Reassessing Central Bank Reputation:

Beyond Long-Run Expectations†

Tomás E. Caravello Alex Carrasco Pedro Martínez-Bruera

September 7, 2026

Click here for the latest version.

We study how shifting perceptions about the central bank's relative preference for inflation stability shape monetary policy. In a New Keynesian model with imperfect information, we show that a hawkish reputation—the perception that the bank places greater weight on inflation relative to other objectives—reduces stabilization costs by dampening the pass-through of shocks to short-run inflation expectations. Using U.S. private-sector forecast data, we find evidence of this mechanism even when long-run expectations are anchored and stable. A central bank that internalizes the value of reputation reacts more strongly to cost-push shocks than a myopic bank. Moreover, appointing a more hawkish but myopic central banker achieves welfare gains close to those under the optimal policy. During the Great Moderation, internalizing the value of reputation is equivalent to appointing a central banker whose weight on inflation is roughly twice the bank's true preference.

†Martínez-Bruera: MIT, pmbruera@mit.edu. Caravello: MIT, tomasec@mit.edu. Carrasco: Department of Economics, University of Chile, acarrascom@fen.uchile.cl. We thank Iván Werning, Christian K. Wolf, and Ricardo Caballero, for guidance throughout this project. Comments and suggestions from Olivier Blanchard, Javier García-Cicco, Emilio Espino, Gonzalo Huertas, Vicente Jimenez G., Jackson Mejia, Valeria Morales, Andy Neumeyer, Mikkel Plagborg-Møller, Jim Poterba, Francisco Roldán, and David Romer, as well as seminar participants at Banca d'Italia, CREI, EPGE-FGV Rio, IMF, MIT Macro Lunch, MIT Macroeconomics Seminar, The University of Chicago, Universidad de Buenos Aires, Universidad de San Andrés, Universidad Torcuato di Tella, and University of Notre Dame are greatly appreciated.

1 Introduction

Central bank reputation is crucial to the effectiveness of monetary policy. When the public believes the central bank is committed to low and stable inflation, disinflation can be achieved at lower output cost. Central banks use long-run inflation expectations as a key indicator of reputation. This practice is grounded in a tradition that identifies time-inconsistency and the resulting inflation bias as the main threat to credibility (Kydland and Prescott 1977; Barro and Gordon 1983; Rogoff 1985).

By this metric, reputation seems a solved problem in advanced economies: long-run expectations have remained anchored and stable for almost three decades, including through the post-COVID inflation. In this paper, we show that anchored long-run expectations are *necessary* but not *sufficient* for a strong reputation. Focusing only on long-run expectations misses an important dimension: a shifting perception of the central bank's priority on inflation stability relative to its other objectives when these objectives conflict. This perception shapes how shocks pass through to short-run inflation expectations. We develop a framework to study the implications of this form of reputation, test its main predictions using US data, and quantify its impact on monetary policy.

We build a tractable New Keynesian model, where the central bank has a dual mandate over inflation and output gap stability. The bank's reaction to shocks is shaped by the relative priority given to these objectives. As in standard discretionary analysis, we assume the central bank behaves myopically, as it optimizes period-by-period without internalizing how its actions affect beliefs. The private sector knows the central bank is myopic but does not know its relative priority among conflicting objectives, and it learns about it from observing the bank's actions. While the central bank's relative priority can take any value, it affects the equilibrium through a single sufficient statistic, which we call reputation. A shift in beliefs toward more hawkish types improves reputation.

Our model delivers two testable implications. First, a hawkish reputation anchors short-run inflation expectations, reducing the pass-through of cost-push shocks and lowering the cost of stabilization. Second, since the private sector learns from the central bank's actions, reputation responds to monetary policy surprises. In response to inflationary pressures, raising rates by more (less) than expected improves (worsens) reputation.

Empirically, testing these predictions requires measuring reputation, which is unobservable from time-series data alone: realized outcomes conflate changes in reputation with changes in the state of the economy. However, our theory is informative about how to measure reputation using cross-sectional variations in private-sector forecasts. When forecasters share the same beliefs about the central bank but disagree about future economic conditions, their forecasts co-move in a way that is informative about reputation. Under the structure of our model, we prove that a simple cross-sectional regression of output gap forecasts on inflation forecasts recovers reputation at any point in time. A smaller slope signals that the private sector believes the central bank places greater priority on inflation stabilization. We estimate the Fed's reputation using the Blue Chip Survey of Financial Forecasts (BCFF), which provides individual forecasts from professional forecasters at major financial and economic institutions for the US economy. Using these estimates, we find evidence consistent with both predictions. A more hawkish reputation reduces the pass-through of oil price news shocks to inflation expectations, and an unexpected monetary tightening improves reputation.

Given that our model is stylized, the slope in the cross-sectional regressions may fluctuate for reasons other than reputation, such as structural breaks or disagreements across forecasters.[1] However, these other forces would not by themselves generate the empirical patterns we document. In particular, we show that if these other forces were driving all the variation in the slope, the slope would not correlate with the pass-through of oil shocks to inflation expectations or respond to monetary policy shocks. The fact that we do observe commovements that align with our model provides evidence supporting the presence of our reputation channel.

Having shown the empirical relevance of this channel, we turn to studying its implications for optimal policy. We do so in two steps: we first study the decision problem of a *strategic* central bank, which optimizes subject to the private-sector learning rule *in the baseline equilibrium*. We then show that this can be approximately implemented by a myopic delegate. Regarding the decision problem, we show that a strategic bank *overreacts*, i.e., reacts more than a myopic bank with the same preferences.[2] It responds more strongly to shocks because it internalizes the effect of its actions on reputation. In response to a positive cost-push shock, the magnitude of the output gap response is larger and the magnitude of the inflation response is smaller than under the myopic policy. Relative to what the public expects, a strategic central bank *manages* its reputation, treating it as an intangible asset. While

[1]We discuss these and other potential confounding channels in Section 4.3.

[2]We use the term "overreaction" to emphasize that policy reacts more than what a myopic central bank with the same preferences would choose, not that the response is excessive relative to the social optimum.

having a hawkish reputation is always beneficial, maintaining it requires very aggressive reactions, which are costly. Thus, when reputation is high, a strategic central bank finds it optimal to spend it by reacting less aggressively than expected. Conversely, when the bank is perceived as dovish, it builds reputation by reacting more aggressively than the market expects.

One natural issue with the above analysis is that, if the private sector correctly anticipates the central bank's incentives, it eventually learns that the bank is strategic, and thus its learning rule would change. However, we show that we can approximate its strategic behavior by appointing a myopic central bank with a larger weight on inflation stabilization. In particular, we prove that steady state losses of an appropriately chosen myopic and a strategic central bank coincide up to fourth order in the size of the shocks. The two policies therefore deliver nearly identical welfare for the shock sizes that dominate in normal times, and diverge only for large shocks, which are infrequent. This shows that the gains from the optimal policy can be obtained in a particularly simple way, that does not require managing expectations in a sophisticated manner. This result echoes Rogoff (1985), who proposes appointing a more conservative central banker to offset the inflation bias that arises from time-inconsistency.

Our theoretical insights extend in two directions. First, overreaction also holds in a model with a continuum of strategic banks and internally consistent beliefs. Second, in the presence of interest-rate adjustment costs, this logic extends beyond cost-push shocks. When the Divine Coincidence fails, it is optimal to overreact to demand disturbances as well. Here the bank signals not a higher priority on inflation stability, but a greater willingness to use rates to stabilize the economy. Overreaction, then, emerges as a general principle of optimal monetary policy when the reputation channel is present.

To quantify the importance of the reputation channel, we fit the model on quarterly U.S. data from the pre-pandemic Great Moderation, 1992:Q1-2019:Q4. We calibrate the central bank's preference for inflation stabilization, the volatility of cost-push innovations, and the parameters governing private-sector learning by simulated method of moments. The calibration is disciplined by the joint behavior of inflation and the output gap and, importantly, by the response of inflation expectations to cost-push shocks, which identifies the strength of the reputation channel. Our benchmark specification closely matches its targeted moments and provides a conservative calibration of the gains from reputation management.

We then quantify the gains from reputation management. Relative to the myopic benchmark, the strategic central bank accepts somewhat greater output-gap volatility in exchange for lower inflation volatility and a more hawkish average reputation. Welfare rises by 2.11%.

A useful summary of the mechanism is the pass-through of cost-push shocks to inflation expectations, which provides a directly measurable counterpart to the underlying preference parameters. In our benchmark calibration, this pass-through falls from 0.67 under myopic policy to 0.54 under optimal policy.

Taking the estimated preference as society's objective, the welfare-maximizing myopic delegate places 1.79 times as much weight on inflation stabilization as society. This delegate reproduces almost exactly the outcomes under optimal reputation management: the pass-through falls to 0.53 and the welfare gain is 2.06%, compared with 0.54 and 2.11% under the strategic policy. We also consider the inverse exercise: if the data was generated by an economy with a hawkish delegate, we ask how much lower the true societal weight on inflation would be. Our results imply an underlying social preference only about half as large, so the observed policymaker is 1.94 times as hawkish as society. Under this interpretation, delegation lowers the pass-through from 0.92 to 0.67. Taken together, the quantitative results imply that the value of reputation management is approximately equivalent to appointing a myopic central banker whose weight on inflation stabilization is about twice that of society.

Related Literature. First, our framework contributes to the literature that studies reputation management when there is uncertainty about the central bank's commitment power (Barro, 1986; Backus and Driffill, 1985; Cukierman and Meltzer, 1986; King et al., 2008; Lu et al., 2016; Kostadinov and Roldán, 2025). In these papers, the central bank's preferences are common knowledge, but there is uncertainty about whether the central bank has commitment. Reputation reduces the inflation bias arising from the central bank's preferences (Kydland and Prescott, 1977; Barro and Gordon, 1983) and is built by signaling a strong commitment to price stability. King and Lu (2022) and King and Lu (2025) measure central bank reputation for commitment in the US and study optimal policy in a New Keynesian economy.[3] We study a complementary problem in which the source of uncertainty is reversed: the central bank has no inflation bias, the commitment technology is common knowledge, and uncertainty concerns its relative preference for inflation stability. As in this literature, optimal policy calls for building reputation through a strong response to shocks. The benefit of a good reputation, however, differs: with no inflation bias, reputation instead reduces the pass-through of shocks to inflation expectations. We also empirically document that this reputation mechanism operates in the US economy, where long-run

[3]A related literature also documents the long-lasting effects of reputation for commitment on the country's borrowing costs (Morelli and Moretti, 2023).

inflation expectations are anchored and stable, and quantify the welfare gains from optimal policy.

The closest paper to ours is the contemporaneous work by Bocola et al. (2025), which studies optimal policy in a similar environment with uncertainty about the central bank's preferences. In their model, the sensitivity of *long-run* inflation expectations to monetary policy surprises is a sufficient statistic for reputational considerations. They use this result to measure reputation from the response of the 5y5y forward inflation rate to monetary policy surprises, which they find to be significant in Brazil. We instead use cross-sectional variation in *short-run* forecasts, as in recent work by Bauer et al. (2024), which is informative about reputation even when long-run inflation expectations are anchored. Using this measure, we test the predictions of our model, and we provide evidence of this mechanism for the US economy, where long-run inflation expectations are anchored and stable.[4]

Our empirical strategy connects directly to recent work showing that perceptions of monetary policy vary over time. This literature documents shifts in the perceived degree of commitment (Debortoli and Lakdawala, 2016; Matthes, 2015) or in perceived Taylor-rule coefficients (Hamilton et al., 2011; Bauer et al., 2024; Bocola et al., 2025). Our paper does not aim to provide a new measurement of perceptions of central bank preferences. Rather, we use a model-implied reputation measure about the bank's preferences to validate the central predictions of our theory.[5]

Our paper is also related to the literature on monetary policy, learning, and the anchoring of inflation expectations, see Eusepi and Preston (2018) for a comprehensive review. Molnár and Santoro (2014) study optimal policy when the central bank internalizes the private sector's adaptive learning about the long run, and reach qualitatively similar conclusions. Relative to this work, learning in our model is Bayesian and arises endogenously from uncertainty about an optimizing central bank's preferences, rather than from an exogenous learning rule.

Another branch studies how learning shapes the sensitivity of long-run expectations to

[4]There are also modeling differences. Bocola et al. (2025) adopt 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, which lets us study an extensive margin of reputation: the policy-relevant question is not whether the central bank is hawkish or dovish, but how hawkish the public believes it to be. In our baseline, we also assume the central bank is myopic, though as shown in Section 5, this distinction is not essential for our results.

[5]Relatedly, Lakdawala (2016) uses a structural model to estimate time-varying preferences of the central bank itself. By contrast, we study time-varying perceptions of those preferences.

short-run forecast errors (Carvalho et al., 2023; Bonomo et al., 2024; Bernanke and Blanchard, 2025; Malmendier and Nagel, 2026). With this literature, we share the view that the anchoring of expectations is endogenous and depends on beliefs about monetary policy. However, our notion of reputation differs. We abstract from uncertainty about the long-run inflation target and instead study uncertainty about the central bank's relative preference for inflation stability. This lets us examine how perceived hawkishness shapes short-run expectations and the welfare gains from optimal policy.

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 presents empirical evidence from U.S. forecast data. Section 5 derives the optimal policy. Section 6 provides a quantitative exercise evaluating the relative performance of the optimal policy.

2 The Reputation Channel in a Three-Period Model

To build intuition for the reputation channel, consider a three period New Keynesian economy. The purpose of this exercise is to isolate how uncertainty about the central bank's preferences changes the intertemporal incentives of monetary policy, leaving the explicit Bayesian learning problem to the general model. Inflation is determined by

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

where y_t is the output gap, π_t is inflation, and ε_t is a cost-push shock. The operator $\mathbb{E}_t^P [\cdot]$ denotes private-sector expectations, and our finite period setup implies $\mathbb{E}_t^P [\pi_3] = 0$. Expectations are formed before the central bank chooses period t policy so that $\mathbb{E}_t^P [\pi_{t+1}]$ is predetermined from the perspective of the current policy. Once current outcomes are observed, however, the private sector updates its beliefs, so current policy can affect the expectations formed in subsequent periods.

The central bank evaluates outcomes according to

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

where $\lambda > 0$ is its relative weight on inflation stability and $\mathbb{E}_t^{CB} [\cdot]$ is the central bank's expectation operator. The distinction between the two expectation operators reflects an informational asymmetry. The central bank knows its own λ, whereas the private sector does not and therefore forms expectations using beliefs about it. In the general model, policy is only an imperfect signal of λ because the private sector does not observe all the information available to the central bank, in particular its assessment of the natural rate.

The central bank chooses output and inflation, without commitment, to maximize (2) subject to the Phillips curve (1). We solve this by backward induction. At $t = 2$ there is no future expectation that policy can influence. The central bank therefore solves the standard static stabilization problem, whose first order condition is

$$y_2 + \kappa\lambda \cdot \pi_2 = 0.$$

Together with the Phillips curve and the terminal condition, this implies

$$\pi_2 = \psi^\pi \cdot \varepsilon_2 \quad \text{and} \quad y_2 = -\kappa\lambda\psi^\pi \cdot \varepsilon_2 \tag{3}$$

where $\psi^\pi \equiv \frac{1}{1+\kappa^2\lambda} > 0$ is the pass-through of an inflationary shock into inflation under the central bank's stabilization policy. Furthermore, in this economy long-run inflation expectations are anchored at zero. Since

$$\frac{\partial \psi^\pi}{\partial \lambda} = -\frac{\kappa^2}{(1+\kappa^2\lambda)^2} < 0,$$

a central bank that places greater relative weight on inflation stabilization allows less of a cost push shock to pass through inflation. This observation provides the link between preferences and reputation that will remain central throughout the paper. Beliefs that place greater weight on high values of λ imply a lower expected ψ^π, and hence a more hawkish reputation. Similarly, reputation is more dovish when the private sector's belief assigns greater probability to low values.

At $t = 1$, private sector expectations of terminal inflation have already been formed using (3) and the beliefs over λ inherited from period 0. They therefore enter the Phillips curve as given when the central bank chooses current policy. Moreover, although the period 1 outcome could in principle lead agents to update their beliefs again, such updating cannot affect any subsequent inflation expectation relevant for equilibrium because $\mathbb{E}_t^P[\pi_3] = 0$. There is consequently no reputational incentive left at time $t = 1$, and the optimal output gap is again characterized by the static stabilization rule,

$$\pi_1 = \psi^\pi \cdot X_1 \quad \text{and} \quad y_1 = -\kappa\lambda\psi^\pi \cdot X_1 \tag{4}$$

where $X_1 \equiv \varepsilon_1 + \beta \cdot \mathbb{E}_1^P[\pi_2]$ summarizes the relevant inflationary pressure in this period. This step is useful because it identifies exactly what matters for the policy decision at period $t = 0$: a change in $\mathbb{E}_1^P[\pi_2]$ shifts the inflationary pressure faced by the central bank at $t = 1$ and therefore changes the severity of the future stabilization problem.

The additional incentive appears at $t = 0$. As before $\mathbb{E}_0^P[\pi_1]$ is predetermined when policy is chosen, so current policy cannot manipulate current expectations. It can, however, affect how the private sector interprets the central bank's preferences and thus alter $\mathbb{E}_1^P[\pi_2]$. Using the period 1 optimality condition (4) and the envelope theorem to account for this continuation effect, the optimal output gap satisfies

$$\pi_0 = \overbrace{\psi^\pi \cdot \left(\varepsilon_0 + \beta\mathbb{E}_0^P[\pi_1]\right)}^{\text{Myopic Stabilization}} + \overbrace{\kappa\psi^\pi\beta^2 \cdot \mathbb{E}_0^{CB}\left[\frac{\partial\mathbb{E}_1^P[\pi_2]}{\partial\pi_0}y_1\right]}^{\text{Reputation Management/Intertemporal Smoothing}} . \tag{5}$$

The first term, *myopic stabilization*, is the policy that would obtain if beliefs were unaffected by current outcomes. The second term is the reputational motive. It shows up because a marginal change in current inflation changes the beliefs with which the private sector enters period 1, and hence the inflation expectations that the central bank will confront at that date. If policy conveys no information about λ, the derivative in the second term is zero and the solution reduces to a sequence of static stabilization problems.[6]

The terminal solutions (3) and (4) clarify how learning affects future expectations. Holding fixed the private sector's forecast of the shock process,

$$\frac{\partial\mathbb{E}_1^P[\pi_2]}{\partial\pi_0} = \frac{\partial\mathbb{E}_1^P[\psi^\pi]}{\partial\pi_0}\mathbb{E}_1^P[\varepsilon_2].$$

Two ingredients are therefore required for the reputation channel to affect policy. First, observed policy must change beliefs about the central bank's preferences. In the general model, relatively low inflation for a given inflationary disturbance is interpreted as evidence

[6]In this section, we focus on the inflation allocation for expositional simplicity, since it provides the most direct characterization of the reputation channel. The corresponding output gap allocation follows immediately from the Phillips curve (1) and therefore inherits the same decomposition between static stabilization and the intertemporal reputational motive.

of a larger λ, which lowers the expected pass through $\mathbb{E}_1^P[\psi^\pi]$. Second, there must be an inflationary pressure that is sufficiently persistent to matter for future expectations, so that $\mathbb{E}_1^P[\varepsilon_2] \neq 0$. Either shutting down learning or eliminating the relevant persistence removes the reputational term.

Consider, in particular, a positive and persistent inflationary shock. The static policy response implies $y_1 < 0$. A tighter policy at $t = 0$ lowers current inflation and, through the learning process, shifts beliefs toward a more hawkish central bank. The resulting decline in expected shock pass through reduces $\mathbb{E}_1^P[\pi_2]$, thereby reducing the inflationary pressure inherited by period 1. The central bank is therefore willing to engineer a larger contraction today than under purely static stabilization because doing so improves the inflation output trade off it expects to face tomorrow. This reputational incentive can be interpreted as an *intertemporal smoothing*: by bearing a larger share of the output cost of disinflation today, the central bank reduces the output cost required to contain future inflationary pressures.

This simple setup clarifies the reputation channel studied in the remainder of the paper. In particular, our paper is not primarily concerned with uncertainty about the long run inflation target, or about the ability to commit. In our case,asymmetric information comes from uncertainty about how aggressively the central bank stabilizes inflation in response to shocks, summarized by the expected pass through of inflationary disturbances under the perceived policy rule. The general model developed below makes the learning process explicit.

3 The Economy

The economy consists of a textbook New Keynesian (Woodford, 2003a; Gali, 2010) model with incomplete information. This section lays out the main equations that define equilibrium in our economy. We provide a detailed description and derivation 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):

$$y_t = \mathbb{E}_t^P[y_{t+1}] - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t^P[\pi_{t+1}] - r_t^n\right) \tag{6}$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t^P[\pi_{t+1}] + \varepsilon_t \tag{7}$$

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:

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

where a_t is (log) TFP, and z_t (log) embodies shocks to the household's discount factor. 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 (6) and (7) fix notation and make clear where private-sector expectations enter the equilibrium.

Central Bank and Monetary Policy Regime. The central bank has a dual mandate over the output gap and inflation, represented by the welfare loss function

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{k=0}^{\infty} \beta^k \left(y_{t+k}^2 + \lambda \pi_{t+k}^2\right)\right] \tag{8}$$

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.

This specification, while stylized, is standard: it arises as a second-order approximation to household welfare around the efficient steady state.[7] Under this microfoundation, λ depends on deep parameters of the economy. Here, we instead treat λ as exogenous and dictated by the policymaker's preferences, which is an assumption widely used in applied policy analysis (e.g., Federal Reserve Board 2016, Barnichon and Mesters 2023).

