1^\pi \tilde{\varphi} + \Psi_2^\pi \mathbb{E}_t^P[\varphi] - \Psi_3^\pi \sum_{s=0}^{\infty} (\beta\psi^\pi)\, \mathbb{E}_t[\varepsilon_{t+s}] = \Psi_1^\pi \tilde{\varphi} + \Psi_2^\pi \mathbb{E}_t^P[\varphi] + \Psi_3^\pi X_t \tag{D.10}$$

The private sector's perception about the inflation target acts as a cost-push shock: it raises inflation expectations and decreases output. In contrast to the case where λ is uncertain, it does not depend on the shocks: according to the private sector, the central bank will try to boost the output gap to inflate the economy. Therefore, they will raise their inflation expectations. In this model we can interpret $\mathbb{E}_t^P[\varphi]$ as the reputation of the central bank: a good reputation is tied to a low inflation target.

Upon observing the final allocations π_t and y_t, the private sector does not know whether the current realizations are due to the forecast error η_t or the myopic central bank's inflation target $\tilde{\varphi}$. From the private sector's point of view, the first-order condition of a myopic central bank with inflation target $\varphi = \tilde{\varphi}$ is

$$\tilde{y}_t(\tilde{\varphi}) + \lambda\kappa(\tilde{\pi}_t(\tilde{\varphi}) - \tilde{\varphi}) = 0$$

Then, the forecast error of a central bank with inflation bias $\varphi = \tilde{\varphi}$ must be

$$\eta_t(\tilde{\varphi}) = \psi^y\left(\tilde{y}_t + \lambda\kappa(\tilde{\pi}_t - \tilde{\varphi}) + \left(1 + \lambda\kappa^2\right)\eta_t\right)$$

We make assume that the forecast error is normally distributed, i.e., $\eta_t \sim \mathcal{N}\left(0, \tau_\eta^{-1}\right)$ for all t; and, the prior belief about φ is also normally distributed, i.e., $\tilde{\varphi} \sim \mathcal{N}\left(\mathbb{E}_t^P[\varphi], \tau_t^{-1}\right)$. Under these assumptions, beliefs have a conjugate prior, and

$$\mathbb{E}_{t+1}^P[\varphi] = \omega_t(\lambda\kappa)^{-1}\left(\tilde{y}_t + \lambda\kappa\tilde{\pi}_t + \left(1 + \lambda\kappa^2\right)\eta_t\right) + (1 - \omega_t)\,\mathbb{E}_t^P[\varphi]$$

where

$$\omega_t = \frac{(\psi^\pi)^2\tau_\eta}{(\psi^\pi)^2\tau_\eta + \tau_t}$$

In contrast to our main model, the attention weight ω_t does not depend on the size of the shocks. To understand the implications of the learning process, rewrite this expression as

$$\mathbb{E}_{t+1}^P[\varphi] = \mathbb{E}_t^P[\varphi] + \omega_t\left(\tilde{y}_t + \lambda\kappa\left(\tilde{\pi}_t - \mathbb{E}_t^P[\varphi]\right) + \left(1 + \lambda\kappa^2\right)\eta_t\right)$$

Then, the only way to improve the central bank can improve its reputation is by acting more hawkish than expected. This way, the central bank signals a lower inflation target and reduces inflation expectations in the following period. As in our main model, we assume the Kalman gain is constant, so $\omega_t = \omega$ for all t. After observing the cost push shocks $\{\mathbb{E}_t[\varepsilon_{t+s}]\}_{s\geqslant 0}$ the central bank maximizes

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s}^2 + \lambda(\tilde{\pi}_{t+s} - \varphi)^2\right)\right]$$

subject to

(i) The New Keynesian Phillips Curve

$$\tilde{\pi}_t = \kappa \tilde{y}_t + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1} \right] + \varepsilon_t$$

(ii) The private sector's learning structure,

$$\begin{aligned} \mathbb{E}_t^P \left[\tilde{\pi}_{t+1} \right] &= \left(\Psi_1^\pi + \Psi_2^\pi \right) \mathbb{E}_t^P \left[\varphi \right] + \Psi_3^\pi \mathbb{E}_t \left[X_{t+1} \right] \\ \mathbb{E}_{t+1}^P \left[\varphi \right] &= \omega_t \left(\lambda \kappa \right)^{-1} \left(\tilde{y}_t + \lambda \kappa \tilde{\pi}_t + \left(1 + \lambda \kappa^2 \right) \eta_t \right) + \left(1 - \omega_t \right) \mathbb{E}_t^P \left[\varphi \right] \end{aligned}$$

where $\omega_t = \frac{(\psi^\pi)^2 \tau_\eta}{(\psi^\pi)^2 \tau_\eta + \tau_t}$.

Let μ_t denote the Lagrange multipliers, the first-order conditions are

$$\begin{aligned} \tilde{y}_t &= \kappa \mu_t + \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{y}_t} \mu_{t+s} \right] \\ \lambda \left(\tilde{\pi}_t - \varphi \right) &= -\mu_t + \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{\pi}_t} \mu_{t+s} \right] \\ \tilde{\pi}_t &= \kappa \tilde{y}_t + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1} \right] + \varepsilon_t \end{aligned}$$

Combining the first and second equation

$$\tilde{y}_t + \lambda \kappa \left(\tilde{\pi}_t - \varphi \right) = \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \left(\frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{y}_t} + \kappa \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{\pi}_t} \right) \mu_{t+s} \right]$$

Using the private sector's learning structure

$$\frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{y}_t} + \kappa \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{\pi}_t} = \frac{\omega}{\psi^y} \left(1 - \omega \right)^{s-1}$$

Combining with the New Keynesian Phillips Curve we have

$$\begin{aligned} \tilde{y}_t &= \tilde{y}_t^M + \frac{\omega}{1-\omega} \sum_{s=1}^{\infty} \left(\beta \left(1 - \omega \right) \right)^s \mathbb{E}_t^{CB} \left[\mu_{t+s} \right] \\ \tilde{\pi}_t &= \tilde{\pi}_t^M + \kappa \frac{\omega}{1-\omega} \sum_{s=1}^{\infty} \left(\beta \left(1 - \omega \right) \right)^s \mathbb{E}_t^{CB} \left[\mu_{t+s} \right] \end{aligned}$$

The unique equilibrium is linear, and given by

$$\tilde{y}_t = \Theta_1^y \varphi - \Theta_2^y \mathbb{E}_t^P [\varphi] - \sum_{s=0}^{\infty} \Theta_{3,s}^y \mathbb{E}_t [\varepsilon_{t+s}] \tag{D.11}$$

$$\tilde{\pi}_t = \Theta_1^\pi \varphi + \Theta_2^\pi \mathbb{E}_t^P [\varphi] + \sum_{s=0}^{\infty} \Theta_{3,s}^\pi \mathbb{E}_t [\varepsilon_{t+s}] \tag{D.12}$$

How does the optimal policy compare to the myopic central bank? First, the inflation target has a smaller influence on the optimal policy, that is, $\Theta_1^y < \Psi_1^y$ and $\Theta_1^\pi < \Psi_1^\pi$. The central bank internalizes that raising output will raise inflation expectations.

Second, the optimal policy for output gap (inflation) is to overreact (underreact) to the central bank's reputation $\mathbb{E}_t^P [\varphi]$, that is, $\Theta_2^y > \Psi_2^y$ and $\Theta_2^\pi < \Psi_2^\pi$. When the private sector learns about the inflation target, inflation expectations are a cost-push shock. The optimal response to a cost-push shock is to induce a recession to mitigate its impact on inflation. For the case of the perceived inflation target, it is a shock that is endogenous to policy. By overreacting, the central bank reduces the future recession's size.