We assume that the Central Bank maximizes (8) by choosing $\tilde{y}_t$ and $\tilde{\pi}_t$ subject to the Phillips curve, taking the path of private sector expectations *as given*, as in the standard discretionary setup of Barro and Gordon (1983). Therefore, everything collapses to a sequence of static problems. We adopt this formulation of central bank behavior as the benchmark for three reasons. First, it delivers a closed-form mapping from beliefs about λ to the allocations (11)-(12), which collapses the impact of beliefs about λ on the equilibrium to a single scalar. Solving instead for a central bank that internalizes the effect of its current choices on *future*

[7] We present this derivation in Appendix A.

beliefs turns the problem into a reputational dynamic game, in which policy today is chosen partly to manipulate tomorrow's posterior; the resulting allocations no longer collapse to a single low-dimensional state variable. Second, and consequently, tractability is what lets us take the model directly to data in Section 4: because reputation is measurable and sufficient under myopia, we can construct its empirical counterpart and test the model's comparative statics without having to first estimate a forward-looking reputational policy function, whose functional form is not pinned down in closed form. Third, myopic discretion is the natural benchmark with no reputation concerns: it describes how a central bank facing period-by-period cost-push shocks would behave absent any incentive to manage beliefs about its type, so this allows us to cleanly isolate our channel as deviations from that benchmark.[8]

To avoid full revelation of the central bank's preferences from its actions, we include a source of noise. At the start of period t, the central bank announces the nominal interest rate, i_t, before the current demand and productivity shocks are realized, that is, before the current value of the natural rate of interest is known.[9] The central bank sets policy based on its private forecast of the natural rate $\mathbb{E}_t^{CB}\left[r_t^n\right]$ which differs from the eventual realization by the forecast error ν_t:

$$r_t^n = \mathbb{E}_t^{CB}\left[r_t^n\right] + \nu_t,$$

with $\mathbb{E}_t^P[\nu_t] = 0$. The private sector does not observe the central bank's forecast, only the realized natural rate. Since $\mathbb{E}_t^P\left[\nu_t\right] = 0$, the private sector believes the central bank is unbiased on average

$$\mathbb{E}_t^P\left[\mathbb{E}_t^{CB}\left[r_t^n\right]\right] = r_t^n.$$

Because the central bank's forecast is unobservable, the noise term ν_t prevents the private sector from perfectly inferring the central bank's preferences from observed policy.[10]

Private Sector Expectations' Formation Process. The private sector knows the full structure of the economy: they correctly believe the central bank optimizes period-by-

[8] In Section 5, we relax myopia and let the central bank choose policy to maximize (8) taking into account its effect on future beliefs, with the private sector correctly anticipating it. The model's qualitative predictions continue to hold.

[9] This friction is similar in spirit to Faust and Svensson (2001), where the central bank has imperfect control over policy outcomes, or to Cukierman and Meltzer (1986), where it has advantageous information about the economy. In both cases, outcome variables are noisy signals about the central bank's preferences.

[10] Another interpretation of ν_t is that it represents a pure monetary policy shock (Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2022)

period without internalizing the private-sector learning. However, they do not know the central bank's relative weight on inflation, λ, nor their forecast for the natural rate $\mathbb{E}_t^{CB}[r_t^n]$. The private sector shares a common prior, μ_t, over possible values of $\lambda \in (0, \infty)$. A myopic central bank of type $\lambda = \tilde{\lambda}$ maximizes(8) subject to the ex-ante equilibrium conditions, (6) and (7)

$$\tilde{y}_t = \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[y_{t+1}]\right] - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[\pi_{t+1}]\right] - \mathbb{E}_t^{CB}[\tilde{r}_t^n]\right) \tag{9}$$

$$\tilde{\pi}_t = \kappa\tilde{y}_t + \beta\mathbb{E}_t^{CB}\left[\mathbb{E}_t^P[\pi_{t+1}]\right] + \varepsilon_t \tag{10}$$

by choosing $\{\tilde{y}_t, \tilde{\pi}_t\}$ taking private-sector expectations as given, 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. This formulation is analogous to assuming the central bank follows a targeting rule in the sense of Giannoni and Woodford (2002), with i_t backed out from (9).

We impose no anticipated learning (Marcet and Sargent, 1989; Kreps, 1998; Evans and Honkapohja, 2001; Cogley and Sargent, 2008; Eusepi and Preston, 2018): 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 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 of type $\lambda = \tilde{\lambda}$ implements:*

$$\tilde{y}_t^M\left(\tilde{\lambda}\right) = -\psi^y\left(\tilde{\lambda}\right)\sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P[\psi^\pi]\right)^s\mathbb{E}_t[\varepsilon_{t+s}] = -\psi^y\left(\tilde{\lambda}\right)X_t \tag{11}$$

$$\tilde{\pi}_t^M\left(\tilde{\lambda}\right) = \psi^\pi\left(\tilde{\lambda}\right)\sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P[\psi^\pi]\right)^s\mathbb{E}_t[\varepsilon_{t+s}] = \psi^\pi\left(\tilde{\lambda}\right)X_t \tag{12}$$

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

Proof. See Appendix A ■

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

$$\mathbb{E}_t^P[y_{t+k}] = -\mathbb{E}_t^P[\psi^y]\,\mathbb{E}_t^P[X_{t+k}] + \mathbb{E}_t^P[\eta_{t+k}] \tag{13}$$

$$\mathbb{E}_t^P\left[\pi_{t+k}\right] = \mathbb{E}_t^P\left[\psi^\pi\right]\mathbb{E}_t^P\left[X_{t+k}\right] + \kappa\mathbb{E}_t^P\left[\eta_{t+k}\right] \tag{14}$$

where $\eta_t := y_t - \tilde{y}_t = \frac{1}{\sigma}\nu_t$ is the central bank's forecast error for the output gap, and $\mathbb{E}_t^P\left[X_{t+k}\right] = \sum_{s=0}^{\infty}\left(\beta\mathbb{E}_t^P\left[\psi^\pi\right]\right)^s\mathbb{E}_t\left[\varepsilon_{t+k+s}\right]$. Under our modeling assumptions, $\mathbb{E}_t^P\left[\eta_{t+k}\right] = 0$, but it need not hold in the data.[11] These forecasts depend on the expected values of ψ^y and ψ^π, the parameters through which reputation shapes the transmission of shocks.

The divine coincidence (Blanchard and Galí, 2007) holds in the baseline model: since demand shocks are fully stabilized, they do not affect inflation expectations. In Appendix B we extend our analysis to a case where demand shocks have a sizeable effect on expectations and show that all of our conclusions hold qualitatively.

The Central Bank's reputation. Private-sector expectations for inflation and the output gap depend on beliefs about λ only through $\mathbb{E}_t^P\left[\psi^y\right]$ 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 entire belief distribution μ_t, not merely the odds of any single type.[12]

Given our assumptions, reputation is a *sufficient statistic* for the private-sector expectations formation. Expectations depend on beliefs only through their effect on reputation. Hence, the infinite-dimensional state variable μ_t can be replaced by the one-dimensional object $\mathbb{E}_t^P[\psi^\pi]$, which renders the problem tractable.

From (13)-(14), 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 in response to a positive cost-push shock.

[11] 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 or long-run inflation target, as in Carvalho et al. (2023).

[12] The continuous-type formulation has two advantages. First, it captures the extensive margin of reputation: changes in reputation reflect movements in the entire belief distribution, rather than the single probability of being a Hawk. Second, it lets us study how higher moments of beliefs shape reputation. As in two-type models, assigning higher odds to more hawkish types improves reputation; in addition, greater dispersion in beliefs worsens it. We formally derive these comparative statics mapping the belief distribution to reputation in Appendix A.

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 $\mathbb{E}_t^P[\psi^\pi]$ is desirable.[13] To see this, assume the cost-push shock is AR(1) with persistence ρ, and combine (10) and (14):

$$\tilde{\pi}_t = \kappa \tilde{y}_t + X_t \qquad X_t = \frac{1}{1 - \beta \mathbb{E}_t^P[\psi^\pi]\rho} \varepsilon_t.$$

In this economy X_t acts as a single cost-push shock, whose magnitude depends on reputation. All else equal, a good reputation shrinks X_t, letting the central bank stabilize inflation at a lower activity cost. However, achieving and maintaining such a reputation is costly. We study how these incentives shape the optimal policy in Section 5.

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. Crucially, decisions are taken before beliefs update, so within-period actions are based on μ_t, which is exogenous. Figure 1 summarizes the sequence.

Given beliefs μ_t, households and firms update beliefs each period as follows. At the end of period t, after all variables have been realized, the public performs the following inference to update beliefs from μ_t to μ_{t+1}:

1. The cost-push shocks are realized $\{\mathbb{E}_t[\varepsilon_{t+s}]\}_{s \geqslant 0}$. 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 (11) and (12), leading to the ex-post realizations

$$\tilde{y}_t^M\left(\tilde{\lambda}\right) + \eta_t \quad \text{and} \quad \tilde{\pi}_t^M\left(\tilde{\lambda}\right) + \kappa \eta_t$$

and analogously at time $t+1$. Given its prior, the private sector forms inflation

[13] This may no longer hold at the Zero Lower Bound, where a more dovish reputation can help the economy escape the ZLB: a larger pass-through of shocks to inflation expectations can offset the effect of a negative natural rate, analogous to Krugman (1998) and Eggertsson and Woodford (2003). We explore this case in Appendix C.

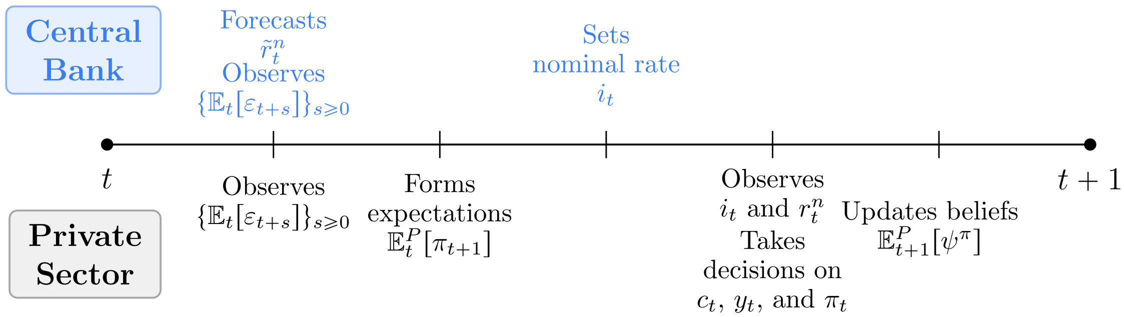

Figure 1: Timing within period

Notes: Blue text denotes central-bank actions and information; black text denotes private-sector actions and information.

expectations as $\mathbb{E}_t^P[\pi_{t+1}] = \int \mathbb{E}_t[\pi_{t+1}^M(\lambda)|\lambda]d\mu_t$.

2. The central bank sets the policy rate i_t to achieve allocations $\tilde{y}_t$ and $\tilde{\pi}_t$. 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, as in Cukierman and Meltzer (1986). 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 this type's forecast error must have been

$$\eta_t\left(\tilde{\lambda}\right) = \kappa^{-1}\left(\pi_t - \psi^\pi\left(\tilde{\lambda}\right)X_t\right) = \kappa^{-1}\left(\tilde{\pi}_t - \psi^\pi\left(\tilde{\lambda}\right)X_t + \kappa\eta_t\right) \tag{15}$$

Beliefs are updated via Bayes' rule:

$$Pr\left(\tilde{\lambda}\middle|\pi_t, y_t\right) = \frac{Pr\left(\pi_t, y_t\middle|\tilde{\lambda}\right)}{Pr\left(\pi_t, y_t\right)}Pr\left(\tilde{\lambda}\right) \Longleftrightarrow \mu_{t+1}\left(\tilde{\lambda}\right) = \frac{f_\eta\left(\eta_t\left(\tilde{\lambda}\right)\right)}{\int f_\eta\left(\eta_t\left(a\right)\right)\mu_t(a)da}\mu_t\left(\tilde{\lambda}\right) \tag{16}$$

where $f_\eta\left(\cdot\right)$ is the time-invariant probability distribution of η_t.[14]

Because ψ^π is a sufficient statistic for how λ affects inflation expectations (Lemma 1), it

[14] Our conclusions are unchanged if forecast errors are expressed in deviations from expected inflation, $\tilde{\pi}_t - \tilde{\pi}_t^M$, expected output gap $\tilde{y}_t - \tilde{y}_t^M$, monetary policy surprises $i_t - \mathbb{E}_t^P\left[i_t\right]$ or the first-order condition $\tilde{y}_t^M + \lambda\kappa\tilde{\pi}_t^M$. We use inflation deviations for simplicity.

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 2. *Under these assumptions, the prior is conjugate, and*

$$\mathbb{E}_t^P\left[\psi^\pi\right] = \overline{\psi}_t - \frac{1}{\sqrt{\tau_t}} \frac{\phi\left(\sqrt{\tau_t}\left(1-\overline{\psi}_t\right)\right) - \phi\left(-\sqrt{\tau_t}\overline{\psi}_t\right)}{\Phi\left(\sqrt{\tau_t}\left(1-\overline{\psi}_t\right)\right) - \Phi\left(-\sqrt{\tau_t}\overline{\psi}_t\right)} \tag{17}$$

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

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

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

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

Lemma 2 establishes that the private sector's Bayesian learning admits a conjugate prior, whose properties we discuss in detail in Appendix B. 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.

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

$$\overline{\psi}_{t+1} = \omega_t X_t^{-1} \overbrace{\left(\tilde{\pi}_t + \kappa\eta_t\right)}^{=\pi_t} + \left(1-\omega_t\right)\overline{\psi}_t \tag{18}$$

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, which is larger when shocks are large relative to noise. The *effective* updating weight $\omega_t X_t^{-1}$, however, is hump-shaped in X_t because large shocks dilute the signal about the central bank's type. Policy affects beliefs through $\pi_t = \tilde{\pi}_t + \kappa\eta_t$: when shocks are inflationary ($X_t > 0$), lower realized inflation shifts beliefs toward more hawkish types. Policy choices therefore affect reputation, which will be central to the optimal policy problem in Section 5.[15]

4 Testing the Reputation Channel in the Data

In this section, we derive and test two main empirical predictions of our model. First, the pass-through of cost-push shocks to inflation expectations is smaller when the central bank has a stronger reputation. Second, reputation responds to monetary policy surprises. To test for them, we need to measure reputation, which is unobservable. Our theory is informative about how to measure it. In this section, we show how to retrieve reputation from forecast data and test both predictions using US data.

A key empirical challenge in measuring reputation is that it evolves endogenously over time. Using time-series variation in *realized* output gaps and inflation would conflate changes in the central bank's reputation with changes in the state of the economy. Our theory provides a way to measure reputation, not from time-series variation of realized outcomes, but from cross-sectional variation in *forecasts*.[16] The reason is that forecasters' expectations reflect the central bank's systematic reaction to shocks (Carvalho and Nechio, 2014), so their cross-sectional variation is informative about reputation.

Suppose at time t, there is a continuum of forecasters indexed by i who share the common model of the economy in Section 3, and identical beliefs about the central bank's preferences. This implies they all have the same prior over the central bank's λ, and use the same estimate of the slope of the NKPC, κ.

At the beginning of period t, each forecaster receives a private signal of the period's shocks and forms forecasts accordingly. This induces cross-sectional variation in forecasts that disappears once the shocks are realized and publicly observed. Therefore, forecast heterogeneity does not impact the equilibrium allocations. Under this assumption, forecaster i's forecasts for the output gap and inflation are

$$\mathbb{E}_t^i\left[y_{t+k}\right] = -\,\mathbb{E}_t^P\left[\psi^y\right]\mathbb{E}_t^i\left[X_{t+k}\right] + \mathbb{E}_t^i\left[\eta_{t+k}\right] \tag{19}$$

[15] The model also rationalizes why a central bank that "falls behind the curve" (Levin and Taylor, 2013; Hakamada and Walsh, 2024) loses reputation: if a central bank repeatedly underestimates demand ($\eta_t > 0$), the resulting inflation overshoots erode its reputation, much like a weak policy reaction to inflationary pressures.

[16] Using cross-sectional dispersion in private forecasts to recover the private sector's perceptions of monetary policy follows Bauer et al. (2024). They estimate the perceived Taylor-rule coefficient. We instead recover perceptions of the central bank's relative weight on inflation, λ. We relate the two objects in Section 4.3.

$$\mathbb{E}_t^i\left[\pi_{t+k}\right] = \mathbb{E}_t^P\left[\psi^\pi\right]\mathbb{E}_t^i\left[X_{t+k}\right] + \kappa\mathbb{E}_t^i\left[\eta_{t+k}\right] \tag{20}$$

where

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

Suppose we have access to survey data on expectations from these forecasters at period t, and consider the regression

$$\mathbb{E}_t^i\left[y_{t+k}\right] = \gamma_{1t} + \gamma_{2t}\mathbb{E}_t^i\left[\pi_{t+k}\right] + u_{t,k}^i \tag{21}$$

Our first result provides conditions under which we can directly recover reputation.

Proposition 1. *Suppose forecasters only disagree about future cost push shocks, so that* $\mathbb{E}_t^i\left[\eta_{t+k}\right] = \mathbb{E}_t^P\left[\eta_{t+k}\right]$. *Then*

$$\gamma_{2t} = -\frac{\mathbb{E}_t^P\left[\psi^y\right]}{\mathbb{E}_t^P\left[\psi^\pi\right]} = -\kappa^{-1}\left(\frac{1}{\mathbb{E}_t^P\left[\psi^\pi\right]} - 1\right) \tag{22}$$

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

Proof. See Appendix D ■

Proposition 1 establishes how our theory allows us to retrieve reputation from the data. When forecasters believe the central bank is myopic and differ only in their expectations about cost-push shocks, we can recover reputation from a simple cross-sectional regression of output-gap forecasts on inflation forecasts.[17] The intuition is straightforward. Consider the optimal policy problem of 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 1 extends this idea to the case where the private sector is uncertain about λ. Beliefs over λ induce a correlation structure between forecasts of the output gap and forecasts of inflation, which is informative about the central bank's reputation.

[17] Since η_t can be interpreted as a pure monetery policy shock, the assumption of homogeneous forecasts of η_{t+k} is analogous to no cross-sectional disagreement about future monetary policy shocks.

An improvement in reputation, captured by a decline in $\mathbb{E}_t^P[\psi^\pi]$, corresponds to a more negative value of $\gamma_{2,t}$. Furthermore, Proposition 1 is sharp. Under those assumptions, $\gamma_{2,t}$ will *only* change in response to changes in reputation. In practice, these assumptions may not hold, resulting in time variation of our estimate independent of changes in reputation. In Section 4.3 we characterize the possible sources of time-variation in $\gamma_{2,t}$ unrelated to reputation and assess their relevance for our tests.

4.1 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 for the U.S. economy. This survey has been used in the literature to measure the private sector's expectations. Each month, forecasters report projections for their current quarter and four to five quarters ahead. We use the sample from January 1992 to December 2024.[18]

The forecasts of output growth and CPI inflation are reported in quarter-over-quarter annualized terms. We follow Bauer et al. (2024) and 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.[19] Potential output is taken from Congressional Budget Office (CBO) projections, also retrieved in real time from ALFRED. Finally, we calculate the output gap as

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

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

We estimate (21) month-by-month using data from January 1992 to December 2024, including forecaster fixed effects to control for forecaster heterogeneity. Figure 2 plots the time series 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, in which

[18]Data are available prior to 1992, but earlier surveys report GNP rather than GDP forecasts, so we restrict the sample to avoid potential inconsistencies.

[19]ALFRED provides the vintages available to forecasters at each survey date. If the exact date is missing, we use the closest vintage.

reputation fluctuates around a long-run mean.

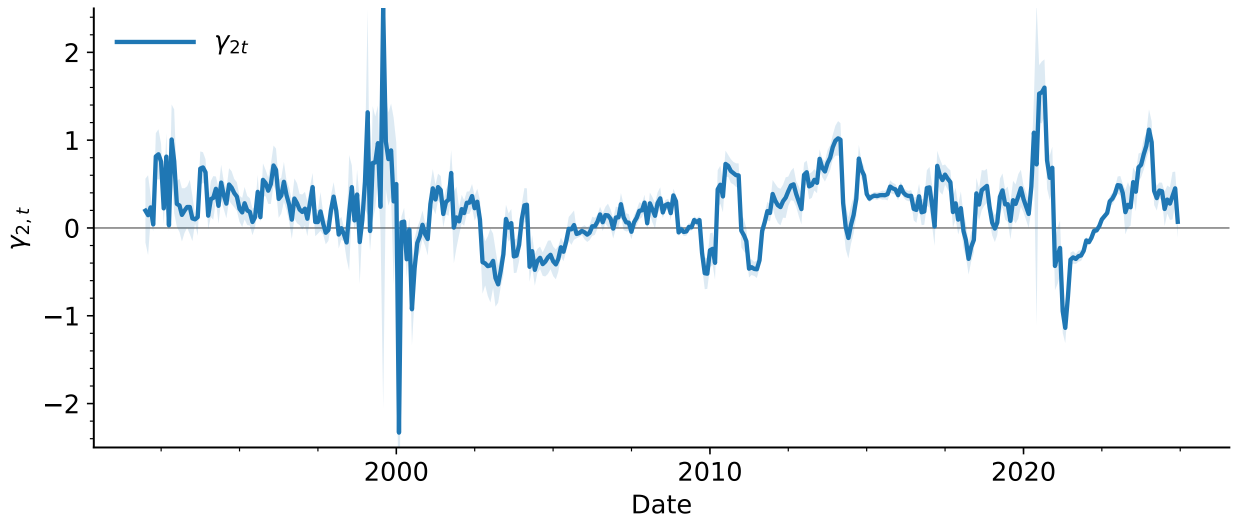

Figure 2: Monthly estimates of the reputation measure $\gamma_{2,t}$

Notes: Monthly BCFF estimates of $\gamma_{2,t}$ from January 1992 to December 2025. Shaded areas denote 90% confidence intervals.

Our measure of perceived hawkishness captures a dimension of reputation distinct from long-run inflation expectations: the correlation between $\gamma_{2,t}$ and the 5-Year, 5-Year forward inflation expectation rate constructed from the SPF is only 0.05. That said, the two are not unrelated: Bocola et al. (2025) document that monetary policy surprises move long-run inflation expectations in Brazil, a setting with greater variation in those expectations, but not for the US.[20] Through the lens of their theory, reputation anchors long-run inflation expectations. We view our findings as complementary. In our model, monetary policy affects reputation, which in turn shapes the pass-through of shocks to short-run inflation expectations. We document both links in the remainder of this section.

As a robustness check, we repeat the exercise using quarterly individual forecasts from the Survey of Professional Forecasters (SPF). Using the same methodology, the correlation coefficient between both estimates of $\gamma_{2,t}$ is 0.60, indicating that both series capture similar underlying variation. In Appendix D, we replicate the empirical analysis below using SPF data and find qualitatively similar results. We now turn to testing the model's two predictions.

[20]We conjecture this reflects Brazil's weaker reputation relative to the advanced economies in their sample.

4.2 Deriving and Testing the Model's Predictions

We derive and test the main model predictions below. We find empirical support for both, which establishes that the channel operates in the U.S. We use these estimates in Section 6 to discipline the quantitative model and measure the welfare gains from the optimal policy.

Prediction #1: Reputation Anchors Inflation Expectations. In our model, the payoff of a good reputation is that the pass-through of shocks to inflation expectations is lower. Suppose for simplicity that the cost-push shock follows an AR(1) process, and forecasters disagree on its persistence ρ_i. From (20), the pass-through of a cost-push shock of size ε on inflation expectations is given by

$$\mathbb{E}_t^i\left[\pi_{t+k} \middle| \varepsilon_t = \varepsilon\right] - \mathbb{E}_t^i\left[\pi_{t+k} \middle| \varepsilon_t = 0\right] = \frac{\mathbb{E}_t^P\left[\psi^\pi\right]}{1 - \beta\mathbb{E}_t^P\left[\psi^\pi\right]\rho_i}\rho_i^k$$

Therefore, a better reputation (lower $\mathbb{E}_t^P\left[\psi^\pi\right]$) reduces the pass-through. We test this prediction using oil price news shocks from Känzig (2021) and the nonlinear local projection

$$\mathbb{E}_t^i\left[\pi_{t+k}\right] = \beta_0^k + \beta_1^k\varepsilon_t + \beta_2^k f\left(\hat{\gamma}_{2,t-1}\right)\varepsilon_t + u_{i,t,k} \tag{23}$$

where $f(\cdot)$ is a positive function that captures how the impulse response may depend on reputation, and $\hat{\gamma}_{2,t-1}$ is our empirical estimate.[21] Oil price news shocks are a natural test: they are widely used cost-push shocks in the literature that raise inflation (Hamilton, 1983; Bernanke et al., 1997; Kilian, 2009) and they also affect short-run inflation expectations (Coibion and Gorodnichenko, 2015). Under the null, $\beta_2^h = 0$ and reputation does not affect pass-through; under the alternative, $\beta_2^h \neq 0$.[22]

Figure 3 plots the impulse response of inflation expectations to an oil price news shock when reputation is high ($\gamma_{2,t}$ at its first quartile) and when it is low ($\gamma_{2,t}$ at its third quartile), using $f\left(\hat{\gamma}_{2,t-1}\right) = \exp\left(-\hat{\gamma}_{2,t-1}\right)$. Pass-through is larger when reputation is dovish and smaller when it is hawkish, consistent with the model's first prediction. The difference vanishes at longer horizons, which may reflect that oil price news shocks are short-lived, or that

[21] We use the lag to avoid simultaneity, and a positive transformation of γ so that the interpretation of the coefficient is clear.

[22] An alternative approach, following Ramey and Zubairy (2018), would be to split the sample into high- and low-reputation regimes. We prefer (23) since our estimate $\gamma_{2,t}$ can vary for reasons unrelated to reputation, which would introduce misclassification error into the regime assignment.

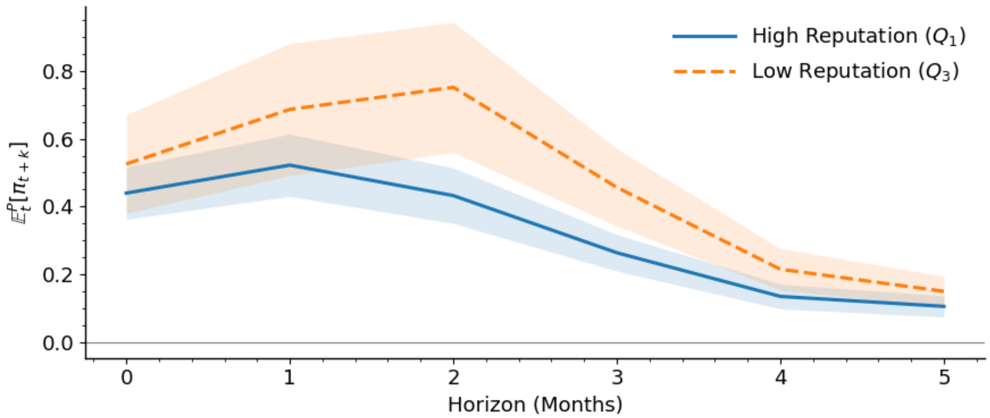

Figure 3: Inflation expectations response to an oil price news shock under high and low reputation

Notes: High and low reputation correspond to the first and third quartiles of $\gamma_{2,t}$, respectively. Shaded areas denote 90% confidence intervals.

reputation matters more for short-run than long-run expectations.

Prediction #2: Reputation Responds to Monetary Policy Surprises. Our second testable prediction is that reputation responds to unexpected changes in the interest rate. Define $mps_t := i_t - \mathbb{E}_t^P[i_t]$ as the monetary policy surprise. From the private sector's learning model, the impulse response of reputation to a monetary policy surprise is given by[23]

$$\mathbb{E}\left[\overline{\psi}_{t+1}\middle| mps_t = \varepsilon^m\right] - \mathbb{E}\left[\overline{\psi}_{t+1}\middle| mps_t = 0\right] = -\mathbb{E}\left[\omega_t X_t^{-1}\right]\frac{\kappa}{\sigma}\varepsilon^m. \tag{24}$$

Crucially, the impact of a monetary policy surprise on reputation depends on the term $\mathbb{E}\left[\omega_t X_t^{-1}\right]$. Intuitively, this term captures the distribution of the cost-push shocks. When there are positive shocks, a positive monetary policy surprise improves reputation. However, when there are negative shocks, the opposite occurs. Thus, this term summarizes what occurs on average. If shocks are mean zero and symmetric, then the average effect of a monetary policy surprise is zero.

Given these observations, we estimate

$$\gamma_{2,t+h} = \alpha_h + \beta_h mps_t + \delta_{1,h}\gamma_{1,t-1} + \delta_{2,h}\gamma_{2,t-1} + u_{t,h} \tag{25}$$

where mps_t is the FF4 monetary policy surprise. In our data, an unexpected tightening

[23]See Appendix B for details.

improves reputation: Figure 4a shows that $\gamma_{2,t}$ declines following a monetary policy surprise. Our estimates are consistent with $\mathbb{E}\left[\omega_t X_t^{-1}\right] > 0$, suggesting that cost-push shocks are either positive on average or negatively skewed.[24] Intuitively, this finding is consistent with the typical cost-push shock calling for a modest tightening, while the rare large realizations call for an aggressive easing.

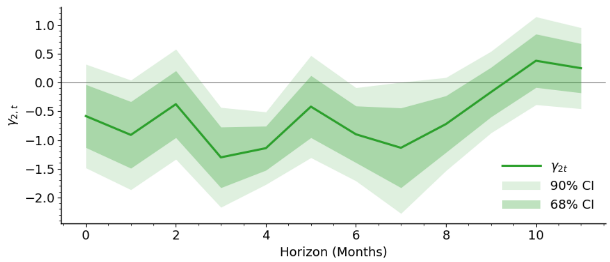

(a) Monetary tightening

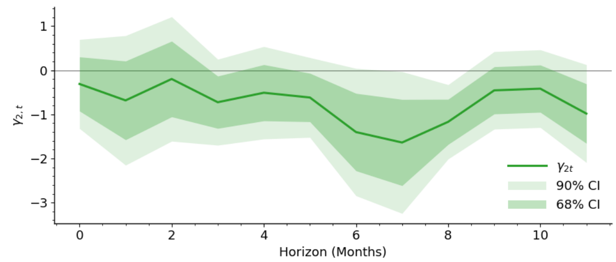

(b) Orthogonalized monetary tightening

Figure 4: Response of the reputation measure to monetary tightening

Notes: Panel (a) uses the FF4 monetary policy surprise; Panel (b) uses the orthogonalized FF4 surprise. Dark and light shaded areas denote 68% and 90% confidence intervals, respectively.

For robustness, we check whether this effect can be explained by the Fed Information Effect (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020) or more generally by the predictability of raw monetary policy surprises. 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 raw shocks using the same six variables in Bauer and Swanson (2022): nonfarm payroll 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 4b plots the impulse response to an orthogonalized monetary tightening. The point estimates are larger in magnitude, but more imprecisely estimated. Overall, our conclusions are not affected by controlling for public information.

4.3 Alternative Sources of Time Variation

In this section, we discuss how departures from our identifying assumptions affect the interpretation of $\gamma_{2,t}$ and the validity of our empirical tests.

[24]Using (18) and $\mathbb{E}[X_t] = 0$, we have $\mathbb{E}\left[\omega_t X_t^{-1}\right] = -\left(\frac{\kappa^{-2}\tau_\eta}{\tau_t}\right)^2 \mathbb{E}\left[X_t^3\right] + \mathbb{E}\left[\mathcal{O}\left(X_t^5\right)\right]$. If cost-push shocks are negatively skewed (i.e., $\mathbb{E}\left[X_t^3\right] < 0$), then $\mathbb{E}\left[\omega_t X_t^{-1}\right] > 0$.

Structural Breaks. Our identification assumes a stable slope of the NKPC, κ. Changes in κ affect our estimate through two channels. First, holding beliefs constant, a change in κ directly alters reputation since $\mathbb{E}_t^P\left[\psi^\pi\right] = \mathbb{E}_t^P\left[\frac{1}{1+\kappa^2\lambda}\right]$. Second, even holding reputation fixed, $\gamma_{2,t} = -\kappa^{-1}\left(\frac{1}{\mathbb{E}_t^P[\psi^\pi]} - 1\right)$ changes mechanically with κ. For both reasons, estimates across periods with different slopes are not directly comparable. However, if these breaks are independent of monetary policy surprises, we would not find that reputation improves with unexpected rate hikes.

Disagreements about forecasts of η_t. Our identification strategy assumes that forecasters do not disagree in their forecasts of the central bank's forecast error, $\mathbb{E}_t^i\left[\eta_{t+k}\right]$. As we show in Appendix D, cross-sectional disagreement about this object induces a positive component in $\gamma_{2,t}$. This bias is analogous to the one monetary policy shocks induce when estimating a Taylor rule (Carvalho et al., 2021). Here, it follows from disagreement about $\mathbb{E}_t\left[\eta_{t+k}\right]$, which can be interpreted as disagreements about future monetary policy shocks.

This bias can be time-varying, as time-varying disagreement generates movements in $\gamma_{2,t}$ unrelated to reputation. However, better reputation always lowers $\gamma_{2,t}$. In addition, if all the time variation in $\gamma_{2,t}$ reflected time-varying disagreement, then our tests would return no effect. Thus, our findings suggest that the reputation channel is present even if some variation in $\gamma_{2,t}$ is driven by this.

Demand Shocks. Our identification strategy assumes the central bank fully stabilizes (its estimate of) demand shocks. This may not hold in the data because the central bank is gradualist. One way to model this is by including an interest rate term in its loss. Then, cross-sectional disagreement about demand shocks will induce a positive correlation structure between forecasts of the output gap and forecasts of inflation. In this case $\gamma_{2,t}$ can be positive, and vary because of time-varying cross-sectional disagreement about demand and cost-push shocks. However, as we show in Appendix D, better reputation still leads to a smaller value of $\gamma_{2,t}$ in this extended setting.

When we allow for demand shocks to have real effects, a new dimension of reputation appears: reputation about the central bank's gradualism. Therefore, reputation about reaction becomes a two-dimensional object: for cost-push shocks, reputation for hawkishness; for demand shocks, for gradualism. We cannot disentangle both dimensions in our dataset, so our estimate $\gamma_{2,t}$ reflects a combination of both.[25]

Belief Heterogeneity. We initially assumed the private sector agreed on the central bank's reputation. Suppose instead that beliefs over λ differ across agents. As we show in Appendix D, our estimate $\gamma_{2,t}$ then reflects both the average belief about reputation and the degree of disagreement about λ. As a result, $\gamma_{2,t}$ can vary even if reputation is unchanged. For instance, lower disagreement about the central bank's preferences would reduce $\gamma_{2,t}$ independently of any aggregate change in reputation.

Our model features a representative agent and is therefore not well-suited to directly model belief heterogeneity about the central bank. However, we can offer the following interpretation. Suppose forecasters share a common prior over λ but draw an individual signal when forming their forecasts, reverting to the common prior once all information is publicly revealed. Under this interpretation, cross-sectional disagreement about the central bank's preferences maps to the precision of beliefs over λ. As we show in Appendix A, higher precision corresponds to better reputation. Therefore, time-varying disagreement about λ can itself be interpreted as variation in reputation, rather than noise.

Model Heterogeneity. When forecasters use heterogeneous models to forecast, derived beliefs over $\psi^\pi = \frac{1}{1+\kappa^2\lambda}$ will be heterogeneous even if they shared a common prior. In particular, each forecaster's estimate of reputation depends on their own value of the NKPC slope, $\mathbb{E}_t^i\left[\psi_i^\pi\right] = \mathbb{E}_t^i\left[\frac{1}{1+\kappa_i^2\lambda}\right]$. Model heterogeneity is therefore a special case of belief heterogeneity, and our estimate $\gamma_{2,t}$ will reflect both average reputation and the degree of model disagreement across forecasters. Time-varying changes in the models used by forecasters can generate spurious variation in $\gamma_{2,t}$. However, the key comparative static is preserved: a shift in beliefs towards a more hawkish central bank will decrease the estimate of $\gamma_{2,t}$.

Alternative Models. In Appendix D we consider how alternative true models of the economy affect our conclusions. First, when the central bank operates under commitment, the population coefficient in (21) is unchanged once we control for lagged output gap forecasts. Second, when the economy is inertial, controlling for lagged inflation forecasts recovers the same population coefficient. Third, to isolate periods in which the ZLB does not bind, we drop observations whose forecasts are at the bound and re-estimate (21). In all three

[25] Separating them would require individual forecasts of demand and cost-push shocks separately, which are not currently available in our data.

cases the resulting series is highly correlated with our baseline estimate of $\gamma_{2,t}$: 0.84, 0.82, and 0.78, respectively.

Continuum of Strategic Central Banks. Our baseline model assumes the central bank is myopic. In Appendix D, we study the case of a continuum of *strategic* central banks, whose incentives are correctly anticipated by households. In this case, we can no longer summarize the entire belief distribution with a single statistic. Instead, our measure $\gamma_{2,t}$ becomes a function of the full belief distribution, and a shift toward more hawkish types generally leads to a more negative $\gamma_{2,t}$, though the magnitude depends on the specific shape of the belief distribution.

Relationship with Perceived Taylor Rule Coefficients. A growing body of literature estimates the private sector's perceived Taylor rule coefficients using survey data (Bauer et al., 2024) or market data (Hamilton et al., 2011; Bocola et al., 2024). How do these estimates relate to perceived central bank preferences? In Appendix D, we establish two results.First, when demand shocks have real effects, the perceived Taylor-rule coefficient on the output gap maps one-to-one to the central bank's reputation for gradualism: a bank perceived as less gradualist carries a higher perceived output-gap coefficient.[26] Second, no such mapping exists for the perceived coefficient on inflation. A central bank with a strong reputation and one with a weak reputation can share the same perceived inflation coefficient. The reason is that a hawkish central bank is perceived as willing to move rates aggressively, but may not *need* to do so as much in response to the same shock, as it feed less into expectations. In practice, this implies there is an inverse-U relationship between reputation and the perceived Taylor rule coefficient on inflation.[27]

So far, we have shown that both predictions of the model hold in the U.S. Having established the empirical relevance of this channel, we now study how it shapes the design of monetary policy.

[26] Moreover, Bauer et al. (2024) document that the perceived Taylor rule coefficient on output gap increases in response to unexpected tightenings when the economy is strong, and decreases when the economy is weak. These findings are in line with the predictions of our model.

[27] This non-monotonicity is similar to the identification problem in McLeay and Tenreyro (2020). Because policy responds optimally, the equilibrium comovement in the data may not be informative about the structural relationships.

5 Optimal Policy

In this section, we study the incentives that this channel generates by focusing on the decision problem of a strategic central bank that internalizes the private sector's expectations formation process and learning. This strategic central bank balances the conventional stabilization trade-off with the benefits from influencing how it is perceived.

In the model, the private sector's prior places zero weight on the strategic central bank, so it would never detect that the central bank was behaving strategically, which is unlikely to hold in practice. For this reason, we also study how to implement a policy that replicates the main features of the strategic one in a simple and internally consistent way.

5.1 Strategic Central Bank's Problem

By definition, a myopic central bank does not internalize the private sector's learning. Consider instead a strategic central bank that does. Given the private sector's beliefs, μ_t, a strategic central bank chooses $\{\tilde{\pi}_t, \tilde{y}_t\}$ to maximize:

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

subject to

(i) the New Keynesian Phillips Curve (NKPC)

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

(ii) The private sector's learning structure,

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

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 $\tau_{t+1} = \tau_t + \left(\kappa^{-1} X_t\right)^2 \tau_\eta$.[28]

Note that given how private sector beliefs are formed, forward guidance is irrelevant. At time t, the private sector's inflation expectations depend only on the central bank's current reputation and the cost-push shock; promises about the future path of policy have no effect, in contrast to the powerful forward guidance of standard models (Del Negro et al., 2023). Furthermore, since inflation expectations are predetermined at t, discretionary policy is time-consistent, in contrast to the traditional literature (Kydland and Prescott, 1977; Barro and Gordon, 1983; Woodford, 2003a).

5.2 Characterizing the Optimal Policy

In this subsection, we study the problem of Central Bank that faces cost-push shocks, and incorporates how private sector beliefs are shaped by its actions. Our analysis yields two main results. First, relative to a myopic central bank, the optimal policy *overreacts*: it raises rates more. The optimal policy manages its reputation: taking actions to improve it when it is low, and spend it when it is high. First, we explain the mechanisms driving each result. Then, we explore how to implement these insights in practice.

Overreaction. Our first result is that the optimal policy reacts more aggressively than the myopic benchmark. Consider a positive cost-push shock. For a myopic central bank, the problem boils down to a static trade-off between inflation and recession. However, it misses that 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. To see how this incentive shapes the optimal policy, notice that the optimal output gap satisfies

$$\tilde{y}_t\left(\tilde{\psi}^\pi\right) = \overbrace{\tilde{y}_t^M\left(\tilde{\psi}^\pi\right)}^{\text{Myopic Stabilization}} + \overbrace{\tilde{\psi}^\pi \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t^{CB}\left[\beta \frac{\partial \mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+1+s}\right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s}\left(\tilde{\psi}^\pi\right)\right]}^{\text{Intertemporal Smoothing}} \tag{26}$$

Intertemporal smoothing means that, if the central bank strengthens its reputation, future shocks will have a smaller impact on inflation expectations. This creates an incentive to

[28]We do not need to write the dynamic IS equation (9) explicitly: given any desired path for ex-ante inflation and the output gap, the central bank can choose the interest rate to implement it. We assume the Zero Lower Bound never binds.

overreact, relative to a myopic central bank. If the current shock is persistent, overreaction reduces its direct future damage. If it is *iid*, it insures against future shocks.

Even though reputation gives rise to the intertemporal smoothing motive, it also affects the myopic stabilization. A worse reputation makes inflation expectations more sensitive, therefore forcing a more aggressive policy response, absent any reputational concerns. A similar result holds in an economy with regime-switching central banks (Debortoli and Nunes, 2014). In that economy, a higher transition probability towards a dovish central bank raises inflation expectations and forces the current appointee to react more strongly to shocks. In our economy, reputation is an endogenous counterpart of such a transition probability.

The policy trade-off is therefore between a deeper *current* recession in response to inflationary pressures, and a milder *future* one. Therefore, 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. Proposition 2 formalizes this intuition

Proposition 2. *Under the optimal policy, output gap overreacts, and inflation underreacts. That is,*

$$|\tilde{y}_t| > \left|\tilde{y}_t^M\right| \quad and \quad |\tilde{\pi}_t| < \left|\tilde{\pi}_t^M\right|$$

Proof. See Appendix B ■

Figure 5 illustrates the intuition behind Proposition 2 through the lens of price theory. Suppose there is a positive cost-push shock at time t. The NKPC plays the role of the budget constraint. The central bank chooses which point along it to sit at, that is, which combination of recession and inflation maximizes its static preferences. Since the central bank has a bliss point at zero output gap and inflation, utility increases toward the origin.

Start with Figure 5a. The black dashed line depicts the NKPC at period t. A myopic central bank chooses the point at which its indifference curve (solid black line) is tangent to the NKPC, trading off recession and inflation without accounting for reputation. This leads to a certain reputation in $t+1$, and consequentially an NKPC at $t+1$, depicted by the dashed black line in Figure 5b. The red dashed indifference curve depicts the preferences of the strategic central bank, which factors in both the current inflation-output trade-off and private-sector learning.[29] Notice that this hypothetical indifference curve is flatter than

[29] We do not claim this is a true indifference curve, since learning occurs in equilibrium and only points on the NKPC have well-defined learning dynamics. We invoke it purely for illustrative purposes.