Third, the optimal policy for current shocks is to react exactly like the myopic central bank, that is, $\Theta_{3,0}^y = \Psi_3^y$ and $\Theta_{3,0}^\pi = \Psi_3^\pi$. This is no longer true when considering the reaction to persistent shocks. When shocks are persistent, a current shock is also informative about a future recession. To smooth the size of the future recession, the central bank overreacts to improve its reputation. Therefore, we have $\Theta_{3,s}^y > \Psi_{3,s}^y$ and $\Theta_{3,1}^\pi < \Psi_{3,s}^\pi$ for $s \geqslant 1$.

As in the main model, the central bank overreacts to persistent shocks. However, there is no insurance. There is no incentive to improve reputation in response to an *iid* shock. In response to a persistent shock, the central bank smooths the size of the recession. In response, the central bank trades-off some of the future recession by inducing a current recession. In the main model, this mechanism holds regardless of the nature of the shock. For that reason, there is no overreaction in response to contemporary cost-push shocks. In this sense, when the central bank is concerned about the private sector's perception of the inflation target, reputation and stabilization are independent from each other. This is no longer true in the main model: there, the role of reputation is precisely to reduce the cost of stabilization.

C.3 Markov Perfect Equilibrium

So far, we assumed for tractability that the private sector held misspecified beliefs. They believe the central bank does not internalize their belief updating process. In reality, financial institutions are aware that the central bank is trying to manipulate its beliefs. To add this dimension, we assume the private sector is fully aware that the central bank internalizes their learning process. For simplicity, we assume the central bank has no commitment. The learning process is still the same as (12), but we do not place any functional form assumptions . The first-order condition of a central bank of type $\tilde{\psi}^{\pi}$ is:

$$\tilde{y}_t\left(\tilde{\psi^{\pi}}\right) = \tilde{y}_t^M\left(\tilde{\psi^{\pi}}\right) + \tilde{\psi}^{\pi}\sum_{s=1}^{\infty}\beta^s\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^P\left[\pi_{t+1+s}\left(\psi^{\pi}\right)\right]}{\partial\tilde{\pi}_t\left(\tilde{\psi^{\pi}}\right)}\tilde{y}_{t+s}\left(\tilde{\psi^{\pi}}\right)\right]$$

where

$$\begin{aligned}\tilde{y}_t^M\left(\tilde{\psi^{\pi}}\right) &= -\,\psi^y\left(\tilde{\psi^{\pi}}\right)\left(\varepsilon_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\left(\tilde{\psi^{\pi}}\right)\right]\right)\\ \mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\left(\tilde{\psi^{\pi}}\right)\right] &= \int_0^1\mathbb{E}_t\left[\tilde{\pi}_{t+1}\left(\tilde{\psi^{\pi}}\right)\right]\mu_t\left(\tilde{\psi^{\pi}}\right)d\tilde{\psi}^{\pi}\end{aligned}$$

Current inflation expectations, $\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\left(\psi^{\pi}\right)\right]$, are endogenous, which means that we cannot interpret $\tilde{y}_t^M$ as the equilibrium allocation implemented by a myopic central bank. However, $\tilde{y}_t^M$ is the *temporary* equilibrium allocation (Hicks (1946), García-Schmidt and Woodford (2019)). That is, the equilibrium allocation taking expectations of the next period $\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right]$ as given. We can define overreaction and underreaction with respect to this new benchmark, which is the rational expectations analog of a myopic central bank.

Since there is a continuum of states $\psi^{\pi} \in (0,1)$, the whole distribution of prior beliefs is a state variable, making the problem intractable. One possibility to make the problem tractable is to assume types are finite, as in Bocola et al. (2025). However, this would eliminate the extensive margin of reputation we consider in our analysis. We take a different route, and consider the three-period version of our model from Section 2. Given any initial distribution of beliefs at $t = 0$, $\mathbb{E}_1^P\left[\tilde{\pi}_2\right] = \mathbb{E}_1^P\left[\psi^{\pi}\right]\mathbb{E}_1\left[\varepsilon_2\right]$. Therefore, all that matters for optimal policy is the central bank's reputation in $t = 1$, $\mathbb{E}_1^P\left[\psi^{\pi}\right]$. Using backward induction, we have

$$\tilde{y}_0\left(\tilde{\psi}^\pi\right) = \tilde{y}_0^M\left(\tilde{\psi}^\pi\right) + \tilde{\psi}^\pi \beta \mathbb{E}_t^{CB}\left[\beta \frac{\partial \mathbb{E}_1^P\left[\tilde{\pi}_2\right]}{\partial \tilde{\pi}_0\left(\tilde{\psi}^\pi\right)} \tilde{y}_1\left(\tilde{\psi}^\pi\right)\right]$$
$$\tilde{\pi}_0\left(\tilde{\psi}^\pi\right) = \tilde{\pi}_0^M\left(\tilde{\psi}^\pi\right) + \kappa\tilde{\psi}^\pi \beta \mathbb{E}_0^{CB}\left[\beta \frac{\partial \mathbb{E}_t^P\left[\tilde{\pi}_2\right]}{\partial \tilde{\pi}_0\left(\tilde{\psi}^\pi\right)} \tilde{y}_1\left(\tilde{\psi}^\pi\right)\right]$$

Let $\mu_0\left(\psi^\pi; \tilde{\psi}^\pi\right)$ denote the prior density of the central bank of type ψ^π, from the point of view of a central bank of type $\tilde{\psi}^\pi$. The posterior density is characterized by

$$\mu_1\left(\psi^\pi; \tilde{\psi}^\pi\right) = \frac{f_\eta\left(\kappa^{-1}\left(\tilde{\pi}_0\left(\psi^\pi\right) - \tilde{\pi}_0\left(\tilde{\psi}^\pi\right)\right) + \eta_0\right)\mu_0\left(\psi^\pi; \tilde{\psi}^\pi\right)}{\int_0^1 f_\eta\left(\kappa^{-1}\left(\tilde{\pi}_0\left(\psi^\pi\right) - \tilde{\pi}_0\left(\tilde{\psi}^\pi\right)\right) + \eta_0\right)\mu_0\left(\psi^\pi; \tilde{\psi}^\pi\right) d\psi^\pi}$$
$$= \frac{f_\eta\left(\psi^\pi; \tilde{\psi}^\pi, \eta_0\right)\mu_0\left(\psi^\pi; \tilde{\psi}^\pi\right)}{\int_0^1 f_\eta\left(\psi^\pi; \tilde{\psi}^\pi, \eta_0\right)\mu_0\left(\psi^\pi; \tilde{\psi}^\pi\right) d\psi^\pi}$$

Then, we have

$$\frac{\partial \mathbb{E}_1^P\left[\psi^\pi\right]}{\partial \tilde{\pi}_0\left(\tilde{\psi}^\pi\right)} = \kappa^{-1}\mathbb{COV}_1^P\left[\psi^\pi, s_\eta\left(\psi^\pi; \tilde{\psi}^\pi, \eta_0\right)\right]$$

where $s_\eta\left(\psi^\pi; \tilde{\psi}^\pi, \eta_0\right)$ is the score of the likelihood of a central bank of type ψ^π from the point of view of central bank of type $\tilde{\psi}^\pi$. This expression formalizes that the central bank's reputation is more sensitive when its actions are more informative about its type. The score is the sensitivity of the likelihood of a type ψ^π to the central bank's actions. The sensitivity of expectations averages over all the hypothetical types ψ^π. This expression holds regardless of the functional form assumptions.