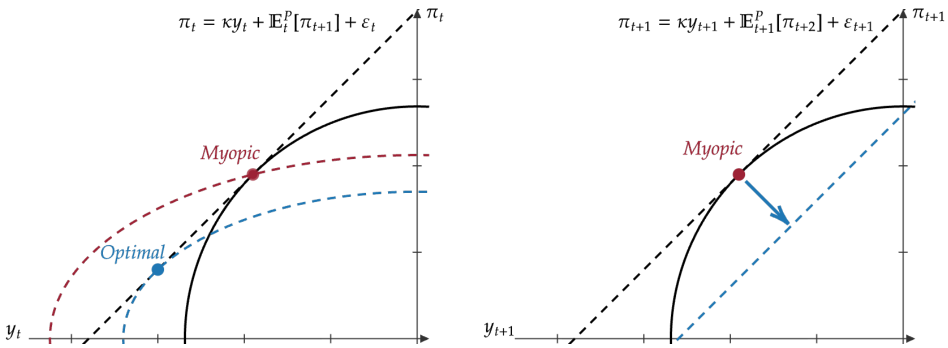

(a) Period-t **(b)** Period-$t+1$

Figure 5: Why optimal policy overreacts to a positive cost push shock

Notes: Panel (a) reports the period-t allocation; Panel (b) reports the period-$t+1$ allocation. Red and blue markers denote the myopic and optimal allocations, respectively.

that of the myopic central bank at the tangency point: the central bank internalizes that being tougher on inflation today improves its reputation, and therefore chooses a point with a larger recession and lower inflation. The payoff comes in period $t+1$: better reputation reduces the pass-through of shocks onto inflation expectations, shifting the NKPC inward in Figure 5b (dashed blue line), a more favorable budget constraint for the future.

Our result shares common features with the literature on reputation for commitment. When the central bank has an inflation bias, it is optimal to overreact to signal commitment, since a committed central bank has less incentive to inflate the economy (King et al., 2008; Lu et al., 2016). In our model, however, there is no inflation bias. We show in Appendix C that, absent an inflation bias, a committed central bank reacts *less* than a discretionary one, smoothing the effect of shocks; in this case, underreaction relative to discretion signals commitment, not overreaction.[30] This shows that, in the absence of inflation bias, uncertainty about the relative weight on inflation yields different predictions that uncertainty about the ability to commit.

A related point is that reputation works in the opposite direction from commitment in a textbook model with complete information (Woodford, 2003a; Giannoni and Woodford, 2002). Under commitment, the central bank smooths the recession over time relative to discretion, trading a smaller current recession for a larger future one. In our model, this logic flips: a larger current contraction signals a high λ, resulting in a softer one in the

[30]In this sense, a committed central bank without an inflation bias need only optimize under commitment for the private sector to learn its type, a *learning Divine Coincidence*.

future. There is also a stark difference with the prescriptions of the earlier literature with incomplete information (Barro and Gordon, 1983; Barro, 1986; Backus and Driffill, 1985). There, a discretionary strategic central bank may have incentives to shift policy *towards* commitment to improve its reputation. Here, the discretionary strategic bank distorts *away* from commitment to improve the inflation-output trade-off in the future.

So far, we have established that the optimal policy always overreacts relative to the myopic benchmark. This raises a distinct question: does the bank always want to *improve* its reputation, or can it be optimal to let it *erode*? The answer depends on where the bank currently stands. The logic, as we show next, is that the central bank treats reputation like an investment problem.

Reputation Management. Overreaction means responding more aggressively than the myopic benchmark, but this need not improve the central bank's reputation. The dynamics of reputation depend on how hawkish or dovish it is perceived to be. From the central bank's perspective, (18), the better the reputation, the lower $\overline{\psi}_t$, the stronger the response needed to maintain it.

By this logic, while a good reputation reduces stabilization costs, sustaining it is itself costly: it requires the central bank to respond aggressively to shocks, potentially at odds with its own preferences. The optimal policy problem can thus be seen as a canonical investment problem (Jorgenson, 1963): reputation is the capital stock, whose return is lower stabilization costs, and whose investment cost is deviating from the central bank's static optimum.[31] Formally, rewrite (26) as

$$\overbrace{\tilde{y}_t^M\left(\tilde{\psi}^\pi\right)-\tilde{y}_t\left(\tilde{\psi}^\pi\right)}^{\text{Cost: Deviation from Myopic}}=\overbrace{-\tilde{\psi}^\pi\sum_{s=1}^{\infty}\beta^s\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+1+s}\right]}{\partial\tilde{\pi}_t}\tilde{y}_{t+s}\left(\tilde{\psi}^\pi\right)\right]}^{\text{Payoff: Lower sensitivity of expectations}}$$

This equation formalizes the trade-off at the heart of optimal policy: the cost of deviating from what static preferences would dictate, against the payoff of a counterfactually better reputation. When reputation is low relative to the bank's static preferences the cost of improving it is small: overreaction is modest, while the payoff from better-anchored short-

[31] This investment logic, in which a cost incurred today yields a return through a more favorable intangible state in the future, is not unique to our setting. Board and Meyer-ter Vehn (2013) apply it to a firm's reputation for quality, Bilal and Rossi-Hansberg (2021) to a household's location choice, and Amador and Phelan (2021) for sovereign default.

run expectations is large. The bank therefore invests in reputation. When reputation is high, however, sustaining it requires a large deviation from the static optimum, and it can be optimal for the policymaker to allow it to erode toward a more sustainable level. Proposition 3 formalizes this:

Proposition 3. *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.*

Proof. See Appendix B. ■

Notice that Proposition 3 is expressed in terms of the central bank's *expected* reputation on the following period. The central bank controls its reputation only in expectation. Realized reputation also depends on its forecast errors. A central bank that systematically underestimates the natural rate will lose reputation despite its best efforts, rationalizing why policymakers who "fall behind the curve" can suffer lasting credibility costs.

Figure 6 illustrates the intuition behind Proposition 3. Suppose there is a positive cost-push shock. The dashed black line depicts the NKPC at period t, and the red dot depicts the allocation the central bank must implement to maintain its current reputation into the following period.

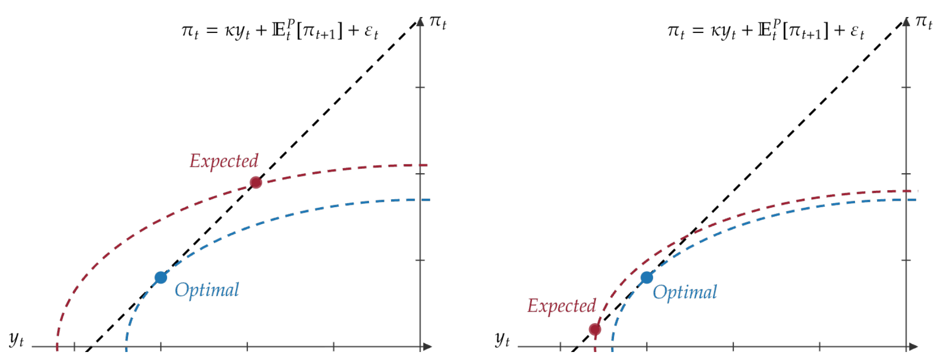

(a) Weak Reputation **(b)** Strong Reputation

Figure 6: Reputation Dynamics: Weak vs. Strong Reputation

Notes: Panels (a) and (b) report weak and strong reputation, respectively. Red and blue markers denote the expected and optimal allocations, respectively.

Consider first Figure 6a, which displays the case of low reputation. Since reputation is low, sustaining it requires only a modest recession paired with relatively high inflation. This is not optimal: at that allocation, the bank's indifference curve (red dashed line) is flatter than the NKPC, signaling a preference for lower inflation and a deeper recession. The bank therefore moves to the tangency of the NKPC and its indifference curve (dashed blue line), implementing a larger recession and lower inflation — and in doing so, improving its reputation.

Figure 6b shows the symmetric case of high reputation. Here, sustaining current reputation demands a very deep recession paired with very low inflation. At that point, the indifference curve reveals a preference for more inflation and a shallower recession, so the bank moves away from the sustainability threshold, accepting a deterioration in reputation in exchange for a more favorable allocation.

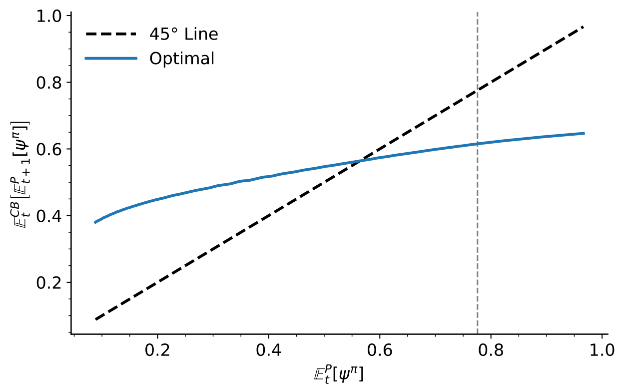

Figure 7: Reputation Dynamics: Law of motion for reputation under optimal policy

Notes: The solid line reports the expected law of motion for reputation under optimal policy. The dashed line denotes the 45-degree line.

Figure 7 shows that the central bank manages its reputation towards a target, building it when perceived as dovish, and letting it decline when perceived as very hawkish. We also observe a direct consequence of our result: our policy function crosses the 45-degree line once, so there are no multiple equilibria.

5.3 Implementing the Optimal Policy via Delegation

This characterization of the optimal policy relies on the private sector assigning zero prior probability to the central bank being strategic. In practice, this is a strong assumption. As long as the market places even a small, nonzero prior on the strategic type, it will eventually learn the truth and detect the strategic bank, which would render the above analysis invalid. The following result shows that we can almost exactly replicate the stabilization under the optimal policy by delegating to a myopic central banker that is more hawkish.

Proposition 4. *Suppose $\tau_t = \tau$ for all t. There exists a preference $\lambda^* > \lambda$ such that*

$$\mathcal{W}^*(\lambda) \;=\; \mathcal{W}^M(\lambda^*, \lambda) + O\left(\sigma_\varepsilon^2 \sigma_{\psi^\pi}^2\right) + O\left(\sigma_\varepsilon^4\right),$$

where $\mathcal{W}^*(\lambda)$ is the steady-state welfare of the strategic central bank and $\mathcal{W}^M(\lambda^*, \lambda)$ is the steady-state welfare generated by a myopic central banker who acts on preference λ^*, both evaluated under the true objective λ, σ_ε is the volatility of the cost-push shock, and σ_{ψ^π} is the volatility of reputation under the strategic central bank's policy.

Proof. See Appendix B ■

Proposition 4 shows that the welfare cost of this delegation relative to the loss of a strategic central bank is fourth order in the volatility of the economy. It is affected by the volatility of shocks and the volatility of reputation. A more volatile economy raises both the cost of the business cycle (lower $\mathcal{W}^*$) and the cost of delegating policy. In addition, central banks more willing to actively manage reputation, with higher $\sigma_{\psi^\pi}^2$, have more to lose when they delegate to a myopic banker who does not do so.

This delegation echoes the seminal contribution of Rogoff (1985), who proposed appointing a more hawkish central banker to offset inflation bias. Here, the delegate instead approximates the policy of a strategic central bank that internalizes the value of reputation. Because the appointed banker is genuinely myopic, the private sector's beliefs are correctly specified, so the equilibrium is internally consistent. One important advantage of this delegation is that it is simple to implement. Furthermore, under delegation, one can partly skip the transition to the new steady state. Under the optimal policy, reputation is built only through the bank's actions. With a myopic delegate, the bank can instead announce the delegate's reaction function ex ante, what Lagarde (2026) recently termed framework guidance. The announcement is incentive-compatible, and can partly avoid the transition costs to the new steady state.

5.4 Extensions

Beyond the baseline results, our framework speaks to several policy questions and extends naturally in a number of directions, as we show next.

State-Dependence. Should the central bank overreact more when reputation is low or high? The answer, shown in Appendix B, is neither. From (26), overreaction is strongest when beliefs are most sensitive to the central bank's actions, a logic similar to the career concerns model of Holmström (1999).[32] Under our functional form assumptions, this occurs at intermediate levels of reputation, when the private sector is uncertain whether the bank is hawkish or dovish. At the extremes, actions are uninformative and beliefs are very rigid. In the middle ground, however, actions are highly informative, and the incentive to overreact peaks. Thus, overreaction follows an inverse-U shape in reputation. Figure 8 plots the policy rate overreaction as a function of the central bank's reputation in response to a positive cost-push shock.

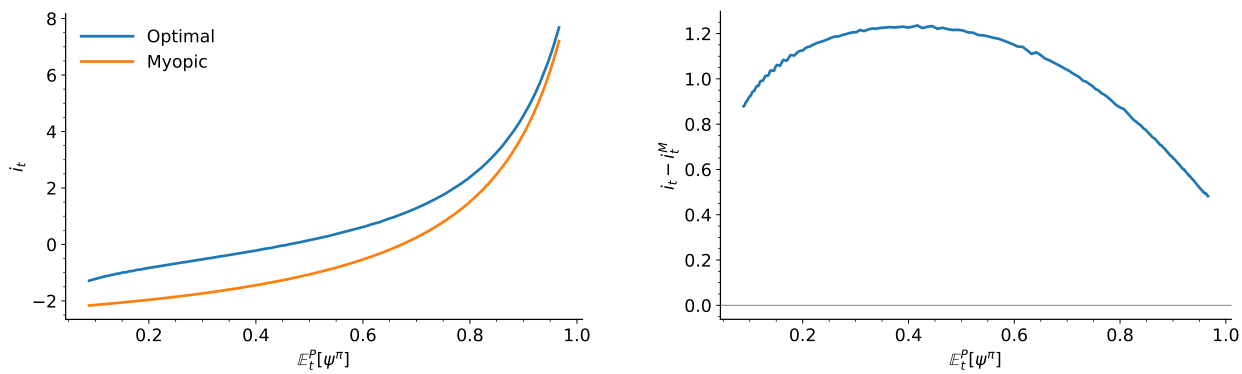

(a) Policy function for the interest rate, i_t

(b) Interest rate overreaction, $i_t - i_t^M$

Figure 8: State Dependence of Policy rate overreaction

Notes: Panel (a) reports the optimal and myopic policy rates; Panel (b) reports their difference. Both panels consider a positive cost-push shock, $\varepsilon_t > 0$.

Good Times vs Bad Times. A long-debated policy question is when the central bank should build reputation and when it should spend it. Is it better to build in good times and spend it during crises, or are bad times precisely when the bank should "show what it is made of"? Our framework clarifies this question. The size of the shock shapes optimal policy through two channels. First, from (18), the effective weight on the bank's action, $\omega_t X_t^{-1}$, is non-monotone in shock size: it increases when the shock is small but decreases when the shock is large, as actions become harder to distinguish from the shock itself. Second, a large

[32]There, the sensitivity depends entirely on the precision of beliefs; in our model, it also depends on the level of reputation.

persistent shock signals a large future recession, raising the incentives to maintain a high reputation.

Taken together, these two channels shape how reactions differ in different times. When a shock is large, the bank's actions are less informative, reducing the incentive to overreact. But if the shock is also persistent, spending reputation is not optimal. The bank correctly anticipates a large recession in the future, and a good reputation remains valuable. Therefore, the only case in which spending reputation is unambiguously optimal is when the shock is both large and short-lived.

Demand Shocks. So far, our discussion has focused on cost-push shocks, which are relatively rare compared to demand-driven fluctuations. Do the same results carry over? To address this, we consider a setting where the Divine Coincidence does not hold, modifying the welfare function to penalize large movements in the policy rate:

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s}^2 + \lambda\tilde{\pi}_{t+s}^2 + \varphi i_{t+s}^2\right)\right] \tag{27}$$

The third term, inspired by Woodford (2003b), captures gradualism. When $\varphi \to 0$, we recover the benchmark where demand shocks are fully stabilized. In Appendix C, we show that, if prices are fixed and the private sector is uncertain about φ, the optimal policy problem is isomorphic to the cost-push case: overreaction in the interest rate, underreaction in the output gap, and all conclusions carry over. In this case, reputation is not about the relative priority given to price stability, but rather the central bank's willingness to move rates aggressively to stabilize the economy. These two interpretations of reputation are complementary, and the logic of optimal policy extends naturally to demand-driven fluctuations.

Communication. In the baseline model, the central bank cannot communicate its internal forecasts to the private sector. The effects of central bank communication and its optimal design has been widely studied as a tool for managing expectations (Blinder et al., 2008, 2024; Gáti and Handlan, 2025; Wabitsch, 2025). Suppose in our model that the central bank were able to communicate its own internal forecasts. What would its incentives be?

Consider a positive cost-push shock, which calls for raising rates. To improve its reputation, the central bank can overreact. But now it has an additional instrument: it can announce that its estimate of the natural rate is low. By doing so, it leads the private sector to believe that demand conditions alone do not warrant a large rate increase. Therefore, if the central bank raises rates aggressively, it must be because it is genuinely tough on inflation. More generally, the central bank has an incentive to bias its announcements in the direction *opposite* to the cost-push shock: downward when the shock is positive, and upward when it is negative. Reputation therefore shapes not only the central bank's *actions* but also its *communication strategy*. We leave their joint study for future research.

Continuum of Strategic Central Banks. Our theoretical insights under a strategic type extend to a case where the private sector correctly anticipates the strategic central bank's incentives. In Appendix C we study an economy where the private sector knows the central bank behaves strategically and it is uncertain about its type, λ, which is assumed to be continuous. We prove that, as long as higher equilibrium inflation signals a lower λ when there are positive shocks (and the opposite when shocks are negative), the strategic central bank overreacts. However, the strength of the signaling may be smaller in equilibrium, a property shared with the career concerns model (Holmström, 1999). Furthermore, in a three-period economy, all of the results in this section continue to hold.

Long-run Inflation Expectations. In our economy, long-run inflation expectations are anchored at zero. Following the traditional literature on reputation and monetary policy, we consider a model where 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.[33] This policy can be interpreted as akin to Average Inflation Targeting (see Powell 2020; Eggertsson and Kohn 2023). We then use the model to study the anchoring of long-run inflation expectations, which we show depends on the central bank's λ.

6 Quantitative Analysis

This section quantifies the importance of the reputation channel. Using U.S. data from the Great Moderation (1992 - 2019), we estimate the model, evaluate the gains from optimal

[33] A direct implication is that the central bank overreacts to persistent shocks.

reputation management, and show that they can be closely replicated by a myopic central banker whose weight on inflation stabilization is about 1.8 to 1.9 times that of society.

6.1 Parameterization and Model Fit

We calibrate the model to quarterly US data. We assume the central bank lacks commitment and behaves myopically. We use this calibration to quantify the strength of the reputation channel and the welfare gains under the optimal policy and delegation.

Our sample covers 1992:Q1 to 2019:Q4, the Great Moderation. We use four US time series from FRED: the Consumer Price Index (CPI), the effective federal funds rate (FFR), real GDP, and the Congressional Budget Office estimate of real potential GDP. The CPI is reported at monthly frequency and the FFR at daily frequency; we aggregate both to quarterly averages.[34] We construct three observable variables from these series. The output gap is

$$y_t^{\text{data}} = 100 \times \log(Y_t/Y_t^n)\,,$$

where Y_t is real GDP and Y_t^n is real potential GDP. Annualized CPI inflation is

$$\pi_t^{\text{data}} = 400 \times [\log(\text{CPI}_t) - \log(\text{CPI}_{t-1})]\,,$$

and the annualized federal funds rate is taken directly from the data.

For our quantitative assessment, we assume the private sector's beliefs have a fixed precision, τ, which keeps them from collapsing to a point mass. This modification allows us to study the reputation channel in a stationary environment. In the baseline learning model, accumulating signals eventually make beliefs arbitrarily precise, causing the response of reputation to new policy outcomes, and hence the reputation channel itself, to vanish asymptotically. Fixing precision abstracts from this long-run convergence while preserving the key mechanism relevant for our quantitative exercises. For a given level of precision, the updating of beliefs and the central bank's incentive to influence future expectations are the same as in the baseline model. The assumption therefore eliminates the secular decline in the strength of learning without altering the local trade-offs that determine our impulse responses and optimal-policy comparisons.

Table 6.1 reports the exogenously calibrated parameters. The discount factor is $\beta = 0.99$,

[34] Results are robust to using end-of-period values of the CPI and FFR rather than quarterly averages.

Parameter	Description	Value	Source
β	Discount factor	0.99	Standard
κ	NKPC slope	0.20/0.40	—
σ	Inverse EIS	1.00	Log utility
υ	Natural rate loading	1.00	Woodford (2003a)
ρ_a	Persistence, productivity	0.95	Standard RBC
σ_{ν^a}	Std. dev., productivity innovation	0.0065	Standard RBC
ρ_ϵ	Persistence, cost-push	0.72	Data

Table 6.1: Calibrated Parameters - Exogenous Block

and log utility implies $\sigma = 1$. We consider two values for the slope of the NKPC. The first, $\kappa = 0.20$, is close to a textbook value (e.g. Galí 2003). The second, $\kappa = 0.40$, makes monetary policy more powerful, letting us evaluate the value of reputation when the inflation-output trade-off is sharper. Productivity and cost-push shocks follow AR(1) processes with normally distributed innovations.

$$\begin{aligned} a_t =& \rho_a a_{t-1} + \nu_t^a \qquad \nu_t^a \sim \mathcal{N}\left(0, \sigma^2_{\nu^a}\right) \\ \varepsilon_t =& \rho_\varepsilon \varepsilon_{t-1} + \nu^{\varepsilon_t} \qquad \nu_t^\varepsilon \sim \mathcal{N}\left(0, \sigma^2_{\nu^\varepsilon}\right) \end{aligned}$$

The persistence and volatility of the productivity shock that drives the natural rate are set to standard RBC values. We calibrate the cost-push shock's persistence to match that of the estimated impulse response of inflation expectations to an oil price news shock. Formally, combining (14) with an AR(1) assumption for the cost-push shock yields

$$\mathbb{E}_t^P\left[\pi_{t+k}\right] = \frac{\mathbb{E}_t^P\left[\psi^\pi\right]}{1 - \beta\rho\mathbb{E}_t^P\left[\psi^\pi\right]}\rho^k\varepsilon_t$$

Consider the Local Projection of inflation expectations on an oil price news shock,

$$\mathbb{E}_t^P\left[\pi_{t+k}\right] = \beta_0^k + \beta_1^k\varepsilon_t + u_{t,k} \tag{28}$$

The model implies that this response decays geometrically, $\beta_1^k = \beta_1^0\rho^k$, so

$$\log\beta_1^k = \log\beta_1^0 + \log\rho \times k. \tag{29}$$

We estimate (29) on the LP estimates from (28), using quarterly SPF data, and recover ρ from the slope. Figure 9 plots the model fit.

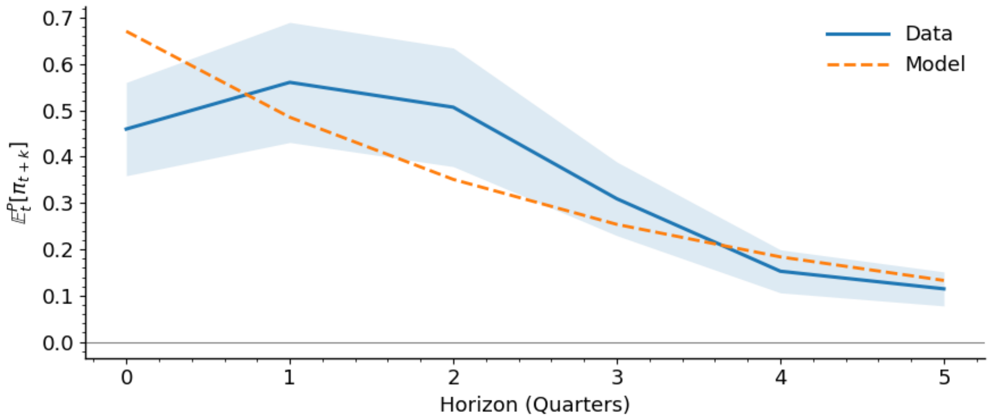

Figure 9: Matching the persistence of the inflation expectations response to cost-push shocks

Notes: The solid line reports the empirical impulse response using quarterly SPF data; the dashed line reports the fitted model response. Shaded areas denote 90% confidence intervals.

The remaining parameters are the central bank's relative preference for inflation stability λ, the volatility of cost-push innovations σ_{ν^ε}, the standard deviation of the central bank's forecast error $\sigma_\eta = \sqrt{1/\tau_\eta}$, and the precision of the private sector's prior τ. We obtain these by the Simulated Method of Moments (SMM). In particular, we calibrate the same set of parameters $(\lambda, \sigma_{\nu^\varepsilon}, \sigma_\eta, \tau)$ separately for each $\kappa \in \{0.2, 0.4\}$ and each of two moment sets. The *small* set contains the volatilities of the output gap and inflation, their correlation, and the impulse response of inflation expectations to an oil price news shock. The *large* set adds the autocorrelation of inflation, the empirical volatility of our estimate $\gamma_{2,t}$, and replaces the unconditional impulse response with its high- and low-reputation counterparts (Appendix D). In our simulation, we recover $\gamma_{2,t}$ from Proposition 1: $\gamma_{2,t} = -\frac{1}{\kappa}\left(\frac{1}{\mathbb{E}_t^P[\psi^\pi]} - 1\right)$. Further details are in Appendix E, and Table 6.2 reports the calibrated parameters.

Table 6.2 reports the parameters under the two values of the NKPC slope, κ, and the small and large sets of matching moments. Increasing κ systematically lowers the value of λ. This is intuitive: given the same comovement between inflation and the output gap, a steeper Phillips curve makes monetary policy more effective at stabilizing inflation, so the central bank needs to place relatively less weight on inflation stabilization to reproduce the same observed behavior.[35,36] Changes in τ and σ_η, by contrast, are best understood through

[35] Accordingly, the level of λ should not be interpreted independently of κ. In the static stabilization problem, their joint effect on inflation is summarized by $\psi^\pi = 1/(1 + \lambda\kappa^2)$.

[36] This mechanism is related to the optimal policy results in Karadi et al. (2024), where a steeper effective NKPC calls for a stronger policy response; here, instead, it implies a smaller calibrated relative weight on

Moment	Data	$\kappa = 0.20$ Small	$\kappa = 0.20$ Large	$\kappa = 0.40$ Small	$\kappa = 0.40$ Large
Parameters					
λ	—	19.76	29.28	12.35	11.43
σ_{ν^ε}	—	2.32e−3	2.44e−3	3.69e−3	4.12e−3
σ_η	—	1.75e−2	1.72e−2	1.28e−2	1.23e−2
τ	—	200.42	293.83	12.18	306.97
κ (fixed)	—	0.20	0.20	0.40	0.40
Sensitivity of $\mathbb{E}_t^P\left[\psi^\pi\right]$	—	0.015	0.010	0.090	0.007
Second moments					
$\sigma(y_t)$	1.756	2.115	2.204	1.756	1.761
$\sigma(\pi_t)$	2.259	1.846	1.666	2.259	2.255
$\sigma(\gamma_{2,t})$	0.343	0.441	0.439	1.855	0.374
$\text{Corr}(\pi_t, \pi_{t-1})$	0.476	0.294	0.220	0.137	0.162
$\text{Corr}(y_t, \pi_t)$	0.357	0.257	0.284	0.357	0.260
Impulse response functions					
cost-push: $E_t^P \pi_{t+h}$ (all), $h = 0$	0.671	0.945	0.718	0.671	0.503
cost-push: $E_t^P \pi_{t+h}$ (all), $h = 1$	0.485	0.683	0.519	0.485	0.364
cost-push: $E_t^P \pi_{t+h}$ (low), $h = 0$	0.663	1.122	0.836	1.040	0.596
cost-push: $E_t^P \pi_{t+h}$ (low), $h = 1$	0.273	0.811	0.605	0.752	0.431
cost-push: $E_t^P \pi_{t+h}$ (high), $h = 0$	0.420	0.775	0.603	0.375	0.412
cost-push: $E_t^P \pi_{t+h}$ (high), $h = 1$	0.140	0.561	0.436	0.271	0.298
SMM objective	—	2.12e−1	4.47e−1	1.90e−8	8.15e−2

Note: Data sample: 1992-Q1 to 2019-Q4 (112 obs.). Each column is one estimation cell (fixed κ × moment set). Model values are cross-simulation point estimates; values in blue were *not* estimation targets in that column. Sensitivity is the response of $\mathbb{E}_t^P[\psi^\pi]$ to a one percentage point unexpected inflation realization; the final row reports the minimized SMM criterion at the estimated parameters.

Table 6.2: Calibrated Parameters - SMM Results

their implications for the strength of learning. A larger τ makes beliefs more anchored in the prior, while a larger σ_η makes policy outcomes less informative about the central bank's preferences. Both therefore tend to reduce the responsiveness of reputation to new inflation outcomes.

inflation stability.

Because these parameters are jointly obtained, however, their movements across specifications should not be interpreted as partial comparative statics. Changing κ alters both the mapping from policy outcomes into information about λ and the calibrated value of λ itself, while changing the set of moments imposes additional restrictions on the combinations of preference, learning, and shock parameters that can reproduce the data. To summarize the net implications of these adjustments, Table 6.2 reports the sensitivity of reputation, $\mathbb{E}_t^P[\psi^\pi]$, to a one-percent unexpected inflation realization, evaluated at $\mathbb{E}_t^P[\psi^\pi] = 1/(1 + \lambda\kappa^2)$. As shown in Lemma 2, this sensitivity is state dependent and depends jointly on κ, τ, σ_η, and the prevailing level of reputation.

Under the small moment set, increasing κ raises the implied sensitivity of reputation to unexpected inflation. Thus, although a larger κ directly makes observed inflation less informative about the central bank's preferences, the associated changes in the other calibrated parameters more than offset this effect. Under the large moment set, instead, the implied sensitivity is essentially unchanged as κ increases. The additional moments therefore substantially restrict the extent to which changes in the assumed Phillips-curve slope translate into changes in the effective strength of learning. While the individual estimates of λ, τ, and σ_η can vary across specifications, the richer calibration delivers a considerably more stable quantitative magnitude for the reputation channel.

Overall, the model reproduces well the volatility of inflation and the output gap and their contemporaneous comovement. As expected from a standard forward-looking New Keynesian model, it generates less inflation persistence than observed in the data. In the counterfactual exercises that follow, we use the $\kappa = 0.40$ specification estimated with the small moment set as our benchmark. Within the small-moment specification, this calibration delivers by far the lowest SMM objective and therefore the closest fit to its targeted moments.[37] Its estimated learning parameters also imply a relatively responsive reputation process and faster convergence across policy regimes, making the transitional and state-dependent mechanisms particularly transparent. At the same time, it provides a conservative benchmark for the quantitative importance of the channel: among our estimation cells, it generates the smallest welfare gain from optimal policy and the smallest reduction in inflation volatility. Our qualitative conclusions do not depend on this choice, and we report the corresponding counterfactuals for the remaining estimation cells in Appendix E.3.

[37]The small-moment specification is just identified in terms of independent restrictions. Exact fit is nevertheless not guaranteed for a given κ: when $\kappa = 0.20$, the empirical moment vector cannot be reproduced closely within the admissible parameter space, and the minimized SMM criterion remains significant.

6.2 Reputation Management and its Implementation

We now quantify both the value of reputation management and a simple way to implement it. Because the model is estimated under the assumption that the myopic policy regime is the data-generating process, our first counterfactual replaces the myopic central bank with the strategic central bank characterized in Section 5, holding fixed the estimated structural parameters and shock processes. We then ask whether the resulting gains can be reproduced by delegating policy to a genuinely myopic but more hawkish central banker. Throughout, we use our benchmark calibration, $\kappa = 0.40$ with the small moment set. We use a relatively steep slope for the Phillips curve because the model primarily captures the reaction to cost-push shocks. Given that these shocks directly affect firms costs, the pass-through to prices (and therefore the slope) can be higher than what standard estimates suggest (Gagliardone et al., 2025).[38]

Optimal Reputation Management. The strategic central bank internalizes how current policy affects future reputation and inflation expectations. Relative to the myopic benchmark, this introduces the intertemporal smoothing motive characterized in Proposition 2: by responding more strongly to current inflationary pressures, the bank improves its reputation, reduces the pass-through of shocks to future inflation expectations, and thereby lowers future stabilization costs.

This strategic policy, however, relies on the private sector continuing to interpret policy through the myopic model. As discussed in Section 5.3, this creates an implementation problem. Proposition 4 shows that the optimal allocation can instead be closely approximated by delegating policy to a genuinely myopic central banker with a larger weight on inflation stabilization. We therefore consider two complementary delegation exercises. First, taking the estimated preference λ as society's objective, we choose the delegate's behavioral weight λ^d to reproduce as closely as possible the welfare attained under optimal reputation management:

$$\lambda^* \equiv \underset{\lambda^d}{\operatorname{argmin}} \, \|\mathcal{W}^M(\lambda^d, \lambda) - \mathcal{W}^*(\lambda)\|.$$

This exercise asks how much more hawkish a myopic delegate must be to obtain the gains generated by a strategic central bank. Second, we consider the inverse interpretation of

[38] When these shocks are large, state-dependence in price setting can further steepen the slope, which is important for optimal monetary policy (Karadi et al., 2025).

our calibration. Our estimate of λ is recovered from observed policy under the maintained assumption that the central bank behaves myopically. If the observed policymaker is itself a delegate, however, the estimated λ corresponds to the delegate's behavioral preference rather than society's underlying objective. We therefore recover the social preference λ^{**} for which the observed behavioral weight λ reproduces the welfare of the corresponding strategic central bank:

$$\lambda^{**} \equiv \underset{\lambda^s}{\operatorname{argmin}} \|\mathcal{W}^M(\lambda, \lambda^s) - \mathcal{W}^*(\lambda^s)\|.$$

The first exercise takes the estimated preference as society's objective and asks how hawkish the delegate should be. The second takes the estimated preference as the behavior of an observed delegate and asks how dovish the underlying social objective must be to rationalize that delegation.[39]

	Society λ			Society λ^{**}	
	Naive myopic(λ)	Optimal optimal(λ)	Delegate myopic(λ^*)	Naive myopic(λ^{**})	Delegate myopic(λ)
Preference weights					
Behavioural weight	12.35	12.35	22.08	6.35	12.35
Ratio to society's weight	—	—	1.79	—	1.94
Volatilities, reputation, and welfare					
$\sigma(y)$	1.78	1.86	1.87	1.65	1.78
$\sigma(\pi)$	2.27	2.14	2.13	2.62	2.27
$\sigma(i)$	5.99	6.02	6.03	5.96	5.99
Reputation, $E^P[\psi^\pi]$	0.38	0.33	0.32	0.48	0.38
Pass-through, $E_t^P[\pi_{t+1}]$	0.67	0.54	0.53	0.92	0.67
Welfare gain over naive (%)	—	2.11	2.06	—	4.03

Note: Great Moderation sample: 1992-Q1 to 2019-Q4 (112 obs.). One estimation cell at its estimated parameters; columns are policy regimes, blocks are societies; levels compare within a block. Welfare is the discounted sum and gains are in percent of the benchmark. The pass-through row is the impact response of expected inflation to a cost-push shock, the $h = 0$ target of the estimation. Here λ^* is the welfare-maximising delegate, not an exact match.

Table 6.3: Optimal, naive and delegated central bank ($\kappa = 0.40$, Small moment set)

[39] The inverse exercise also provides a robustness check on the interpretation underlying our calibration. Rather than assuming that the observed myopic central bank shares society's objective, it allows the observed policymaker to be a delegate with different preferences.

Quantitative Results. Table 6.3 first compares the myopic and optimal policies. Consistent with the theoretical characterization, optimal policy raises output-gap volatility from 1.78 to 1.86 while reducing inflation volatility from 2.27 to 2.14. Average reputation also becomes more hawkish, with $\mathbb{E}^P[\psi^\pi]$ falling from 0.38 to 0.33. The pass-through of the shock to inflation expectations provides a direct measure of the payoff from this improvement in reputation: it falls from 0.67 under myopic policy to 0.54 under optimal policy. Thus, the strategic bank not only stabilizes realized inflation more effectively, but also weakens the propagation of inflationary disturbances through expectations, which is the central benefit of reputation in the model.

The additional stabilization does not require substantially greater policy-rate volatility, which rises only from 5.99 to 6.02. The stronger response needed to manage reputation is partly offset by the resulting stabilization of inflation expectations, which reduces the policy adjustment required subsequently. Under the central bank's objective, the reduction in inflation losses more than compensates for the additional output-gap losses, yielding a welfare gain of 2.11% relative to the myopic benchmark.

The delegation exercise delivers almost the same outcomes as the optimal policy regime. Taking $\lambda = 12.35$ as society's preference, the implied delegate has $\lambda^* = 22.08$, or approximately 1.79 times society's weight on inflation stabilization. The delegate generates inflation volatility of 2.13, average reputation of 0.32, and a cost-push pass-through of 0.53, compared with 2.14, 0.33, and 0.54 under the strategic policy. Its welfare gain is 2.06%, compared with 2.11% under optimal reputation management. In short, appointing a myopic central banker about 79% more hawkish than society reduces the pass-through of cost-push shocks to inflation expectations from 0.67 to 0.53, essentially reproducing the 0.54 pass-through under optimal policy. Delegation therefore reproduces not only the welfare consequences of the strategic policy, but also its effect on the propagation of inflationary shocks through expectations.

The inverse exercise yields a similar magnitude. Interpreting the estimated $\lambda = 12.35$ as the behavioral preference of an existing delegate implies an underlying social preference of $\lambda^{**} = 6.35$. The observed policymaker is therefore approximately 1.94 times as hawkish as society. Relative to a myopic policymaker sharing this underlying social preference, delegation lowers the cost-push pass-through from 0.92 to 0.67 and generates a welfare gain of 4.03%. Thus, a myopic central banker about 94% more hawkish than society lowers the pass-through of cost-push shocks to inflation expectations by roughly one quarter, from 0.92 to 0.67. Whether we interpret the estimated preference as society's objective or as the preference of an existing delegate, the quantitative implication is therefore similar: internalizing the value of reputation is equivalent to appointing a myopic central banker whose weight on inflation stabilization is roughly twice the underlying social weight.[40]

Taken together, these exercises suggest that the gains from optimal reputation management can be implemented almost entirely through a simple delegation scheme. The logic echoes Rogoff (1985), but the source of the delegation motive is different. Rather than offsetting an inflationary bias generated by time inconsistency, the more hawkish delegate replicates the additional stabilization incentives created by the value of reputation. Because the delegate is genuinely myopic, the private sector's model of policy is correctly specified, eliminating the detection problem associated with the strategic policy. The corresponding results for the remaining estimation cells are reported in Appendix E.3.

7 Concluding Remarks

This paper studies the design of monetary policy when the private sector is uncertain about the central bank's relative preference for inflation stability. We develop a tractable framework in which this uncertainty—summarized by reputation—shapes both the propagation of shocks and the optimal policy response.

Our main theoretical results follow from two properties of reputation. First, a hawkish reputation reduces stabilization costs by dampening the pass-through of cost-push shocks to short-run inflation expectations. Second, reputation evolves endogenously: the private sector learns from the bank's actions, so policy choices shape future beliefs. Because of this, the optimal central bank overreacts to cost-push shocks relative to a myopic benchmark—it internalizes the value of its reputation. Furthermore, the bank manages reputation as an asset: it invests in it when perceived as dovish, and draws it down when perceived as hawkish.

Empirically, we construct a measure of reputation from cross-sectional variation in U.S. professional forecasts and find broad support for the theoretical mechanisms: a hawkish reputation dampens the pass-through of cost-push shocks to inflation expectations, and unexpected monetary tightenings improve it. Our measure captures a dimension of credibility distinct from long-run inflation expectations, the standard empirical proxy.

[40] Appendix E.2 reports the transition dynamics and state-dependent responses of reputation under the alternative policy regimes.

Quantitatively, the optimal reputation management raises welfare by 2.11% in our benchmark Great Moderation calibration. Importantly, these gains can be closely approximated by delegating policy to a conservative but myopic central banker whose weight on inflation stabilization is about 1.8 times that of society. This delegation provides a simple and robust institutional solution in the spirit of Rogoff (1985).