For any value of η_0, the score is nondecreasing in ψ^π. As a result, the covariance will be positive. Since $\mathbb{E}_1\left[\varepsilon_2\right]\tilde{y}_1 < 0$, then the insurance principle from 1 still holds, and there will be overreaction in the output gap and underreaction in inflation. Overreaction will be smaller compared to the case where the private sector believes the central bank is myopic: since the private sector internalizes that the central bank would like to increase its reputation, overreaction becomes less informative about the central bank's type. As a result, overreaction is smaller.

Figure 19a plots the overreaction as a function of the central bank's type, ψ^π. The magnitude of overreaction is u-shaped: reputation allows the central bank to improve its inflation–output trade-off, whose importance increases when the central bank's preferences are more balanced. Figure 19b displays overreaction of a central bank with $\lambda = 10$ as a function of reputation. Consistent with our main model, overreaction is also u-shaped as a function of reputation. Finally, Figure 19c compares the extent of overreaction across priors. When the private sector's prior is uniform, uncertainty is maximal—every type is equally likely—and the central bank overreacts more than under the truncated normal prior with the same mean.

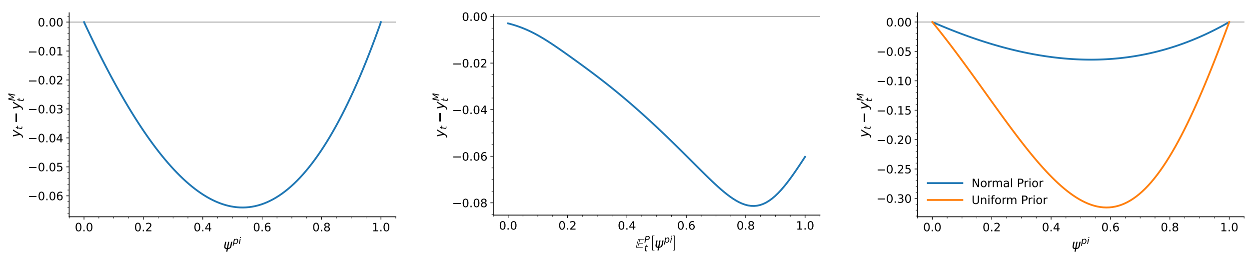

(a) Overreaction and ψ^{π} **(b)** Overreaction and $\mathbb{E}_t^P[\psi^{\pi}]$ **(c)** Normal vs Uniform priors

Figure 19: Overreaction in the Markov equilibrium

Figure 20 then plots the dynamics of reputation at $t = 1$. The pattern mirrors the main model: reputation improves when the central bank is perceived as dovish, and deteriorates when it is perceived as hawkish.

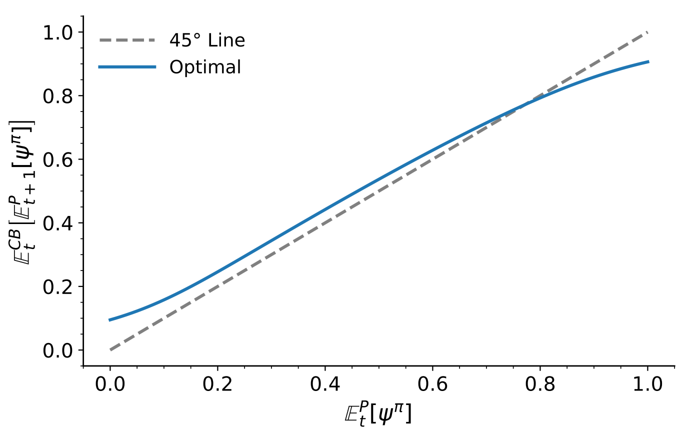

Figure 20: Reputation dynamics in the Markov equilibrium

Taken together, these figures show that the qualitative results of our economy do not change if we consider a Markov perfect equilibrium. Overreaction remains a robust feature, its magnitude varies with type and reputation in a u-shaped fashion, and the dynamics of credibility follow the same logic as in the baseline environment. Even though the equilibrium in this finite-horizon economy is unique, we cannot rule out the possibility of multiple equilibria in the infinite-horizon case. If the central bank's type λ changes over time, then the posterior distribution does not converge to a singleton around its true value, but rather to a stationary distribution. In that setting, if private sector beliefs are sufficiently rigid, there may exist a steady state with hawkish reputation and another with dovish reputation. When reputation is dovish, the central bank lacks strong incentives to improve it; when it is hawkish, the cost of sustaining credibility is not high enough to induce deteriorating it. Multiple equilibria can therefore arise. In this sense, our assumption that the private sector believes the central bank is myopic can be viewed as an equilibrium refinement that restores uniqueness.

C.4 Contemporaneous Belief Updating

C.5 Zero Lower Bound

If the ZLB binds then

$$\tilde{y}_t + \lambda\kappa\tilde{\pi}_t = -\mu_t$$

Where μ_t is the shadow value of decreasing interest rates below 0. Not being able to decrease rates below zero leads to recession and, through the NKPC, deflation. In this environment, a dovish reputation can allow the central bank to get out of the ZLB. Recall the IS curve is given by

$$\tilde{y}_t = \mathbb{E}_t^P\left[\tilde{y}_{t+1}\right] - \frac{1}{\sigma}\left(\mathbb{E}_t^{CB}\left[\tilde{r}_t^n\right] - \mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right]\right)$$

Using (9) and (10)

$$\tilde{y}_t = \left(\frac{1}{\sigma}\mathbb{E}_t^P\left[\psi^\pi\right] - \mathbb{E}_t^P\left[\psi^y\right]\right)\mathbb{E}_t^P\left[X_{t+1}\right] - \frac{1}{\sigma}\mathbb{E}_t^{CB}\left[\tilde{r}_t\right]$$

Define $\Psi := \mathbb{E}_t^P\left[\psi^\pi\right] - \sigma\mathbb{E}_t^P\left[\psi^y\right])$, it is increasing in $\mathbb{E}_t^P\left[\psi^\pi\right]$. If the central bank has sufficiently dovish reputation then $\Psi > 0$, which can help the central bank get out of the ZLB in response to expansionary shocks. This echoes Krugman (1998), which proposes central banks to "credibly promisse to be irresponsible".

D Proofs and Robustness Checks for Empirical Section

D.1 Proof of Proposition 6

Proof. The estimand of $\gamma_{2,t}$ is

$$\gamma_{2,t} = \frac{Cov\left(\mathbb{E}_t^i\left[y_{t+k}\right], \mathbb{E}_t^i\left[\pi_{t+k}\right]\right)}{Var\left(\mathbb{E}_t^i\left[\pi_{t+k}\right]\right)} = -\frac{\mathbb{E}_t^P\left[\psi^y\right]\mathbb{E}_t^P\left[\psi^\pi\right]Var\left(\mathbb{E}_t^i\left[X_{t+k}\right]\right)}{\mathbb{E}_t^P\left[\psi^\pi\right]^2 Var\left(\mathbb{E}_t^i\left[X_{t+k}\right]\right)} = -\frac{\mathbb{E}_t^P\left[\psi^y\right]}{\mathbb{E}_t^P\left[\psi^\pi\right]} \tag{B.1}$$

Finally, from the definition of ψ^π and ψ^y we know that $\psi^y = \kappa^{-1}\left(1-\psi^\pi\right)$. Replacing into B.1 completes the proof. ■