References

Amador, M. and C. Phelan (2021): "Reputation and Sovereign Default," *Econometrica*, 89, 1979–2010. 31

Backus, D. and J. Driffill (1985): "Inflation and Reputation," *The American Economic Review*, 75, 530–538. 4, 31

Barnichon, R. and G. Mesters (2023): "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, 113, 2809–2845. 10

Barro, R. J. (1986): "Reputation in a Model of Monetary Policy with Incomplete Information," *Journal of Monetary Economics*, 17, 3–20. 4, 31

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. 1, 4, 10, 28, 31

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024): "Perceptions About Monetary Policy," *The Quarterly Journal of Economics*, 139, 2227–2278. 5, 17, 19, 26, 90, 108

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. 11, 23

——— (2023): "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, 113, 664–700. 23

Bernanke, B. and O. Blanchard (2025): "What Caused the US Pandemic-Era Inflation?" *American Economic Journal: Macroeconomics*, 17, 1–35. 6

Bernanke, B. S., M. Gertler, and M. Watson (1997): "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 28, 91–157. 21

Bilal, A. and E. Rossi-Hansberg (2021): "Location as an Asset," *Econometrica*, 89, 2459–2495. 31

Blanchard, O. and J. Galí (2007): "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 39, 35–65. 13

Blinder, A. S., M. Ehrmann, J. De Haan, and D.-J. Jansen (2024): "Central bank communication with the general public: Promise or false hope?" *Journal of Economic Literature*, 62, 425–457. 36

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008): "Central bank communication and monetary policy: A survey of theory and evidence," *Journal of Economic Literature*, 46, 910–945. 36

Board, S. and M. Meyer-ter Vehn (2013): "Reputation for Quality," *Econometrica*, 81, 2381–2462. 31

Bocola, L., A. Dovis, K. Jørgensen, and R. Kirpalani (2025): "Monetary Policy without an Anchor," Working Paper 34436, National Bureau of Economic Research. 5, 20, 82

Bocola, L., A. Dovis, K. Jørgensen, and R. Kirpalani (2024): "Bond Market Views of the Fed," Working Paper 32620, National Bureau of Economic Research. 26

Bonomo, M., C. Carvalho, S. Eusepi, M. Perrupato, D. Abib, J. Ayres, and S. Matos (2024): "Abrupt Monetary Policy Change and Unanchoring of Inflation Expectations," *Journal of Monetary Economics*, 145, 103576. 6

Caballero, R. J. and A. Simsek (2022): "Monetary Policy with Opinionated Markets," *American Economic Review*, 112, 2353–92. 60

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023): "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 15, 1–47. 6, 13

Carvalho, C. and F. Nechio (2014): "Do People Understand Monetary Policy?" *Journal of Monetary Economics*, 66, 108–123. 17

Carvalho, C., F. Nechio, and T. Tristão (2021): "Taylor Rule Estimation by OLS," *Journal of Monetary Economics*, 124, 140–154. 24

Cogley, T. and T. J. Sargent (2008): "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 49, 185–221. 12

Coibion, O. and Y. Gorodnichenko (2015): "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 7, 197–232. 21

Cukierman, A. and A. H. Meltzer (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54, 1099–1128. 4, 11, 15

Debortoli, D. and A. Lakdawala (2016): "How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-optimizations," *American Economic Journal: Macroeconomics*, 8, 42–76. 5

Debortoli, D. and R. Nunes (2014): "Monetary Regime Switches and Central Bank Preferences," *Journal of Money, Credit and Banking*, 46, 1591–1626. 29

Del Negro, M., M. P. Giannoni, and C. Patterson (2023): "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics*, 1, 43–79. 28

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. 37

Eggertsson, G. B. and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34, 139–235. 14

Eusepi, S. and B. Preston (2018): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56, 3–59. 5, 12

Evans, G. W. and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*, Princeton University Press. 12, 58

Faust, J. and L. E. O. Svensson (2001): "Transparency and Credibility: Monetary Policy with Unobservable Goals," *International Economic Review*, 42, 369–397. 11

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. 10

Gagliardone, L., M. Gertler, S. Lenzu, and J. Tielens (2025): "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," *American Economic Review*, 115, 3941–3974. 43

Gali, J. (2010): *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, third ed. 9

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. 39

García-Schmidt, M. and M. Woodford (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109, 86–120. 81

Gáti, L. and A. Handlan (2025): "Reputation for Confidence," CEPR Discussion Paper 20734, Centre for Economic Policy Research. 36

Giannoni, M. P. and M. Woodford (2002): "Optimal Interest-Rate Rules: I. General Theory," NBER Working Paper 9419, National Bureau of Economic Research. 12, 30

Hakamada, M. and C. E. Walsh (2024): "The Consequences of Falling Behind the Curve: Inflation Shocks and Policy Delays Under Rational and Behavioral Expectations," IMF Working Paper 2024/042, International Monetary Fund, accessed May 20, 2026. 17

Hamilton, J. D. (1983): "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, 91, 228–248. 21

Hamilton, J. D., S. Pruitt, and S. Borger (2011): "Estimating the Market-Perceived Monetary Policy Rule," *American Economic Journal: Macroeconomics*, 3, 1–28. 5, 26

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. 81

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies*, 66, 169–182. 35, 37

Jarociński, M. and P. Karadi (2020): "Deconstructing Monetary Policy Surprises—The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 12, 1–43. 11, 23

Jorgenson, D. W. (1963): "Capital Theory and Investment Behavior," *The American Economic Review*, 53, 247–259, papers and Proceedings of the Seventy-Fifth Annual Meeting of the American Economic Association. 31

Känzig, D. R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 111, 1092–1125. 21, 96

Karadi, P., A. Nakov, G. Nuño, E. Pasten, and D. Thaler (2024): "Strike while the Iron Is Hot: Optimal Monetary Policy under State-Dependent Pricing," CEPR Discussion Paper 19339, CEPR Press, Paris and London. 40

Karadi, P., A. Nakov, G. Nuño, E. Pastén, and D. Thaler (2025): "Strike While the Iron Is Hot: Optimal Monetary Policy under State-Dependent Pricing," Working Paper Series 3068, European Central Bank. 43

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053–1069. 21

King, R. G. and Y. K. Lu (2022): "Credibility and Explicit Inflation Targeting," NBER Working Paper 30012, National Bureau of Economic Research. 4

——— (2025): "Evolving Reputation for Commitment: Understanding Inflation and Inflation Expectations," Manuscript. 4

King, R. G., Y. K. Lu, and E. S. Pastén (2008): "Managing Expectations," *Journal of Money, Credit and Banking*, 40, 1625–1666. 4, 30, 89

Kostadinov, R. and F. Roldán (2025): "Reputation and the credibility of inflation plans," *Journal of Economic Theory*, 229, 106062, elsevier. 4

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. 12

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. 14, 85

Kydland, F. E. and E. C. Prescott (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473–491. 1, 4, 28

Lagarde, C. (2026): "Back to Basics in an Uncertain Environment," Introductory speech at the ECB Forum on Central Banking, Sintra, Portugal. 34

Lakdawala, A. (2016): "Changes in Federal Reserve Preferences," *Journal of Economic Dynamics and Control*, 70, 124–143. 5

Levin, A. and J. B. Taylor (2013): "Falling Behind the Curve: A Positive Analysis of Stop-Start Monetary Policies and the Great Inflation," in *The Great Inflation: The Rebirth of Modern Central Banking*, ed. by M. D. Bordo and A. Orphanides, University of Chicago Press. 17

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. 4, 30, 89

Malmendier, U. and S. Nagel (2026): "Seemingly Anchored Inflation Expectations," NBER Working Paper 35395, National Bureau of Economic Research. 6, 81

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. 12

Matthes, C. (2015): "Figuring Out the Fed: Beliefs about Policymakers and Gains from Transparency," *Journal of Money, Credit and Banking*, 47, 1–29. 5

McLeay, M. and S. Tenreyro (2020): "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 34, 199–255. 26

Miranda-Agrippino, S. and G. Ricco (2021): "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13, 74–107. 11

Molnár, K. and S. Santoro (2014): "Optimal Monetary Policy When Agents Are Learning," *European Economic Review*, 66, 39–62. 5

Morelli, J. M. and M. Moretti (2023): "Information Frictions, Reputation, and Sovereign Spreads," *Journal of Political Economy*, 131, 3066–3102. 4

Nakamura, E. and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133, 1283–1330. 23

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. 37

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. 21

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *The Quarterly Journal of Economics*, 100, 1169–1189. 1, 3, 34, 46, 47

Sastry, K. A. (2026): "Disagreement about Monetary Policy," *American Economic Journal: Macroeconomics*, 18, 486–522. 60

Wabitsch, A. (2025): "The Messenger Matters," University of Oxford, working paper. Latest version: September 2025. 36

Werning, I. (2022): "Expectations and the Rate of Inflation," Working Paper 30260, National Bureau of Economic Research. 100

Woodford, M. (2003a): *Interest and prices: foundations of a theory of monetary policy*, Princeton University Press. 9, 10, 28, 30, 39, 58

——— (2003b): "Optimal Interest-Rate Smoothing," *The Review of Economic Studies*, 70, 861–886. 36

Reassesing Central Bank Reputation: Beyond Long-Run Expectations

Online Appendix

Contents

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.[41] 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.

[41] 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.[42] 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), $\pi_t = \log P_t - \log P_{t-1}$ 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,

[42] 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[43]:

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

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.

Second-order Approximation of Welfare. We approximate welfare to second order following Woodford (2003a). A key observation is that incomplete information affects only the transmission of monetary policy — it has no bearing on the flexible-price equilibrium, and therefore does not affect the non-stochastic steady state. This allows us to proceed as in the textbook case. Since there is a production subsidy that offsets firms' market power, the steady state is efficient, and only second-order terms survive the approximation.

The relevant object is the evolution of price dispersion, $\Delta_t \equiv \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} dj$, which follows

$$\Delta_t = \theta \Delta_{t-1} + (1-\theta)\left(\frac{P_t^*}{P_t}\right)^{-\epsilon}$$

Crucially, this law of motion is purely backward-looking and depends only on realized prices — it does not involve households' expectations or beliefs. As a result, the dynamics of price dispersion are independent of the private sector's information set and learning parameters. Following Woodford (2003a), the second-order approximation of welfare is then

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

where λ^{Struct} is a function of the preference parameters of the household, and the average frequency of price adjustment. This expression is identical to the full-information case — learning does not alter the welfare criterion.

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

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 $\lambda \geqslant 0$. Notice that λ is arbitrary and idiosyncratic to the central banker, and may or may not be equal to λ^{Struct}.

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 (A.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. 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.[44] Under the learning structure of our model (described

[44]The Central Bank's ex-ante expectations of the private sector's forward looking terms is present in the

below), it turns out that the terms with double expectations are simply the private sector's expectation of the Central Bank's choices:

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

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 (8) subject to

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

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 (8) 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.

literature of monetary policy with disagreements. See, e.g. Caballero and Simsek (2022) and Sastry (2026).

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

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

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,

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

■

A.4 Comparative Statics

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

Lemma 3. *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, $\mathbb{E}_t^P[\psi^\pi]$ is lower under μ_t' than under μ_t.*

Proof. Rewrite the expected value as

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

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 $F_{\mu_t'}^{\psi^\pi}(\psi^\pi) \geqslant F_{\mu_t}^{\psi^\pi}(\psi^\pi)$ and therefore $\mathbb{E}_t^P[\psi^\pi]$ is lower under μ_t' than under μ_t ■

Lemma 3 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 10 plots the case where first-order stochastic dominance leads to an improvement in reputation.

Lemma 4. *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, $\mathbb{E}_t^P[\psi^\pi]$ is lower under μ_t that under μ_t'*

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

Lemma 4 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.

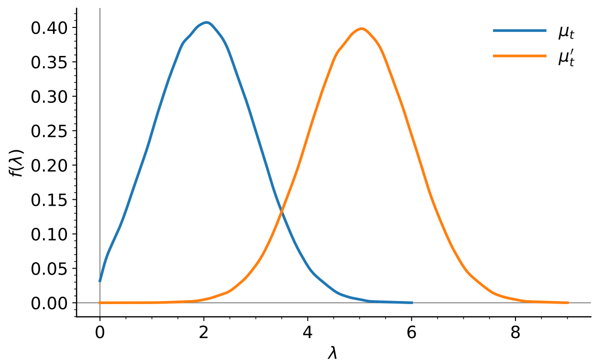

(a) Beliefs over λ

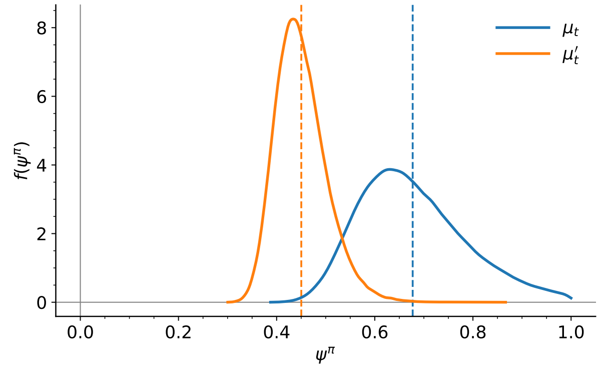

(b) Beliefs over ψ^π

Figure 10: First order stochastic dominance in beliefs and implied reputation: $\mu'_t \overset{FOSD}{\gtrsim} \mu_t$

Notes: Panel (a) reports beliefs over λ; Panel (b) reports the corresponding beliefs over ψ^π. Dashed vertical lines denote the respective means of ψ^π.

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 (15) 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 $\tau_{t+1} = \tau_t + \left(\kappa^{-1}X_t\right)^2 \tau_\eta$.

The properties of this learning structure are illustrated in Figure 11. The first panel depicts that the evolution of beliefs has a mean-reversion component: $\mathbb{E}_t^P\left[\psi^\pi\right] > \overline{\psi}_t$ only when $\overline{\psi}_t < \frac{1}{2}$. To see its effect on the learning structure, rewrite (18) as

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

The first term in this equation captures the effect of mean reversion on the dynamics of reputation. When reputation is poor ($\overline{\psi}_t > \frac{1}{2}$), this force pushes toward an improvement, and the opposite occurs when reputation is good ($\overline{\psi}_t < \frac{1}{2}$).

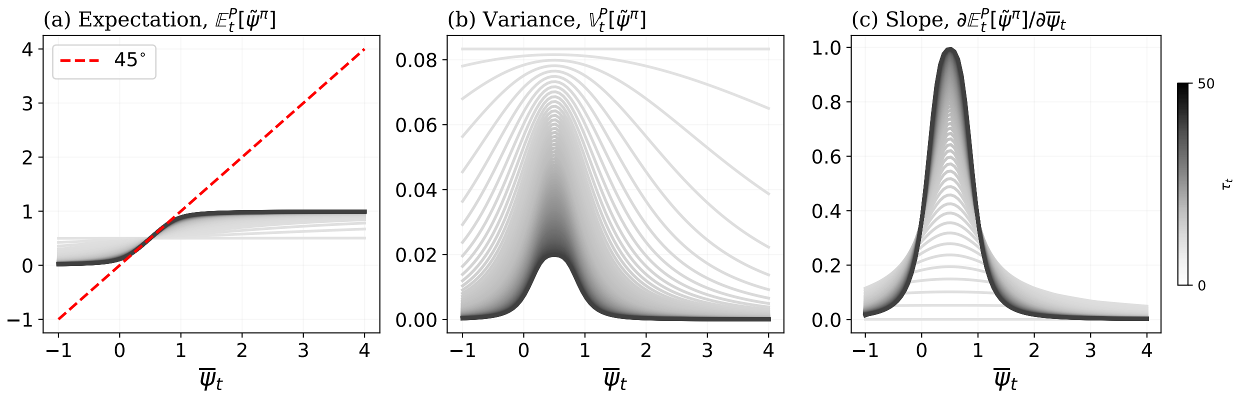

Figure 11: Moments and sensitivity of the truncated normal prior for $\overline{\psi}_t$

Notes: Panels (a)-(c) report the expectation, variance, and slope of expected ψ^π as functions of $\bar{\psi}_t$. Curves correspond to different values of prior precision τ_t, as indicated by the scale.

The second term reflects the effect of policy surprises on reputation. In the presence of inflationary shocks, $X_t > 0$, inflation below what the private sector expected, $\pi_t < \mathbb{E}_t^P[\psi^\pi] X_t$, counterfactually shifts beliefs toward more hawkish types. However, this by itself is not sufficient to improve reputation. If reputation is already high, mean reversion implies that only a sufficiently large hawkish surprise can shift beliefs further. By contrast, when reputation is low, mean reversion works in the opposite direction: reputation may improve even if inflation turns out to be higher than expected.

The second and third panels of Figure 11 show that uncertainty and sensitivity in the private sector's beliefs go hand in hand: states in which beliefs have higher variance are also those in which they are more sensitive to the central bank's actions.

B.2 The Central Bank's Problem

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

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

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,

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

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

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

$$\tilde{\pi}_t = \tilde{\pi}_t^M + \kappa\psi^\pi\sum_{s=1}^{\infty}\beta^s\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{\pi}_{t+s+1}\right]}{\partial\tilde{\pi}_t}\tilde{y}_{t+s}\right] \tag{B.3}$$

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

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

We can rewrite the problem recursively as

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

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

$$g\left(\overline{\psi}\right) = \overline{\psi} - \frac{1}{\sqrt{\tau}}\frac{\phi\left(\sqrt{\tau}\left(1-\overline{\psi}\right)\right) - \phi\left(-\sqrt{\tau}\overline{\psi}\right)}{\Phi\left(\sqrt{\tau}\left(1-\overline{\psi}\right)\right) - \Phi\left(-\sqrt{\tau}\overline{\psi}\right)}$$

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

We discretize ε using the Rouwenhorst 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 $\overline{\psi}'\left(y,\eta\right)$.
   iv. Solve
$$V\left(\overline{\psi}_i,\varepsilon_i\right) = \min_{y}\left\{\frac{1}{2}\left(y^2+\lambda\left(\kappa y+G\right)^2\right)+\beta\sum_{\eta} interp\left(\tilde{\boldsymbol{V}},\overline{\psi}'\left(y,\eta\right)\right)\boldsymbol{F}_{\eta}\right\}$$
   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 2

Proof. From (B.2) and (B.3), it suffices to show that

$$\mathcal{I}_t = \sum_{s=1}^{\infty}\beta^{s}\mathbb{E}_t^{CB}\left[\beta\frac{\partial\mathbb{E}_{t+s}^{P}\left[\tilde{\pi}_{t+s+1}\right]}{\partial\tilde{\pi}_t}\tilde{y}_{t+s}\right] \tag{B.4}$$

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

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

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 (18) 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 State-Dependence

Proposition 5. *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. Recall that the sensitivity of beliefs is given by

$$\frac{\partial \mathbb{E}_{t+s}^{P}\left[\tilde{\pi}_{t+s+1}\right]}{\partial \tilde{\pi}_{t}}=\frac{\partial \mathbb{E}_{t+s}^{P}\left[\tilde{\pi}_{t+s+1}\right]}{\partial \mathbb{E}_{t+s}^{P}\left[\psi^{\pi}\right]}\times\frac{\partial \mathbb{E}_{t+s}^{P}\left[\psi^{\pi}\right]}{\partial \overline{\psi}_{t+s}}\times\frac{\partial \overline{\psi}_{t+s}}{\partial \tilde{\pi}_{t}}. \tag{B.5}$$

Under our functional form assumptions, (B.5) becomes

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

where $\mathbb{V}_{t+s}^{P}\left[\psi^{\pi}\right]$ 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. Since $\frac{\partial \mathbb{E}_{t+s}^{P}[\tilde{\pi}_{t+s+1}]}{\partial \mathbb{E}_{t+s}^{P}[\psi^{\pi}]}$ varies smoothly with reputation, and $\frac{\overline{\psi}_{t+s}}{\partial \tilde{\pi}_{t}}$ is also inverse U-shaped, then $\frac{\partial \mathbb{E}_{t+s}^{P}[\tilde{\pi}_{t+s+1}]}{\partial \tilde{\pi}_{t}}$ is inverse U-shaped. 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 Proof of Proposition 3

Proof. Suppose for simplicity that $X_t > 0$. From the private sector learning, we have

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

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 5, there exist a value $\widehat{\psi} < \psi^\pi$ such that the term in brackets is positive 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.7 Proof of Proposition 4

Proof. The strategic central bank's optimal allocations are given by (B.2) and (B.3)

$$\begin{aligned} \tilde{y}_t =& - \psi^y X_t + \psi^\pi \mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi \right], \varepsilon_t \right) \\ \tilde{\pi}_t =& \psi^\pi X_t + \kappa \psi^\pi \mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi \right], \varepsilon_t \right) \end{aligned}$$

where

$$\mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi \right], \varepsilon \right) = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^{CB} \left[\beta \frac{\partial \mathbb{E}_{t+s}^P \left[\tilde{\pi}_{t+s+1} \right]}{\partial \tilde{\pi}_t} \tilde{y}_{t+s} \right]$$

From (18), since shocks are AR(1) we can reparametrize $\mathcal{I}_t$ as $\mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi \right], X_t \right)$, with $\mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi \right], 0 \right) = 0$. Let $\overline{\mathbb{E}}_t^P \left[\psi^\pi \right]$ denote the long-run mean of reputation around the stochastic steady-state. From Taylor's theorem

$$\mathcal{I}_t = \mathcal{I}_1' \left(\mathbb{E}^P \left[\psi^\pi \right] - \overline{\mathbb{E}}^P \left[\psi^\pi \right] \right) + \mathcal{I}_2' X_t + O \left(\left\| \mathbb{E}_t^P \left[\psi^\pi \right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi \right], \varepsilon_t \right\|^2 \right)$$

where

$$\begin{aligned} \mathcal{I}_1' =& \frac{\partial}{\partial \mathbb{E}_t^P \left[\psi^\pi \right]} \mathcal{I}_t \left(\overline{E}_t^P \left[\psi^\pi \right], 0 \right) \\ \mathcal{I}_2' =& \frac{\partial}{\partial \varepsilon_t} \mathcal{I}_t \left(\overline{E}_t^P \left[\psi^\pi \right], 0 \right) \end{aligned}$$

From (18), there is no learning when $\varepsilon_t = 0$. Therefore $\mathcal{I}_1' = 0$. More generally, for any $k \geqslant 1$

$$\frac{\partial^k}{\partial \mathbb{E}_t^P \left[\psi^\pi\right]^k} \mathcal{I}_t \left(\mathbb{E}_t^P \left[\psi^\pi\right], 0\right) = 0 \tag{B.6}$$

By Proposition 2, $\mathcal{I}_2' < 0$. Then, we have

$$\begin{aligned}
\tilde{y}_t =& -\left(\psi^y - \psi^\pi \mathcal{I}_2'\right) X_t + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^2\right) \\
\tilde{\pi}_t =& \psi^\pi \left(1 + \kappa \mathcal{I}_2'\right) X_t + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^2\right)
\end{aligned}$$

Define

$$\lambda^* = \frac{\lambda - \frac{\mathcal{I}_2'}{\kappa}}{1 + \kappa \mathcal{I}_2'}$$

Then, $\lambda^* > \lambda$ and

$$\begin{aligned}
\psi^y \left(\lambda^*\right) =& \frac{\lambda^* \kappa}{1 + \kappa^2 \lambda^*} = \psi^y - \psi^\pi \mathcal{I}_2' \\
\psi^\pi \left(\lambda^*\right) =& \frac{1}{1 + \kappa^2 \lambda^*} = \psi^\pi \left(1 + \kappa \mathcal{I}_2'\right)
\end{aligned}$$

Thus

$$\begin{aligned}
\tilde{y}_t =& -\psi^y \left(\lambda^*\right) X_t + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^2\right) \\
\tilde{\pi}_t =& \psi^\pi \left(\lambda^*\right) X_t + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^2\right)
\end{aligned}$$

From these expressions, up to first order, the optimal policy is equivalent to the policy of a more conservative myopic central banker.

Let $\mathcal{W}^* \left(\lambda\right)_t$ denote the period-t welfare under the optimal policy, then

$$\begin{aligned}
\mathcal{W}_t^* \left(\lambda\right) = -\frac{1}{2} \left[\tilde{y}_t^2 + \lambda \tilde{\pi}_t^2\right] = & \overbrace{-\frac{1}{2} \left[\tilde{y}_t^M \left(\lambda^*\right)^2 + \lambda \tilde{\pi}_t^M \left(\lambda^*\right)^2\right]}^{\mathcal{W}^M\left(\lambda^*, \lambda\right)} + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^4\right) \\
& + O\left(\left\|\mathbb{E}_t^P \left[\psi^\pi\right] - \overline{\mathbb{E}}_t^P \left[\psi^\pi\right], \varepsilon_t\right\|^2\right) \left[\tilde{y}_t^M \left(\lambda^*\right) + \lambda \tilde{\pi}_t^M \left(\lambda^*\right)\right]
\end{aligned}$$

Since ε_t is symmetric and by (B.6), then

$$\mathbb{E}\left[O\left(\left\|\mathbb{E}_t^P\left[\psi^\pi\right]-\overline{\mathbb{E}}_t^P\left[\psi^\pi\right],\varepsilon_t\right\|^2\right)\left[\tilde{y}_t^M\left(\lambda^*\right)+\lambda\tilde{\pi}_t^M\left(\lambda^*\right)\right]\right]=\mathbb{E}\left[O\left(\left\|\mathbb{E}_t^P\left[\psi^\pi\right]-\overline{\mathbb{E}}_t^P\left[\psi^\pi\right],\varepsilon_t\right\|^4\right)\right]$$

and,

$$\mathbb{E}\left[O\left(\left\|\mathbb{E}_t^P\left[\psi^\pi\right]-\overline{\mathbb{E}}_t^P\left[\psi^\pi\right],\varepsilon_t\right\|^4\right)\right]=O\left(\sigma_\varepsilon^2\sigma_{\psi^\pi}^2\right)+O\left(\sigma_\varepsilon^4\right)$$

where $\sigma_{\psi^\pi}^2$ is the variance of reputation in the stationary distribution. Defining

$$\mathcal{W}^*\left(\lambda\right)=\frac{1}{1-\beta}\sum_{t=0}^{\infty}\beta^t\mathcal{W}_t\left(\lambda\right)$$

we have

$$\mathcal{W}^*\left(\lambda\right)=\mathcal{W}^M\left(\lambda^*,\lambda\right)+O\left(\sigma_\varepsilon^2\sigma_{\psi^\pi}^2\right)+O\left(\sigma_\varepsilon^4\right)$$

■

B.8 Sensitivity of expectations to τ

Recall the sensitivity of beliefs is given by

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

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

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

which, for $s \geqslant 2$ increases in τ when τ is small, and decreases when τ is large.[45] 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 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.

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

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

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 (27):

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

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

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

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

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

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

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

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

Matching coefficients for $s = 0$

$$\Theta_0 = 1$$

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

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

For any arbitrary s, we have

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

Putting everything together,

$$\tilde{y}_t(\tilde{\varphi}) = \psi^y(\tilde{\varphi})\sum_{s=0}^{\infty}\left(\sigma\mathbb{E}_t^P\left[\psi^y\right]\right)^s\mathbb{E}_t\left[r_{t+s}^n\right] = \psi^y(\tilde{\varphi})X_t \tag{C.2}$$

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

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

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 (C.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 (C.2) and (C.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

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

Then, the same properties will hold. After observing the demand shocks $\left\{\mathbb{E}_t\left[r^n_{t+s}\right]\right\}_{s\geqslant 0}$ the central bank maximizes
$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s} + \varphi\tilde{i}^2_{t+s}\right)\right]$$
subject to

(i) The Euler Equation
$$\tilde{y}_t = \mathbb{E}^P_t\left[\tilde{y}_{t+1}\right] - \frac{1}{\sigma}\left(\tilde{i}_t - \mathbb{E}_t\left[r^n_t\right]\right)$$

(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 $\omega_t = \frac{X_t^2\tau_\eta}{\tau+X_t^2\tau_\eta}$.

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

$$\tilde{y}_t = \tilde{y}_t^M - \psi^y\sum_{s=1}^{\infty}\mathbb{E}_t^{CB}\left[\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{y}_{t+s+1}\right]}{\partial\tilde{y}_t}\tilde{i}_{t+s}\right] \tag{C.4}$$

$$\tilde{i}_t = \tilde{i}_t^M + \sigma\psi^y\sum_{s=1}^{\infty}\mathbb{E}_t^{CB}\left[\frac{\partial\mathbb{E}_{t+s}^P\left[\tilde{y}_{t+s+1}\right]}{\partial\tilde{y}_t}\tilde{i}_{t+s}\right] \tag{C.5}$$

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

$$\mathcal{W}_t = -\frac{1}{2}\mathbb{E}_t^{CB}\left[\sum_{s=0}^{\infty}\beta^s\left(\tilde{y}_{t+s}^2 + \lambda\left(\pi_{t+s}-\varphi\right)^2\right)\right] \tag{C.6}$$

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 private 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

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

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

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

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

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

Then,

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

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

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,

$$\tilde{y}_t(\tilde{\varphi}) = \Psi_1^y\tilde{\varphi} - \Psi_2^y\mathbb{E}_t^P[\varphi] - \Psi_3^y\sum_{s=0}^{\infty}(\beta\psi^\pi)^s\mathbb{E}_t[\varepsilon_{t+s}] = \Psi_1^y\tilde{\varphi} - \Psi_2^y\mathbb{E}_t^P[\varphi] - \Psi_3^y X_t \tag{C.7}$$

$$\tilde{\pi}_t(\tilde{\varphi}) = \Psi_1^\pi\tilde{\varphi} + \Psi_2^\pi\mathbb{E}_t^P[\varphi] + \Psi_3^\pi\sum_{s=0}^{\infty}(\beta\psi^\pi)^s\mathbb{E}_t[\varepsilon_{t+s}] = \Psi_1^\pi\tilde{\varphi} + \Psi_2^\pi\mathbb{E}_t^P[\varphi] + \Psi_3^\pi X_t \tag{C.8}$$

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

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

We 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

$$\mathbb{E}_{t+1}^P[\varphi] = \mathbb{E}_t^P[\varphi] + \omega_t (\lambda\kappa)^{-1} \left(\tilde{y}_t + \lambda\kappa \left(\tilde{\pi}_t - \mathbb{E}_t^P[\varphi]\right) + \left(1+\lambda\kappa^2\right)\eta_t\right) \tag{C.9}$$

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

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

subject to

(i) The New Keynesian Phillips Curve

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

(ii) The private sector's learning structure,

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

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

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

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

$$\tilde{y}_t = \Theta_1^y \varphi - \Theta_2^y \mathbb{E}_t^P\left[\varphi\right] - \sum_{s=0}^{\infty}\Theta_{3,s}^y \mathbb{E}_t\left[\varepsilon_{t+s}\right] \tag{C.10}$$

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

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 $\mathbb{E}_t^P\left[\varphi\right]$, 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 $\Theta^y_{3,s} > \Psi^y_{3,s}$ and $\Theta^\pi_{3,s} < \Psi^\pi_{3,s}$ 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 Long-run Inflation Expectations

In our baseline model, long-run inflation expectations are anchored at zero. The reason is that the central bank's inflation target is zero. When the central bank has a nonzero inflation target, then the sensitivity of inflation expectations will depend on the central bank's reputation. To see this, consider the economy of Appendix C.2 with no discounting, so $\beta = 1$. In that case, long-run inflation expectations are

$$\mathbb{E}^P_t\left[\pi_\infty\right] = \mathbb{E}^P_t\left[\varphi\right]$$

From (C.9), long-run inflation expectations evolve as

$$\mathbb{E}^P_{t+1}\left[\pi_\infty\right] - \mathbb{E}^P_t\left[\pi_\infty\right] = \omega_t\left(\lambda\kappa\right)^{-1}\left[y_t + \lambda\kappa\left(\pi_t - \mathbb{E}^P_t\left[\pi_\infty\right]\right)\right] \tag{C.12}$$

Equation (C.12) shows that, in our model, long-run inflation expectations react to unexpected deviations to the *perceived targeting rule*. Such deviations can come from two sources. First, the central bank's forecast errors, η_t. Second, deviations from the central bank's long-run inflation target.

The sensitivity of long-run inflation expectations to innovations depends on the effective weight $\omega_t/(\lambda\kappa)$, which is non-monotone in λ. In particular, the effective weight is increasing in λ for $\lambda \leqslant \lambda^*$, where

$$\lambda^* = \sqrt{\frac{\tau_t}{\kappa^2\left(\tau_\eta + \tau\kappa^2\right)}}$$

and decreasing thereafter. In the calibrated models of Section 6, λ^* ranges from 0.4 to 1.4, whereas the calibrated values of λ range from 11.4 to 29.3. Since the calibrated λ values lie well above λ^*, the economy operates in the region where the effective weight is *decreasing* in λ: long-run inflation expectations are therefore less sensitive to innovations for larger values of λ.

This offers a complementary explanation for the seemingly anchored inflation expectations documented for the US by Malmendier and Nagel (2026). There, the authors show that a regime of low perceived persistence generates insensitive long-run inflation expectations, and that this apparent anchoring can unravel if the economy transitions to a regime of high persistence. In our model, it is instead the central bank's preferences that generate this pattern: long-run inflation expectations are insensitive in periods when the central bank is hawkish, and can become de-anchored if the central bank turns dovish.

C.4 Continuum of Strategic Central Banks

Consider an economy with a continuum of strategic central banks. The private sector correctly anticipates the banks' incentives. For simplicity, we assume there is no commitment. The learning process is still the same as (16), 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

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

$$\mathbb{E}_t^P\left[\tilde{\pi}_{t+1}\left(\tilde{\psi^\pi}\right)\right] = \int_0^1 \mathbb{E}_t\left[\tilde{\pi}_{t+1}\left(\tilde{\psi^\pi}\right)\right]\mu_t\left(\tilde{\psi^\pi}\right) d\tilde{\psi}^\pi$$

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 analog of a myopic central bank when beliefs are correctly specified.

Since there is a continuum of states $\psi^\pi \in (0,1)$, the whole belief distribution 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 $t = 0$, $\mathbb{E}_1^P[\tilde{\pi}_2] = \mathbb{E}_1^P[\psi^\pi]\,\mathbb{E}_1[\varepsilon_2]$. Therefore, all that matters for optimal policy is the central bank's reputation in $t = 1$, $\mathbb{E}_1^P[\psi^\pi]$. Using backward induction, we have

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

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

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

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

Then, we have

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

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 $\tilde{\psi}^\pi$. Since because central banks with a larger value of ψ^π, will choose higher (lower) inflation in response to a positive (negative) cost-push shock, $\tilde{\pi}_0(\psi^\pi)$ increases (decreases) with ψ^π. 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 when the cost-push shock is positive, and negative when it is negative. As a result, the covariance will be positive when the cost-push shock is positive, and negative otherwise. Since $\mathbb{E}_1[\varepsilon_2]\,\tilde{y}_1 < 0$, then Proposition 2 still holds, and there will be overreaction in the output gap and underreaction in inflation.

Figure 12a 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 12b 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 12c 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.

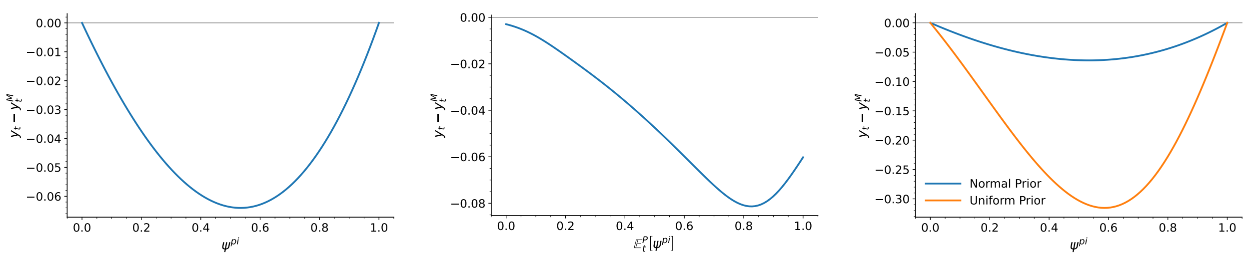

(a) Overreaction and ψ^π **(b)** Overreaction and $\mathbb{E}_t^P[\psi^\pi]$ **(c)** Normal vs Uniform priors

Figure 12: Output Gap Overreaction in the Markov equilibrium

Notes: Panel (a) reports output-gap overreaction across central-bank types; Panel (b) reports overreaction across reputation for $\lambda = 10$; Panel (c) compares the truncated-normal and uniform priors.

Figure 13 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.

Taken together, these figures show that the qualitative results of our economy do not change if we consider an economy with a continuum of strategic central banks. 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

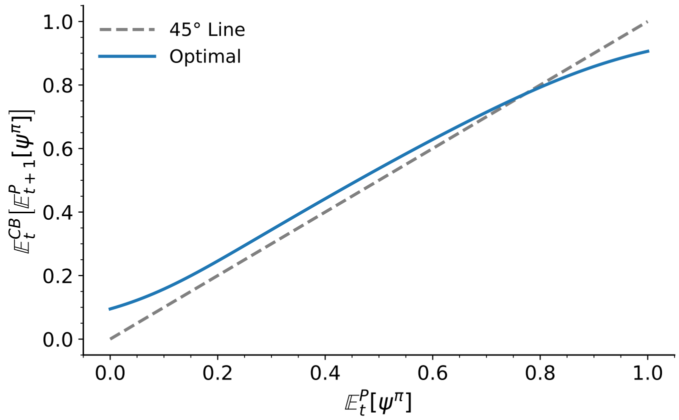

Figure 13: Reputation dynamics in the Markov equilibrium

Notes: The solid line reports the expected law of motion for reputation in the Markov equilibrium. The dashed line denotes the 45-degree line.

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.

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 (13) and (14)

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

Define $\Psi := \mathbb{E}_t^P[\psi^\pi] - \sigma\mathbb{E}_t^P[\psi^y])$, it is increasing in $\mathbb{E}_t^P[\psi^\pi]$. 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".

C.6 Inflation Bias and the Overreaction Result

The overreaction result in the reputation-for-commitment literature is often stated as a general property of optimal policy under type uncertainty. We show that it is not: it rests on the presence of an inflation bias. To make the point transparently, we work in an expositional environment in which the private sector uses the commitment and discretion policy functions to form beliefs and learns about the central bank's type via Bayes' rule, denoting by p the probability it assigns to the commitment type. A strategic committed bank can shape inflation expectations to the extent that it is believed to be committed—that is, through p.

The first-order condition of the discretionary central bank is

$$y_t^D - \overline{y} = \lambda\kappa\pi_t^D,$$

while that of the committed central bank, which for simplicity we assume optimizes under the timeless perspective, is

$$y_t^C - p\, y_{t-1} = \lambda\kappa\pi_t^C.$$

The contrast is already visible in these conditions. The discretionary bank is pulled toward $\overline{y}$—the inflation-bias temptation—whereas the committed bank is pulled toward the inherited gap $p\, y_{t-1}$, which it honors only to the degree the public believes it committed.

Using the NKPC,

$$\begin{aligned} y_t^D &= \psi^\pi\overline{y} - \psi^y\left(\varepsilon_t + \beta\mathbb{E}_t^P[\pi_{t+1}]\right) \\ \pi_t^D &= \kappa\psi^\pi\overline{y} + \psi^\pi\left(\varepsilon_t + \beta\mathbb{E}_t^P[\pi_{t+1}]\right) \end{aligned} \tag{C.13}$$

and

$$\begin{aligned} y_t^C &= p\psi^\pi y_{t-1} - \psi^y\left(\varepsilon_t + \beta\mathbb{E}_t^P[\pi_{t+1}]\right) \\ \pi_t^C &= p\kappa\psi^\pi y_{t-1} + \psi^\pi\left(\varepsilon_t + \beta\mathbb{E}_t^P[\pi_{t+1}]\right), \end{aligned} \tag{C.14}$$

where the private sector's inflation forecast averages over types,

$$\mathbb{E}_t^P[\pi_{t+1}] = p\,\mathbb{E}_t^P\left[\pi_{t+1}^C\right] + (1-p)\,\mathbb{E}_t^P\left[\pi_{t+1}^D\right]. \tag{C.15}$$

Conjecture policy functions of the form

$$\begin{aligned}
y_t^C &= -\varphi_1^{C,y}\overline{y} + p\varphi_2^{C,y}y_{t-1} - \varphi_3^{C,y}\sum_{s=0}^{\infty}\Theta_s^C\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\
\pi_t^C &= \varphi_1^{C,\pi}\overline{y} + p\varphi_2^{C,\pi}y_{t-1} + \varphi_3^{C,\pi}\sum_{s=0}^{\infty}\Theta_s^C\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\
y_t^D &= \varphi_1^{D,y}\overline{y} - p\varphi_2^{D,y}y_{t-1} - \varphi_3^{D,y}\sum_{s=0}^{\infty}\Theta_s^D\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\
\pi_t^D &= \varphi_1^{D,\pi}\overline{y} + p\varphi_2^{D,\pi}y_{t-1} + \varphi_3^{D,\pi}\sum_{s=0}^{\infty}\Theta_s^D\mathbb{E}_t\left[\varepsilon_{t+s}\right].
\end{aligned}$$

The symmetry between the two problems—the committed bank's loading on $\overline{y}$ mirrors the discretionary bank's loading on the lagged gap, and vice versa—implies

$$\varphi_2^{C,x} = \varphi_1^{D,x} \qquad \varphi_1^{C,x} = \varphi_2^{D,x} \qquad x = y, \pi. \tag{C.16}$$

Define the belief-weighted shock loading

$$\tilde{\varphi}_3^{\pi} = p\,\varphi_3^{C,\pi} + (1-p)\,\varphi_3^{D,\pi} = \mathbb{E}_t^P\left[\varphi_3^{\bullet,\pi}\right].$$

Consistency of the forecast then requires the shock to decay at a common rate across types,

$$\Theta_s^C = \Theta_s^D = (\beta\tilde{\varphi}_3^{\pi})^s \equiv \Theta^s,$$

so that, substituting into (C.15),

$$\mathbb{E}_t^P\left[\pi_{t+1}\right] = \tilde{\varphi}_1^{\pi}\overline{y} + \tilde{\varphi}_2^{\pi}\mathbb{E}_t^P\left[y_t\right] + \tilde{\varphi}_3^{\pi}\sum_{s=0}^{\infty}\Theta^s\mathbb{E}\left[\varepsilon_{t+s+1}\right].$$

Replacing this forecast into (C.13) and (C.14) gives

$$\begin{aligned}
y_t^C &= p\psi^{\pi}y_{t-1} - \beta\psi^{y}\left[\tilde{\varphi}_1^{\pi}\overline{y} + \tilde{\varphi}_2^{\pi}\left(py_t^C + (1-p)\,y_t^D\right)\right] + \psi^{y}\sum_{s=0}^{\infty}\Theta^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]\\
\pi_t^C &= p\kappa\psi^{\pi}y_{t-1} + \beta\psi^{\pi}\left[\tilde{\varphi}_1^{\pi}\overline{y} + \tilde{\varphi}_2^{\pi}\left(py_t^C + (1-p)\,y_t^D\right)\right] - \psi^{\pi}\sum_{s=0}^{\infty}\Theta^s\mathbb{E}_t\left[\varepsilon_{t+s}\right]
\end{aligned}$$

$$
\begin{aligned}
y_t^D =&\psi^\pi \overline{y} - \beta\psi^y \left[\tilde{\varphi}_1^\pi \overline{y} + \tilde{\varphi}_2^\pi \left(p y_t^C + (1-p)\, y_t^D\right)\right] - \psi^y \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t \left[\varepsilon_{t+s}\right] \\
\pi_t^D =&\kappa\psi^\pi \overline{y} + \beta\psi^\pi \left[\tilde{\varphi}_1^\pi \overline{y} + \tilde{\varphi}_2^\pi \left(p y_t^C + (1-p)\, y_t^D\right)\right] + \psi^\pi \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t \left[\varepsilon_{t+s}\right].
\end{aligned}
$$

Matching the shock terms yields a common loading across types,

$$
\varphi_3^{C,y} = \varphi_3^{D,y} = \frac{\psi^y}{1+\beta\tilde{\varphi}_2^\pi \psi^y}, \qquad \varphi_3^{C,\pi} = \varphi_3^{D,\pi} = \frac{\psi^\pi}{1+\beta\tilde{\varphi}_2^\pi \psi^y},
$$

and matching the remaining coefficients, using (C.16),

$$
\begin{aligned}
\varphi_1^{C,y} =&\beta\psi^y \left[\tilde{\varphi}_1^\pi + \tilde{\varphi}_2^\pi \left(p\varphi_1^{C,y} - (1-p)\,\varphi_2^{C,y}\right)\right] \\
\varphi_2^{C,y} =&p\psi^\pi - \beta\psi^y \tilde{\varphi}_2^\pi \left(p\varphi_2^{C,y} - (1-p)\,\varphi_1^{C,y}\right) \\
\varphi_1^{C,\pi} =&\beta\psi^\pi \left[\tilde{\varphi}_1^\pi + \tilde{\varphi}_2^\pi \left((1-p)\,\varphi_2^{C,y} - p\varphi_1^{C,y}\right)\right] \\
\varphi_2^{C,\pi} =&p\kappa\psi^\pi + \beta\psi^\pi \tilde{\varphi}_2^\pi \left(p\varphi_2^{C,y} - (1-p)\,\varphi_1^{C,y}\right),
\end{aligned}
\tag{C.17}
$$

where

$$
\tilde{\varphi}_1^\pi = p\varphi_1^{C,\pi} + (1-p)\,\varphi_2^{C,\pi}, \qquad \tilde{\varphi}_2^\pi = p\varphi_2^{C,\pi} + (1-p)\,\varphi_1^{C,\pi}.
$$

This pins down $\left\{\varphi_1^{C,y}, \varphi_2^{C,y}, \varphi_1^{C,\pi}, \varphi_2^{C,\pi}\right\}$.

We verify that the system nests discretion and commitment at the two extremes of belief.