D.2 Proof of Proposition 5

Proof. The first-order condition of a central bank with parameters λ and φ is

$$\tilde{y}_t + \lambda\kappa\tilde{\pi}_t = \varphi i_t$$

Define $\widehat{\psi}^y := \frac{\lambda\kappa}{1+\lambda\kappa^2+\sigma\varphi}$ and $\widehat{\psi}^\pi := \frac{1}{1+\lambda\kappa^2+\sigma\varphi}$. Taking expectations as given, the central bank implements

$$\begin{aligned}
\tilde{y}_t\left(\tilde{\lambda},\tilde{\varphi}\right) =& \widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)\left(r_t^n + \sigma\mathbb{E}_t^P\left[\tilde{y}_{t+1}\right] + \mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right]\right) - \widehat{\psi}^y\left(\tilde{\lambda},\tilde{\varphi}\right)\beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] \\
\tilde{\pi}_t\left(\tilde{\lambda},\tilde{\varphi}\right) =& \kappa\widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)\left(r_t^n + \sigma\mathbb{E}_t^P\left[\tilde{y}_{t+1}\right] + \mathbb{E}_t\left[\tilde{\pi}_{t+1}\right]\right) + \widehat{\psi}^\pi\left(\tilde{\lambda},\tilde{\varphi}\right)\beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right]
\end{aligned}$$

We conjecture that the Central Bank follows a linear policy rule

$$\begin{aligned}
\tilde{y}_t\left(\tilde{\lambda},\tilde{\varphi}\right) &= \sum_{s=0}^{\infty}\Theta_s^y\left(\tilde{\lambda},\tilde{\varphi}\right)\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right] \\
\tilde{\pi}_t\left(\tilde{\lambda},\tilde{\varphi}\right) &= \sum_{s=0}^{\infty}\Theta_s^\pi\left(\tilde{\lambda},\tilde{\varphi}\right)\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right]
\end{aligned}$$

Matching coefficients for $k=0$

$$\Theta_0^y\left(\tilde{\lambda},\tilde{\varphi}\right) = \widehat{\psi}\left(\tilde{\lambda}\right) \qquad \Theta_0^\pi\left(\tilde{\lambda},\tilde{\varphi}\right) = \kappa\Theta_0^y\left(\tilde{\lambda},\tilde{\varphi}\right)$$

For $k = 1$

$$\Theta_1^y\left(\tilde{\lambda},\tilde{\varphi}\right) = \widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)(\sigma+\kappa)\,\mathbb{E}_t^P\left[\widehat{\psi}\right] - \kappa\beta\widehat{\psi}^y\left(\tilde{\lambda},\tilde{\varphi}\right)\mathbb{E}_t^P\left[\widehat{\psi}\right]$$
$$\Theta_1^\pi\left(\tilde{\lambda},\tilde{\varphi}\right) = \kappa\widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)(\sigma+\kappa)\,\mathbb{E}_t^P\left[\widehat{\psi}\right] + \kappa\beta\widehat{\psi}^\pi\left(\tilde{\lambda},\tilde{\varphi}\right)\mathbb{E}_t^P\left[\widehat{\psi}\right]$$

For an arbitrary s, we have

$$\Theta_s^y\left(\tilde{\lambda},\tilde{\varphi}\right) = \widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)\left(\sigma\mathbb{E}_t^P\left[\Theta_{s-1}^y\right] + \mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]\right) - \widehat{\psi}^y\left(\tilde{\lambda},\tilde{\varphi}\right)\beta\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right] \tag{B.2}$$
$$\Theta_s^\pi\left(\tilde{\lambda},\tilde{\varphi}\right) = \kappa\widehat{\psi}\left(\tilde{\lambda},\tilde{\varphi}\right)\left(\sigma\mathbb{E}_t^P\left[\Theta_{s-1}^y\right] + \mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]\right) + \widehat{\psi}^\pi\left(\tilde{\lambda},\tilde{\varphi}\right)\beta\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right] \tag{B.3}$$

Then, k-period ahead forecast is given by

$$\mathbb{E}_t^i\left[y_{t+k}\right] = \sum_{s=0}^{\infty}\mathbb{E}_t^P\left[\Theta_s^y\right]\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right] + \mathbb{E}_t^P\left[\eta_{t+k}\right]$$
$$\mathbb{E}_t^i\left[\pi_{t+k}\right] = \sum_{s=0}^{\infty}\mathbb{E}_t^P\left[\Theta_s^\pi\right]\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right] + \kappa\mathbb{E}_t^P\left[\eta_{t+k}\right]$$

The estimand of $\gamma_{2,t}$ from equation (22) is

$$\gamma_{2,t} = \frac{\sum_{s=0}^{\infty}\mathbb{E}_t^P\left[\Theta_s^\pi\right]\mathbb{E}_t^P\left[\Theta_s^y\right]Var\left(\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right]\right)}{\sum_{s=0}^{\infty}\mathbb{E}_t^P\left[\Theta_s^\pi\right]\mathbb{E}_t^P\left[\Theta_s^\pi\right]Var\left(\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}^n\right]\right)}$$

Assume that $Var\left(\mathbb{E}_t^i\left[\tilde{r}_{t+k+s}\right]\right) = \rho^s Var\left(\mathbb{E}_t^i\left[\tilde{r}_{t+k}\right]\right)$ with $\rho < 1$. Implicitly, we assume that the private sector reaches consensus about the long run. Then, we can rewrite the estimand as

$$\gamma_{2,t} = \sum_{s=0}^{\infty}\omega_s\frac{\mathbb{E}_t^P\left[\Theta_s^y\right]}{\mathbb{E}_t^P\left[\Theta_s^\pi\right]} \qquad \omega_s = \frac{\mathbb{E}_t^P\left[\Theta_s^\pi\right]^2}{\sum_{s=0}^{\infty}\mathbb{E}_t^P\left[\Theta_s^\pi\right]^2}$$

Now we study the sign of $\gamma_{2,t}$ by studying each one of the terms in the sum. First,

$$\frac{\mathbb{E}_t^P\left[\Theta_0^y\right]}{\mathbb{E}_t^P\left[\Theta_0^\pi\right]} = \frac{1}{\kappa} > 0$$

which does not depend on the beliefs. This is the result of the contemporaneous effect of a demand shock that is not fully stabilized: an unit increase of output gap directly implies a contemporaneous increase in inflation of κ. For $k = 1$,

$$\frac{\mathbb{E}_t^P\left[\Theta_1^y\right]}{\mathbb{E}_t^P\left[\Theta_1^\pi\right]} = \frac{\left(1+\sigma\frac{\mathbb{E}_t^P\left[\Theta_{s-1}^y\right]}{\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]}\right) - \beta\frac{\mathbb{E}_t^P\left[\widehat{\psi}^y\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]}}{\kappa\left(1+\sigma\frac{\mathbb{E}_t^P\left[\Theta_{s-1}^y\right]}{\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]}\right) + \beta\frac{\mathbb{E}_t^P\left[\widehat{\psi}^\pi\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]}}$$

Observe that

$$\lim_{\lambda\to 0}\frac{\mathbb{E}_t^P\left[\widehat{\psi}^y\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]} = 0 \quad \lim_{\lambda\to \infty}\frac{\mathbb{E}_t^P\left[\widehat{\psi}^y\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]} \to \infty \quad \lim_{\lambda\to 0}\frac{\mathbb{E}_t^P\left[\widehat{\psi}^\pi\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]} > 0 \quad \lim_{\lambda\to \infty}\frac{\mathbb{E}_t^P\left[\widehat{\psi}^\pi\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]} > 0$$

Then, we have

$$\lim_{\lambda\to 0}\frac{\mathbb{E}_t^P\left[\Theta_1^y\right]}{\mathbb{E}_t^P\left[\Theta_1^\pi\right]} > \frac{\mathbb{E}_t^P\left[\Theta_0^y\right]}{\mathbb{E}_t^P\left[\Theta_0^\pi\right]} \qquad \lim_{\lambda\to \infty}\frac{\mathbb{E}_t^P\left[\Theta_1^y\right]}{\mathbb{E}_t^P\left[\Theta_1^\pi\right]} \to -\infty$$

With anticipated demand shocks there are two opposite effects. First, the expansionary effect of a demand shock that is not fully stabilized. Second, the increase in inflation expectations acts as a cost-push shock for inflation. When the Central Bank is dovish, the first effect dominates. When the Central Bank is hawkish, the second effect dominates.

This logic also holds for an arbitrary horizon s. To see this, use (B.2) and (B.3)

$$\frac{\mathbb{E}_t^P\left[\Theta_s^y\right]}{\mathbb{E}_t^P\left[\Theta_s^\pi\right]} = \frac{\left(1+\sigma\frac{\mathbb{E}_t^P\left[\Theta_{s-1}^y\right]}{\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]}\right) - \beta\frac{\mathbb{E}_t^P\left[\widehat{\psi}^y\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]}}{\kappa\left(1+\sigma\frac{\mathbb{E}_t^P\left[\Theta_{s-1}^y\right]}{\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]}\right) + \beta\frac{\mathbb{E}_t^P\left[\widehat{\psi}^\pi\right]}{\mathbb{E}_t^P\left[\widehat{\psi}\right]}}$$

When $\lambda \to 0$ we have

$$\lim_{\lambda\to 0}\frac{\mathbb{E}_t^P\left[\Theta_s^y\right]}{\mathbb{E}_t^P\left[\Theta_s^\pi\right]} > \lim_{\lambda\to 0}\frac{\mathbb{E}_t^P\left[\Theta_{s-1}^y\right]}{\mathbb{E}_t^P\left[\Theta_{s-1}^\pi\right]}$$

and for $\lambda \to \infty$ we have

$$\lim_{\lambda\to\infty}\frac{\mathbb{E}_t^P\left[\Theta_s^y\right]}{\mathbb{E}_t^P\left[\Theta_s^\pi\right]} = \begin{cases} \frac{1}{\kappa}\left(1+\frac{\beta}{\sigma}\right) & s \in \{2, 4, \ldots, 2n\} \\ -\infty & s \in \{1, 3, \ldots, 2n+1\} \end{cases}$$

Then, we conclude that $\gamma_{2t} > 0$ as $\lambda \to 0$, and $\gamma_{2t} \to -\infty$ as $\lambda \to \infty$. This completes the proof. ■

D.3 Robustness to the Fed Information Effect

For robustness, we follow Bauer and Swanson (2022) and Bauer and Swanson (2023), and orthogonalize the shocks with respect to the public information that became available between FOMC meetings. In particular, we orthogonalize the shocks using the same six variables in Bauer and Swanson (2022): Nonfarm payrolls surprise, employment growth, change in the S&P 500, change in the slope of the Yield curve, change in Commodity prices, and implied skewness of the ten-year Treasury yield.

Figure 21 plots the impulse response to an orthogonalized monetary tightening. The effect prevails.

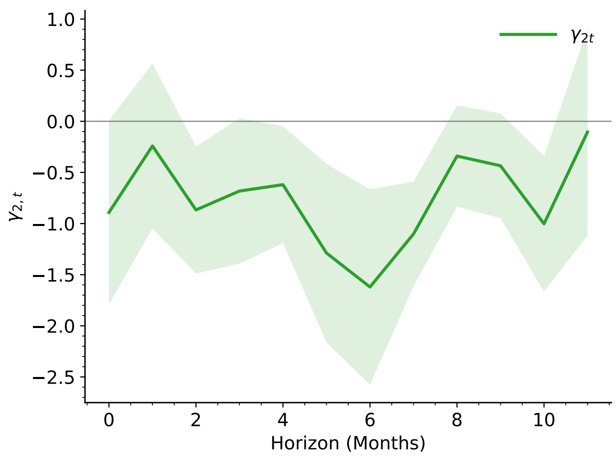

Figure 21: Impulse response to an orthogonalized monetary tightening.
Shaded areas denote 68% confidence intervals.

D.4 Proof of Proposition 6

Proof. For simplicity, suppose there is no disagreement about the forecast of demand. Taking expectations as exogenous, the Central Bank implements

$$\begin{aligned}\tilde{y}_t =& \psi^\pi y_{t-1} - \psi^y \left(\varepsilon_t + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1}\right]\right) \\ \tilde{\pi}_t =& \kappa \psi^\pi y_{t-1} + \psi^\pi \left(\varepsilon_t + \beta \mathbb{E}_t^P \left[\tilde{\pi}_{t+1}\right]\right)\end{aligned}$$

Conjecture the Central Bank follows a linear rule

$$\begin{aligned}\tilde{y}_t =& \varphi^y y_{t-1} - \alpha\psi^y \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\ \tilde{\pi}_t =& \varphi^\pi y_{t-1} + \alpha\psi^\pi \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

Replacing into both equations

$$\begin{aligned}&\varphi^\pi y_{t-1} + \alpha\psi^\pi \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\ &\qquad = \kappa\psi^\pi y_{t-1} + \psi^\pi \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] + \psi^\pi\beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right] y_{t-1} - \alpha\psi^\pi\beta\mathbb{E}_t^P\left[\varphi^\pi\psi^y\right] \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right) \mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

$$\begin{aligned}&\varphi^y y_{t-1} - \alpha\psi^y \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\ &\qquad = \psi^\pi y_{t-1} - \psi^y \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] - \psi^y\beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right] y_{t-1} + \alpha\psi^y\beta\mathbb{E}_t^P\left[\varphi^\pi\psi^y\right] \sum_{s=0}^{\infty} \left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s \mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

Matching coefficients yields

$$\begin{aligned}\alpha =& \frac{1}{1+\beta\mathbb{E}^P\left[\varphi^\pi\psi^y\right]} \\ \varphi^\pi =& \left(\kappa + \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]\right)\psi^\pi \\ \varphi^y =& \left(1 - \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]\right)\psi^y\end{aligned}$$

Notice that $\alpha > 1$. From the equations for φ^π and φ^y, $\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]$ is pinned down by

$$\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right] = \left(\kappa + \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]\right)\left(1 - \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]\right)\mathbb{E}_t^P\left[\psi^\pi\psi^y\right]$$

We conjecture φ^π and φ^y are always positive. This pins down a unique solution for $\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]$. Defining the right hand side as a polynomial in $\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]$, it has one positive and one negative root. Also, it is positive at zero. Since the right hand side is a linear function, it follows there is a unique value $\mathbb{E}_t^P\left[\varphi^\pi\varphi^y\right]$ such that the equality holds. Plugging into our expression for φ^π and φ^y completes the proof. ■

D.5 Proof of Proposition 7

Proof. For simplicity, suppose there is no disagreement about the forecast of demand. Taking expectations as exogenous, the Central Bank implements

$$\begin{aligned}\tilde{y}_t =& -\psi^y \left(\varepsilon_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] + \delta\pi_{t-1}\right)\\ \tilde{\pi}_t =& \psi^\pi \left(\varepsilon_t + \beta\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\right] + \delta\pi_{t-1}\right)\end{aligned}$$

Conjecture the Central Bank follows a linear rule

$$\begin{aligned}\tilde{y}_t =& -\varphi^y\pi_{t-1} - \alpha\psi^y\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\ \tilde{\pi}_t =& \varphi^\pi\pi_{t-1} + \alpha\psi^\pi\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

Replacing into both equations

$$\begin{aligned}&\varphi^\pi\pi_{t-1} + \alpha\psi^\pi\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\ &\quad = \psi^\pi\delta\pi_{t-1} + \psi^\pi\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right] + \psi^\pi\beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^\pi\right]\pi_{t-1} + \alpha\psi^\pi\beta\mathbb{E}_t^P\left[\varphi^\pi\psi^\pi\right]\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)\mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

$$\begin{aligned}&-\varphi^y\pi_{t-1} - \alpha\psi^y\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\ &\quad = -\psi^y\delta\pi_{t-1} - \psi^y\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right] - \psi^y\beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^\pi\right]y_{t-1} - \alpha\psi^y\beta\mathbb{E}_t^P\left[\varphi^\pi\psi^\pi\right]\sum_{s=0}^{\infty}\left(\alpha\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\end{aligned}$$

Matching coefficients yields

$$\begin{aligned}\alpha =& \frac{1}{1-\beta\mathbb{E}_t^P\left[\varphi^\pi\psi^\pi\right]}\\ \varphi^\pi =& \left(\delta + \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^\pi\right]\right)\psi^\pi\\ \varphi^\pi =& \left(\delta + \beta\mathbb{E}_t^P\left[\varphi^\pi\varphi^\pi\right]\right)\psi^y\end{aligned}$$

Therefore we have $\alpha > 1$. From the second and third equation, notice that $\varphi^\pi = \phi\psi^\pi$ and $\varphi^y = \phi\psi^y$. Solving for ϕ yields

$$\phi = \delta + \beta\mathbb{E}_t^P\left[\psi^\pi\psi^\pi\right]\phi^2$$

We can rewrite this expression to obtain

$$\phi = \frac{\delta}{1 - \beta \mathbb{E}_t^P \left[\varphi^\pi \psi^\pi\right]} = \alpha\delta > \delta$$

Then, we can find α as follows

$$\alpha = 1 + \beta \mathbb{E}_t^P \left[\psi^\pi \psi^\pi\right] \delta \alpha^2$$

There are two solutions, but only one of them ensures $\alpha\beta\mathbb{E}_t^P \left[\psi^\pi\right] < 1$, which is needed for a well-defined policy function. To see this, define $\tilde{\alpha} := \alpha\beta\mathbb{E}_t^P \left[\psi^\pi\right]$. It solves

$$\tilde{\alpha} = \beta \mathbb{E}_t^P \left[\psi^\pi\right] + \delta \frac{\mathbb{E}_t^P \left[\psi^\pi \psi^\pi\right]}{\mathbb{E}_t^P \left[\psi^\pi \psi^\pi\right]} \tilde{\alpha}^2$$

This equation has two solutions $0 < \tilde{\alpha}_1 < 1 < \tilde{\alpha}_2$. We pick the first one to ensure the policy functions are well-defined. This completes the proof. ■

D.6 Proof of Proposition 8

Proof. The estimate of the slope of the Taylor Rule is given by

$$\alpha_{2,t} = \frac{Cov \left(\mathbb{E}_t^i \left[i_{t+k}\right], \mathbb{E}_t^i \left[\pi_{t+k}\right]\right)}{Var \left(\mathbb{E}_t^i \left[\pi_{t+k}\right]\right)}$$

The IS curve implies the following equilibrium behavior of the policy rate

$$\begin{aligned} \mathbb{E}_t^i \left[i_{t+k}\right] =& \mathbb{E}_t^i \left[r_{t+k}^n\right] + \mathbb{E}_i^P \left[\pi_{t+k+1}\right] + \sigma \left(\mathbb{E}_t^i \left[y_{t+k+1}\right] - \mathbb{E}_t^i \left[y_{t+k}\right]\right) \\ =& \mathbb{E}_t^P \left[r_{t+k}^n\right] + \mathbb{E}_t^P \left[X_{t+k+1}\right] - \sigma \mathbb{E}_t^P \left[\psi^y\right] \left(\mathbb{E}_t^P \left[X_{t+k+1}\right] - \mathbb{E}_t^P \left[X_{t+k}\right]\right) \end{aligned}$$

For simplicity assume that inflation forecasts are not correlated with forecasts of the natural rate, we have

$$\begin{aligned} Cov \left(\mathbb{E}_t^i \left[i_{t+k}\right], \mathbb{E}_t^i \left[\pi_{t+k}\right]\right) =& \mathbb{E}_t^P \left[\psi^\pi\right] Cov \left(\mathbb{E}_t^i \left[X_{t+k+1}\right], \mathbb{E}_t^i \left[X_{t+k}\right]\right) \\ & - \sigma \mathbb{E}_t^P \left[\psi^y\right] \mathbb{E}_t^P \left[\psi^\pi\right] \left(Cov \left(\mathbb{E}_t^i \left[X_{t+k+1}\right], \mathbb{E}_t^i \left[X_{t+k}\right]\right) - Var \left(\mathbb{E}_t^i \left[X_{t+k}\right]\right)\right) \end{aligned}$$

Since $\mathbb{E}_t^i[X_{t+k}] = \sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P[\psi^\pi]\right)^s \mathbb{E}_t^i[\varepsilon_{t+k+s}]$ then we have

$$\begin{aligned} Var\left(\mathbb{E}_t^i[X_{t+k}]\right) =& Var\left(\mathbb{E}_t^i[\varepsilon_{t+k}]\right) + \left(\beta\mathbb{E}_t^P[\psi^\pi]\right)^2 Var\left(\mathbb{E}_t^i[X_{t+k+i}]\right) \\ Cov\left(\mathbb{E}_t^i[X_{t+k+1}], \mathbb{E}_t^i[X_{t+k}]\right) =& \beta\mathbb{E}_t^P[\psi^\pi] Var\left(\mathbb{E}_t^i[X_{t+k+1}]\right) \end{aligned}$$

Assume there is no heteroskedasticity in the dispersion of forecasts, so that $Var\left(\mathbb{E}_t^i[\varepsilon_{t+k}]\right)$ does not depend on the horizon k. Therefore, we have

$$Cov\left(\mathbb{E}_t^i[X_{t+k+1}], \mathbb{E}_t^i[X_{t+k}]\right) = \beta\mathbb{E}_t^P[\psi^\pi] Var\left(\mathbb{E}_t^i[X_{t+k}]\right)$$

Then, we have

$$\begin{aligned} Cov\left(\mathbb{E}_t^i[i_{t+k}], \mathbb{E}_t^i[\pi_{t+k}]\right) = & \left(\beta\mathbb{E}_t^P[\psi^\pi]\right)\mathbb{E}_t^P[\psi^\pi] Var\left(\mathbb{E}_t^i[X_{t+k}]\right) \\ & \sigma\mathbb{E}_t^P[\psi^y]\mathbb{E}_t^P[\psi^\pi]\left(1-\beta\mathbb{E}_t^P[\psi^\pi]\right) Var\left(\mathbb{E}_t^i[X_{t+k}]\right) \end{aligned}$$

And the estimand of the taylor rule coefficient is

$$\alpha_{2,t} = \beta\mathbb{E}_t^P[\psi^\pi] + \sigma\frac{\mathbb{E}_t^P[\psi^y]}{\mathbb{E}_t^P[\psi^\pi]}\left(1-\beta\mathbb{E}_t^P[\psi^\pi]\right)$$

The first term is increasing in $\mathbb{E}_t^P[\psi^\pi]$, whereas the second one is decreasing. The derivative of the term is constant and equal to β. Let $f\left(\mathbb{E}_t^P[\psi^\pi]\right)$ denote the second term. We compute the derivative

$$\begin{aligned} f'\left(\mathbb{E}_t^P[\psi^\pi]\right) =& \sigma\frac{\partial\left(\frac{\mathbb{E}^P[\psi^y]}{\mathbb{E}_t^P[\psi^\pi]}\right)}{\partial\mathbb{E}_t^P[\psi^\pi]}\left(1-\beta\mathbb{E}_t^P[\psi^\pi]\right) - \beta\frac{\mathbb{E}_t^P[\psi^y]}{\mathbb{E}_t^P[\psi^\pi]} \\ =& -\kappa^{-1}\left(\sigma\frac{1}{\mathbb{E}_t^P[\psi^\pi]}\left(\frac{1}{\mathbb{E}_t^P[\psi^\pi]}-\beta\right)+\beta\left(\frac{1}{\mathbb{E}_t^P[\psi^\pi]}-1\right)\right) < 0 \end{aligned}$$

and $f'\left(\mathbb{E}_t^P[\psi^\pi]\right)$ is increasing in $\mathbb{E}_t^P[\psi^\pi]$. When $\mathbb{E}_t^P[\psi^\pi] \to 0$ then $f'\left(\mathbb{E}_t^P[\psi^\pi]\right) \to -\infty$, and when $\mathbb{E}_t^P[\psi^\pi] \to 1$ then $f\left(\mathbb{E}_t^P[\psi^\pi]\right) \to -\kappa^{-1}\sigma(1-\beta) < -\beta$ under out assumptions. Thus, there exist a value x^* such that $f(x^*) + \beta = 0$ that maximizes $\alpha_{2,t}$.

Under heteroskedascitiy of forecasts errors, the dispersion of the forecast changes with the horizon. This adds an extra term to the estimand of $\alpha_{2,t}$ that varies both because of the change in reputation and because of changes in the dispersion of short-term forecasts; further complicating the structural interpretation of the Taylor rule coefficient. ■

D.7 Perceived Taylor-rule Coefficients

We show in Appendix C that an improvement in reputation about the central bank's responsiveness to demand-driven fluctuations leads to an increase in the perceived Taylor rule coefficient of the output gap. We compare our monthly estimand for reputation with the perceived Taylor rule coefficient for the output gap in Bauer et al. (2024)

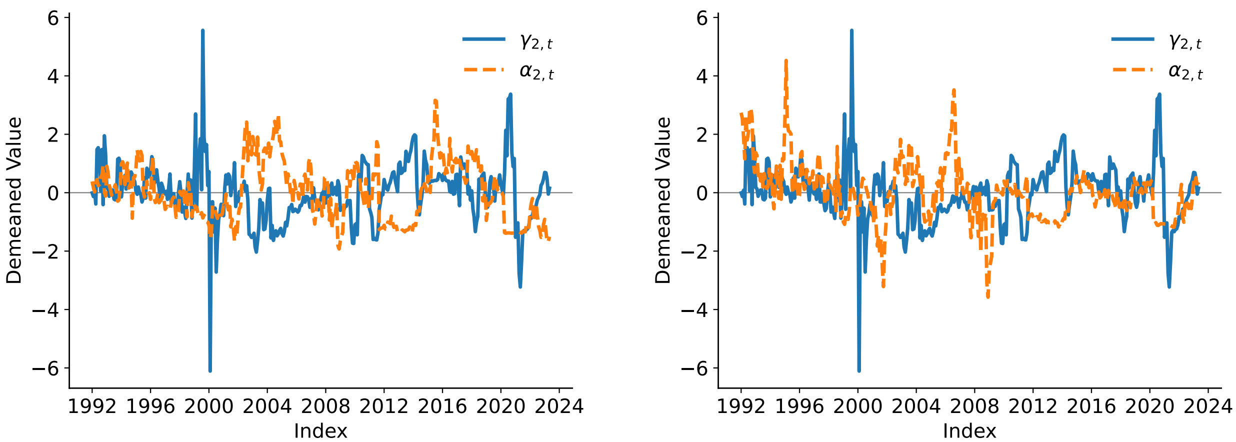

(a) Baseline **(b)** Inertial

Figure 22: Perceived Taylor-rule inflation coefficient for output gap ($\alpha_{2,t}$) from Bauer et al. (2024) and our reputation measure ($\gamma_{2,t}$).

Figure 22 displays a negative comovement of our estimate and the perceived Taylor rule coefficient of output gap, in line with our theory. The correlation coefficient is -0.23 and -0.03, for the baseline and inertial rules, respectively.

D.8 Proof of Proposition 9

Proof. Since the cross-sectional variation in beliefs about the Central Bank is orthogonal to the cross-sectional variation in forecasts of the shocks we have

$$
\begin{aligned}
Cov\left(\mathbb{E}_t^i\left[y_{t+k}\right], \mathbb{E}_t^i\left[\pi_{t+k}\right]\right) &= -Cov\left(\mathbb{E}_t^i\left[\psi^y\right] \sum_{s=0}^{\infty}\left(\beta \mathbb{E}_t^i\left[\psi^\pi\right]\right)^s \mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right], \mathbb{E}_t^i\left[\psi^\pi\right] \sum_{s=0}^{\infty}\left(\beta \mathbb{E}_t^i\left[\psi^\pi\right]\right)^s \mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right) \\
&= -\sum_{s=0}^{\infty} \beta^s \overline{\mathbb{E}}\left[\mathbb{E}_t^i\left[\psi^\pi\right]^{2s+1} \mathbb{E}_t^i\left[\psi^y\right]\right] Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right) \\
&= -\overline{\mathbb{E}}\left[\left(\sum_{s=0}^{\infty}\left(\beta^s \mathbb{E}_t^P\left[\psi^\pi\right]^{s+1}\right)^2 Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right)\right) \frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right]
\end{aligned}
$$

$$
\begin{aligned}
Var\left(\mathbb{E}_t^i\left[\pi_{t+k}\right]\right) =& Cov\left(\mathbb{E}_t^i\left[\psi^\pi\right]\sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^i\left[\psi^\pi\right]\right)^s\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right],\mathbb{E}_t^i\left[\psi^\pi\right]\sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^i\left[\psi^\pi\right]\right)^s\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right)\\
=& \sum_{s=0}^{\infty}\beta^s\overline{\mathbb{E}}\left[\mathbb{E}_t^i\left[\psi^\pi\right]^{2s+2}Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right)\right]\\
=& \overline{\mathbb{E}}\left[\left(\sum_{s=0}^{\infty}\left(\beta^s\mathbb{E}_t^P\left[\psi^\pi\right]^{s+1}\right)^2Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right)\right)\right]
\end{aligned}
$$

Then we have

$$
\frac{Cov\left(\mathbb{E}_t^i\left[y_{t+k}\right],\mathbb{E}_t^i\left[\pi_{t+k}\right]\right)}{Var\left(\mathbb{E}_t^i\left[\pi_{t+k}\right]\right)} = -\overline{\mathbb{E}}\left[\omega^i\frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right]
$$

which depends on the horizon k. Under the assumption $Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right) = \phi^s Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k}\right]\right)$ that expression does not depend on the horizon k and therefore $\gamma_{2,t} = -\overline{\mathbb{E}}\left[\omega^i\frac{\mathbb{E}_t^i[\psi^y]}{\mathbb{E}_t^i[\psi^\pi]}\right]$. Since $\overline{\mathbb{E}}\left[\omega^i\right] = 1$ we have

$$
\gamma_{2,t} = -\overline{\mathbb{E}}\left[\frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right] + Cov\left(\omega^i,-\frac{\mathbb{E}_t^i\left[\psi^y\right]}{\mathbb{E}_t^i\left[\psi^\pi\right]}\right) \tag{B.4}
$$

Notice that ω^i is increasing in $\mathbb{E}_t^P\left[\psi^\pi\right]$, which implies $\gamma_{2,t}$ puts more weight on individuals for which the Central Bank has worse reputation. Then, from (B.4), $\gamma_{2,t}$ is biased towards worse reputation compared to the cross-sectional average reputation.

Focusing on the second term, $Cov\left(\omega^i,-\frac{\mathbb{E}_t^i[\psi^y]}{\mathbb{E}_t^i[\psi^\pi]}\right)$. It is positive, and increasing in the cross-sectional dispersion of the Cenral Bank's reputation. Then, when the disagreement between forecasters increases so does $\gamma_{2,t}$. ■

E Quantitative Exercise

For any beliefs, μ_t, the policy function for the output gap is given by

$$
\tilde{y}_t = \tilde{y}_t^M + \psi^\pi\sum_{s=1}^{\infty}\beta^s\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^P\left[\pi_{t+1+s}\right]}{\partial\tilde{\pi}_t}\tilde{y}_{t+s}\right]
$$

where $\tilde{y}_t^M = -\psi^y X_t$, and $X_t = \sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]$. Let f_η denote the distribution of the forecast error, and $\mu_t\left(s\right)$ denote the prior beliefs over types $s\in[0,1]$. We will study the term in brackets without placing functional form assumptions on f_η or μ_t.

Start with $s=1$. Beliefs evolve as

$$
\mu_{t+1}\left(s\right) = \frac{f_\eta\left(\kappa^{-1}\left(\tilde{\pi}_t^M\left(s\right)-\tilde{\pi}_t\right)+\eta_t\right)\mu_t\left(s\right)}{\int_0^1 f_\eta\left(\kappa^{-1}\left(\tilde{\pi}_t^M\left(x\right)-\tilde{\pi}_t\right)+\eta_t\right)\mu_t\left(x\right)dx} = \frac{f_\eta\left(s;\tilde{\pi}_t,\eta_t\right)\mu_t\left(s\right)}{\int_0^1 f_\eta\left(x;\tilde{\pi}_t,\eta_t\right)\mu_t\left(s\right)}
$$

Notice that

$$\frac{\partial}{\partial\tilde{\pi}_t}\mu_{t+1}(s) = -\kappa^{-1}\left(\frac{f'_\eta(s;\tilde{\pi}_t,\eta_t)}{f_\eta(s;\tilde{\pi}_t,\eta_t)} - \int_0^1 \frac{f'_\eta(x;\tilde{\pi}_t,\eta_t)}{f_\eta(x;\tilde{\pi}_t,\eta_t)}\mu_{t+1}(x)\,dx\right)\mu_{t+1}(s)$$

where $\frac{f'_\eta(s;\tilde{\pi}_t,\eta_t)}{f_\eta(s;\tilde{\pi}_t,\eta_t)} = s_\eta(s;\tilde{\pi}_t,\eta_t)$ is the score of the likelihood of the forecast error, from a statistical model whose location parameter is $\tilde{\pi}_t$. In addition, we have

$$\frac{\partial}{\partial\tilde{\pi}_t}\mathbb{E}^P_{t+1}[\tilde{\pi}_{t+s}] = \left(1 + \frac{\frac{\partial\mathbb{E}^P_{t+1}[X_{t+2}]}{\partial\mathbb{E}^P_{t+1}[\psi^\pi]}}{\mathbb{E}^P_{t+1}[X_{t+2}]}\right)\mathbb{E}^P_{t+1}[\tilde{\pi}_{t+2}]\frac{\partial\mathbb{E}^P_t[\psi^\pi]}{\partial\tilde{\pi}_t} \tag{D.1}$$

Then, the sensitivity of inflation expectations at $t+1$ is

$$\begin{aligned}\frac{\partial\mathbb{E}^P_{t+1}[\psi^\pi]}{\partial\tilde{\pi}_t} &= \frac{\partial}{\partial\tilde{\pi}_t}\int_0^1 s\mu_{t+1}(s)\,ds = \int_0^1 s\frac{\partial}{\partial\tilde{\pi}_{t+1}}\mu_{t+s}(s)\,ds\\ &= -\kappa^{-1}\left(\mathbb{E}^P_{t+1}[\psi^\pi s_\eta(s;\tilde{\pi}_t,\eta_t)] - \mathbb{E}^P_{t+1}[\psi^\pi]\,\mathbb{E}^P_{t+1}[s_\eta(s;\tilde{\pi}_t,\eta_t)]\right)\\ &= -\kappa^{-1}\mathbb{COV}^P_{t+1}[\psi^\pi, s_\eta(\psi^\pi;\tilde{\pi}_t,\eta_t)]\end{aligned}$$

Then, the key statistic to determine the sensitivity of expectations to the central bank's actions is the covariance between the score of the empirical likelihood, and the type. This covariance represents how informative the allocation is of the central bank's type.

When $X_t > 0$, this covariance is negative: starting from $\tilde{\pi}_t > 0$, increasing inflation shifts beliefs towards higher-ψ^π types. When $X_t < 0$, this covariance is positive: starting from $\tilde{\pi}_t < 0$, increasing inflation shifts beliefs towwards lower-ψ^π types.

We now turn to study the first term on (D.1). Assume that shocks obey an AR(1) process with persistence equal to ρ:

$$1 + \frac{\frac{\partial\mathbb{E}^P_{t+s}[X_{t+s+1}]}{\partial\mathbb{E}^P_{t+s}[\psi^\pi]}}{\mathbb{E}^P_{t+s}[X_{t+s+1}]} = 1 + \frac{\frac{\partial}{\partial\mathbb{E}^P_{t+s}[\psi^\pi]}\frac{1}{1-\beta\mathbb{E}^P_{t+s}[\psi^\pi]\rho}}{\frac{1}{1-\beta\mathbb{E}^P_{t+s}[\psi^\pi]\rho}} = 1 + \frac{\beta\rho}{1-\beta\mathbb{E}^P_{t+s}[\psi^\pi]\rho}$$

which is positive for any s. Since $\mathbb{E}^P_{t+1}[\tilde{\pi}_{t+2}]\,\tilde{y}_{t+1} < 0$, $\frac{\partial\mathbb{E}^P_{t+1}[\tilde{\pi}_{t+1}]}{\partial\tilde{\pi}_t}\tilde{y}_{t+1}$ is negative when $X_t > 0$ and positive when $X_t < 0$.

Now we proceed with $s+2$. Belieifs evolve as

$$\mu_{t+2}(s) = \frac{f_\eta(s;\tilde{\pi}_{t+1},\eta_{t+1})\,f_\eta(s;\tilde{\pi}_t,\eta_t)\,\mu_t(s)}{\int_0^1 f_\eta(x;\tilde{\pi}_{t+1},\eta_{t+1})\,f_\eta(x;\tilde{\pi}_t,\eta_t)\,\mu_t(x)\,dx}$$

Then, we have

$$\frac{\partial}{\partial \tilde{\pi}_t}\mu_{t+2}(s) = -\kappa^{-1}\left(s_\eta\left(s;\tilde{\pi}_t,\eta_t\right) - \mathbb{E}^P_{t+2}\left[s_\eta\left(s;\tilde{\pi}_t,\eta_t\right)\right]\right)\mu_{t+s}(s)$$
$$-\kappa^{-1}\left(s_\eta\left(s;\tilde{\pi}_{t+1},\eta_{t+1}\right) - \mathbb{E}^P_{t+2}\left[s_\eta\left(s;\tilde{\pi}_t,\eta_{t+1}\right)\right]\right)\frac{\partial \tilde{\pi}_{t+1}}{\partial \tilde{\pi}_t}\mu_{t+2}(s)$$

Then, there are two ways in which current inflation, $\tilde{\pi}_t$ influence reputation at $t+2$: a direct effect, represented by the first term, and an indirect effect. The indirect effect partially dampens the direct effect: if the central bank improves its reputation then inflation in the following period will be lower. However, it never offsets it: the central bank could achieve the same outcome by overreacting less, which is Pareto dominant. In addition, this dampening is second-order for policy. We can extend this conclusion for any arbitrary horizon s, so that $\frac{\partial \mathbb{E}_{t+s}[\psi^\pi]}{\partial \tilde{\pi}_t} > 0$ when $X_t > 0$, and the opposite when $X_t < 0$.