Full credibility ($p \to 1$). Here $\tilde{\varphi}_1^\pi = \varphi_1^{C,\pi}$ and $\tilde{\varphi}_2^\pi = \varphi_2^{C,\pi}$. The first and third equations of (C.17) become

$$
\begin{aligned}
\varphi_1^{C,y} =&\beta\psi^y \left[\varphi_1^{C,\pi} + \varphi_2^{C,\pi}\varphi_1^{C,y}\right], \\
\varphi_1^{C,\pi} =&\beta\psi^\pi \left[\varphi_1^{C,\pi} - \varphi_2^{C,\pi}\varphi_1^{C,y}\right].
\end{aligned}
$$

Combining them,

$$
(1-\beta)\,\varphi_1^{C,\pi} + \kappa\varphi_1^{C,y} = -(1-\kappa)\,\varphi_2^{C,\pi}\varphi_1^{C,y},
$$

and since every coefficient is positive, the only solution is $\varphi_1^{C,y} = \varphi_1^{C,\pi} = 0$: a fully credible bank places no weight on $\overline{y}$. The second and fourth equations,

$$
\varphi_2^{C,y} =\psi^\pi - \beta\psi^y \varphi_2^{C,\pi}\varphi_2^{C,y},
$$

$$\varphi_2^{C,\pi} = \kappa\psi^\pi + \beta\psi^\pi\varphi_2^{C,\pi}\varphi_2^{C,y},$$

combine to $\varphi_2^{C,\pi} = \kappa\varphi_2^{C,y} + \beta\varphi_2^{C,\pi}\varphi_2^{C,y}$, which admits a nonzero solution. The policy thus loads only on the inherited gap and coincides with commitment.

No credibility ($p \to 0$). Here $\tilde{\varphi}_1^\pi = \varphi_2^{C,\pi}$ and $\tilde{\varphi}_2^\pi = \varphi_1^{C,\pi}$. The economy is believed to operate under discretion with certainty, and the solution collapses to the discretionary baseline.

For any $p \in (0,1)$, the solution is

$$\begin{aligned}
y_t^C &= -\varphi_1^{C,y}\overline{y} + \varphi_2^{C,y}\, p\, y_{t-1} - \alpha\psi^y \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\
\pi_t^C &= \varphi_1^{C,\pi}\overline{y} + \varphi_2^{C,\pi}\, p\, y_{t-1} + \alpha\psi^\pi \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t\left[\varepsilon_{t+s}\right], \\
y_t^D &= \varphi_2^{C,y}\overline{y} - \varphi_1^{C,y}\, p\, y_{t-1} - \alpha\psi^y \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t\left[\varepsilon_{t+s}\right] \\
\pi_t^D &= \varphi_2^{C,\pi}\overline{y} + \varphi_1^{C,\pi}\, p\, y_{t-1} + \alpha\psi^\pi \sum_{s=0}^{\infty} \Theta^s \mathbb{E}_t\left[\varepsilon_{t+s}\right],
\end{aligned}$$

where

$$\alpha = \frac{1}{1+\beta\tilde{\varphi}_2^\pi\psi^y} < 1.$$

(The shock loading $\varphi_3^{\bullet,x}$ derived above equals $\alpha\psi^x$; we write it as α here to isolate the dampening that beliefs induce.)

Three features of this solution drive the result.

1. **The impact response to shocks is identical across types.** Both load on the shock sum with the same coefficient $\alpha\psi^x$. Type uncertainty leaves the contemporaneous response untouched; any signaling content of policy must therefore operate through the *dynamic* response, carried by the $p\, y_{t-1}$ term.

2. **Inflation bias enters with opposite signs across types.** For the committed bank, $\overline{y}$ is the temptation to inflate—something it must visibly resist. For the discretionary bank, the same $\overline{y}$ acts merely as a cost-push shock to be accommodated. The two types pull $\overline{y}$ in opposite directions, which is exactly what makes it a dimension along which they separate.

3. **Commitment enters with opposite signs for the same reason.** The inherited gap y_{t-1} that the committed bank honors is, to the discretionary bank, just another cost-push term it would rather offset.

A shock opens a window in which the bank can signal its type, and a strategic committed bank exploits it. But the *direction* of that signal is not a primitive of reputation building—it is inherited from whatever distinguishes the types, and that, in turn, depends on the inflation bias.

Consider first the case studied in the literature, $\overline{y} > 0$. The discretionary type's defining vice is the temptation to inflate, so the committed type separates by doing the opposite: resisting accommodation and thus *overreacting* to a cost-push shock relative to the myopic benchmark. By the second observation, this overreaction is simply the committed bank treating its own inflation-bias term as if it were a cost-push shock to be leaned against. The result is a levels phenomenon, driven by $\overline{y}$ rather than by the stabilization problem.

Now let $\overline{y} \to 0$. The inflation bias vanishes, and with it the motive just described: there is nothing left to disown on the inflation-temptation margin. Yet the types still differ—now only through the *stabilization bias*. Unable to commit to future policy, the discretionary bank responds too forcefully to a cost-push shock, while the committed bank smooths the response over time. The separating action therefore reverses: to signal commitment, the bank must now react *more slowly* than the discretionary type would—it must *underreact*. The mechanism compounds dynamically, since under uncertainty the discretionary bank treats the inherited gap $p\, y_{t-1}$ as an additional cost-push term (third observation) and responds even more aggressively, widening the gap the committed type signals against.

The presence of an inflation bias is therefore essential to the overreaction result in the existing literature (King et al., 2008; Lu et al., 2016). It is not a general property of reputation building under commitment uncertainty, but a consequence of the particular distortion the public is learning about. Absent an inflation bias, the same separation logic delivers the opposite prescription.

D Proofs and Robustness Checks for Empirical Section

D.1 Proof of Proposition 1

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

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

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

D.2 Disagreement about $\mathbb{E}_t^P\left[\eta_{t+k}\right]$

Suppose there is disagreement about the future path of η_{t+k}. Suppose the dispersion in the forecasts of η_{t+k} does not depend on the horizon k. The population coefficient of (21) is

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

A good reputation, a decrease in $\mathbb{E}_t^P\left[\psi^\pi\right]$ has two effects. First, it decreases $-\frac{\mathbb{E}_t^P\left[\psi^y\right]}{\mathbb{E}_t^P\left[\psi^\pi\right]}$. Second, it also decreases ω_t. The two effects go in opposite directions. However, as long as reputation is not good enough, the first effect dominates.

D.3 Demand Shocks

To allow for situations where the Divine Coincidence does not hold, suppose the central bank follows the triple mandate in (27)

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

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.[46] In this environment, the

[46] 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).

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* (27) *and there are only demand shocks. Define* $\widehat{\psi} := \frac{\varphi\sigma}{1+\lambda\kappa^2+\sigma^2\varphi}$. *Then, the estimand of* $\gamma_{2,t}$ *in equation* (21) *can be positive or negative, and increasing in* $\mathbb{E}_t^P\left[\widehat{\psi}\right]$

Proof. The first-order condition with respect to output gap, $\tilde{y}_t$, of a central bank with parameters λ and φ is

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

Define $\widehat{\psi} = \frac{\varphi\sigma}{1+\lambda\kappa^2+\varphi\sigma^2}$, $\widehat{\psi}^y := \frac{\lambda\kappa}{1+\lambda\kappa^2+\varphi\sigma^2}$ and $\widehat{\psi}^\pi := \frac{1}{1+\lambda\kappa^2+\varphi\sigma^2}$. Our new measure of reputation will be $\mathbb{E}_t^P\left[\widehat{\psi}\right]$. As elsewhere, a lower $\mathbb{E}_t^P\left[\widehat{\psi}\right]$ signals a more hawkish central bank. 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 $k = 1$

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

For an arbitrary s, we have

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

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

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 (21) 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 (D.2) and (D.3)

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

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. ∎

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-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.

D.4 Belief Heterogeneity

Suppose 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 6. *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. 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

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

and

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

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

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

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 (D.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 $\gamma_{2,t}$. ■

Lemma 6 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 4.

D.5 Model Heterogeneity

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.

Under this assumption, Lemma 6 holds: the estimand of $\gamma_{2,t}$ is a weighted average of the individual estimands, $\gamma_{2,t}^i = -\frac{\mathbb{E}_t^i[\psi^y]}{\mathbb{E}_t^i[\psi^\pi]}$, placing more weight on forecaster with a flatter NKPC. Since beliefs over λ are still homogeneous, the comparative statics from Lemma 3 and Lemma 4 continue to hold. A shift in beliefs toward a more hawkish central bank makes each $\gamma_{2,t}^i$ more negative, and therefore their weighted average, $\gamma_{2,t}$, shifts in the same direction.

D.6 Survey of Professional Forecasters

We repeat the empirical exercise using the Survey of Professional Forecasters (SPF), which collects individual forecasts at the quarterly frequency. In particular, we collect the quarterly estimates of $\gamma_{2,t}$ from (21). The correlation coefficient between both series is 0.60, showing that both series capture similar time-series variation in the private sector's beliefs.[47] Figure 14 plots the estimates from both datasets. Qualitatively, both series look similar after the dot-com bubble.

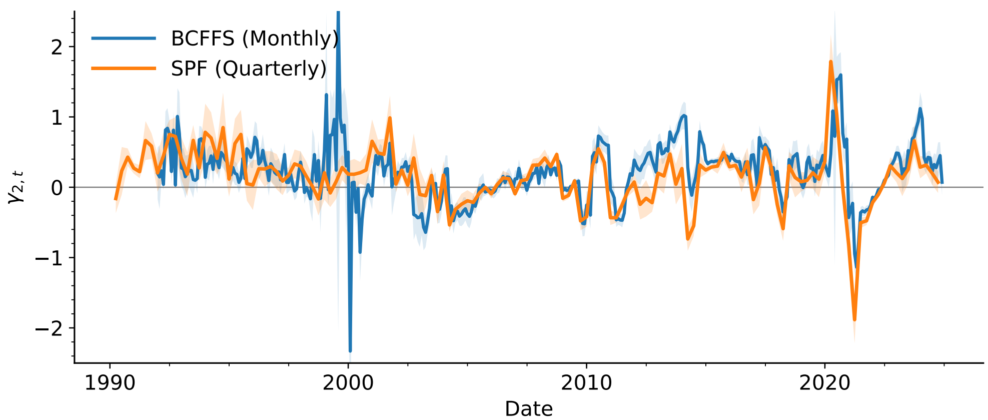

Figure 14: Reputation estimates from BCFF and SPF

Figure 15 plots the impulse response to an oil price news cost-push shock from Känzig (2021). Consistent with the estimates from BCFF, a more hawkish reputation dampens the pass-through of cost-push shocks on inflation expectations. However, the point estimates suggest that when reputation is high, a positive cost-push shock may even reduce inflation expectations, which is implausible. We interpret this as evidence that while the qualitative pattern—stronger reputation leads to lower pass-through—is robust across datasets, the quantitative magnitudes in SPF should be interpreted with caution.

Figure 16 plots the impulse response to an unexpected monetary tightening. We obtain a quarterly time series for monetary surprises by aggregating the monthly series. We find no effect of a monetary policy shock on reputation. This contrasts with the significant effect found in the estimates using BCFF data. One potential explanation is that aggregating

[47] We take the quarterly mean of the monthly estimates from BCFF.

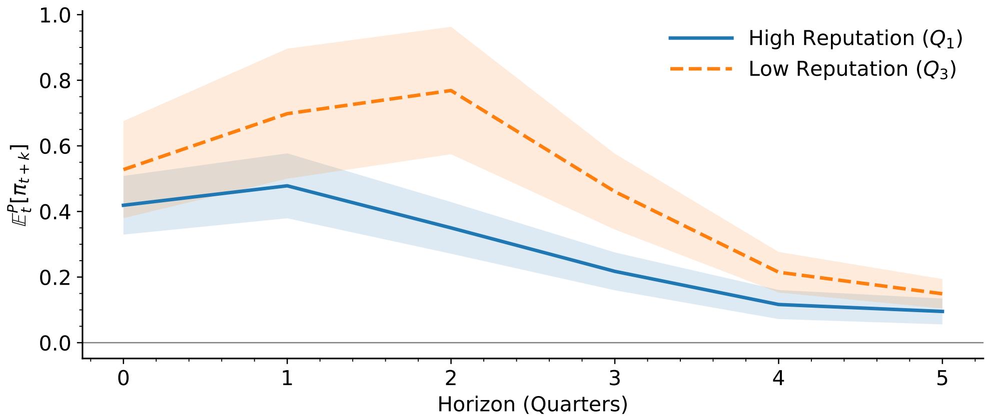

Figure 15: Impulse response of inflation expectations to an oil price news shock for high and low reputation.

Notes: Shaded areas denote 90% confidence intervals.

monthly monetary surprises to the quarterly frequency may obscure the learning dynamics documented in the main text, particularly if belief updating occurs primarily in response to individual FOMC announcements rather than cumulative quarterly policy changes.

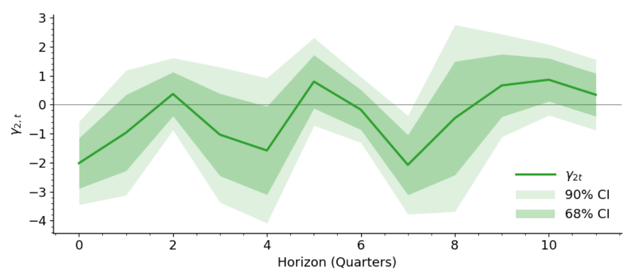

Figure 16: Impulse response to a monetary tightening.

Notes: Dark and light shaded areas denote 68% and 90% confidence intervals, respectively.

Taken together, these findings support that the facts we document are consistent across datasets, and time-varying perceptions about the central bank's hawkishness are present in the data.

D.7 Commitment

In the baseline model we assume the central bank optimizes period-by-period under discretion. Suppose instead that the bank operates under commitment but is still myopic. That is, the bank optimizes under the timeless perspective, but does not internalize the effects of its actions on the private sector's beliefs. The first-order condition is, then

$$\tilde{y}_t - y_{t-1} + \kappa\lambda\tilde{\pi}_t = 0 \tag{D.5}$$

Lemma 7. *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. 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

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

and

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

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. In addition, it is positive at zero. Since the left 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. ■

Lemma 7 highlights two differences relative to the myopic discretionary 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

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

By the Frisch-Waugh-Lovell Theorem, $\gamma_{2,t}$ coincides with the estimand in (21), so Proposition 1 continues to hold. In practice, the correlation coefficient between the commitment-based and baseline estimates is 0.84, confirming the robustness of our measure to the assumed policy behavior.

D.8 Inertia

Suppose that past inflation feeds into current inflation, so that the NKPC exhibits inertia. In this case, the New Keynesian Phillips Curve (7) becomes

$$\tilde{\pi}_t = \kappa \tilde{y}_t + \delta \pi_{t-1} + \beta \mathbb{E}_t^P [\tilde{\pi}_{t+1}] + \varepsilon_t \quad \delta + \beta < 1.^{48}$$

Lemma 8. *Under inertia, the k-periods ahead forecast is*

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

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

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

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

Conjecture the Central Bank follows a linear rule

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

Replacing into both equations

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

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

and

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

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 $\alpha\beta\mathbb{E}_t^P[\psi^\pi] < 1$, which is needed for a well-defined policy function. To see this, define $\tilde{\alpha} := \alpha\beta\mathbb{E}_t^P[\psi^\pi]$. It solves

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

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. ■

Lemma 8 highlights two differences relative to the 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.

Empirically, this implies we can recover reputation in the presence of inertia by estimating

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

By the Frisch-Waugh-Lovell theorem, the coefficient γ_{2t} coincides with the estimand in (21), so Proposition 1 continues to hold. In practice, the correlation coefficient between the inertia and baseline estimates is 0.84, confirming the robustness of our measure to the presence of inertia.

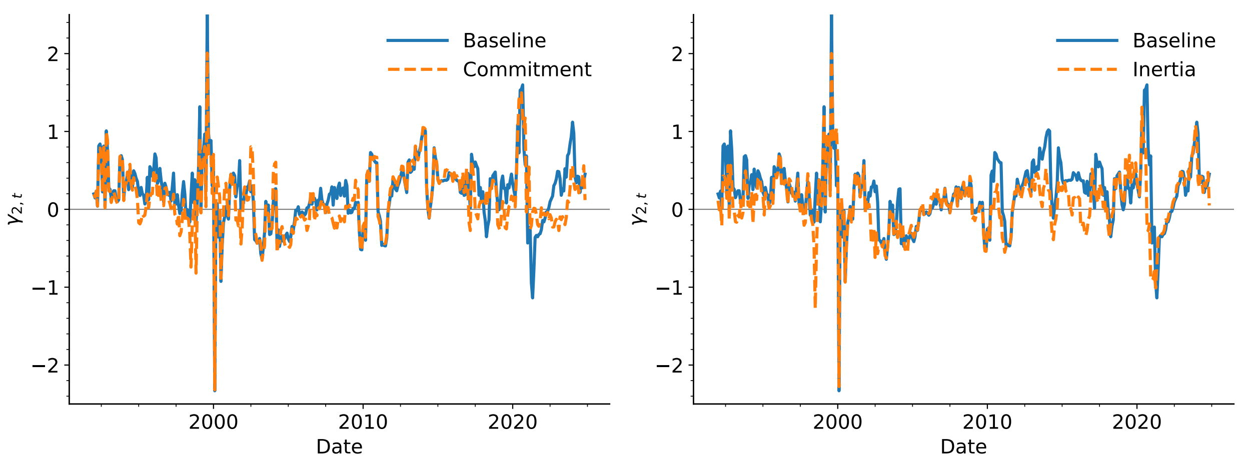

(a) Commitment　　**(b)** Inertia

Figure 17: Robustness of the monthly reputation estimate $\gamma_{2,t}$ to commitment and inflation inertia

Notes: Panel (a) compares the baseline and commitment estimates of $\gamma_{2,t}$; Panel (b) compares the baseline and inertia estimates.

D.9 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 (13) and (14). To address this issue, we re-estimate our baseline specification (21) after excluding forecasts that assume the policy rate will remain at the ZLB—that is, those for which $\mathbb{E}_t^i [i_{t+k}] = 0$— during the 2009-2015 period, when the constraint was binding.[49] Both series display a correlation coefficient of 0.78, confirming robustness with respect to the ZLB. What matters for our measure is whether forecasters *expect* to be at the ZLB, not whether the economy *currently* is. For this reason, both estimates coincide.

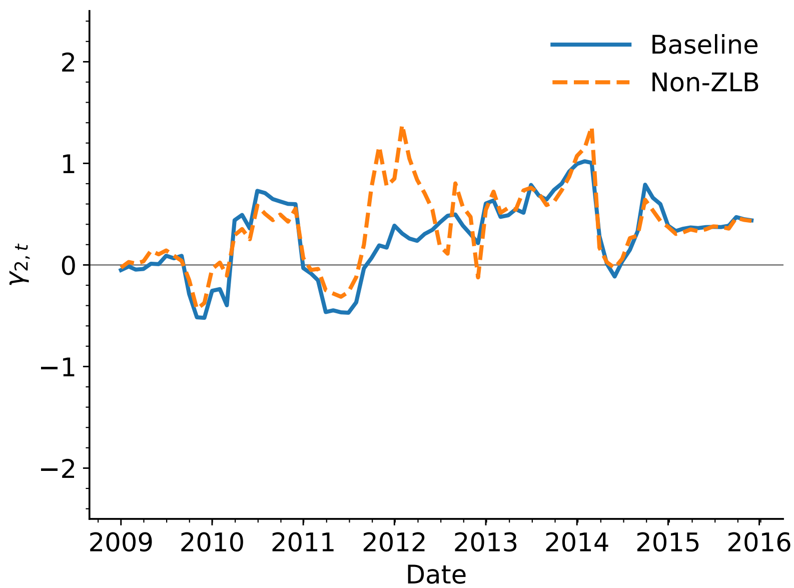

Figure 18: Robustness of the monthly reputation estimate $\gamma_{2,t}$ to zero lower bound constraints

Notes: The figure compares the baseline estimate of $\gamma_{2,t}$ with estimates excluding forecasts with expected policy rates below 20 basis points during 2009-2015.

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

D.10 Strategic Central Bank

A strategic central bank of type λ has policies

$$y_t\left(\lambda, \left\{\mathbb{E}_t^{CB}\left[\varepsilon_{t+s}\right]\right\}_{t+s}\right) \quad \pi_t\left(\lambda, \left\{\mathbb{E}_t^{CB}\left[\varepsilon_{t+s}\right]\right\}_{t+s}\right)$$

Forecasters differ in their forecasts of shocks, $\{\mathbb{E}_t^i\left[\varepsilon_{t+s}\right]\}$ and have the same beliefs over the central bank's type, $\mu_t\left(\lambda\right)$. Assume further that, conditional on the type, forecasts are 'certainty equivalent', namely

$$\begin{aligned}
\mathbb{E}_t^i\left[y_{t+k}\left(\lambda\right)\right] =& \mathbb{E}_t^i\left[y_{t+k}\left(\lambda, \left\{\mathbb{E}_{t+s}^{CB}\left[\varepsilon_{t+k+s}\right]\right\}_{s\geqslant 0}\right)\right] = y_{t+k}\left(\lambda, \left\{\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right\}_{s\geqslant 0}\right)\\
\mathbb{E}_t^i\left[\pi_{t+k}\left(\lambda\right)\right] =& \mathbb{E}_t^i\left[\pi_{t+k}\left(\lambda, \left\{\mathbb{E}_{t+s}^{CB}\left[\varepsilon_{t+k+s}\right]\right\}_{s\geqslant 0}\right)\right] = \pi_{t+k}\left(\lambda, \left\{\mathbb{E}_t^i\left[\varepsilon_{t+k+s}\right]\right\}_{s\geqslant 0}\right)
\end{aligned}$$

Then

$$\begin{aligned}
\mathbb{E}_t^i\left[y_{t+k}\right] &= \int_0^\infty \mathbb{E}_t^i\left[y_{t+k}\left(\lambda\right)\right]\mu_t\left(\lambda\right)d\lambda\\
\mathbb{E}_t^i\left[\pi_{t+k}\right] &= \int_0^\infty \mathbb{E}_t^i\left[\pi_{t+k}\left(\lambda\right)\right]\mu_t\left(\lambda\right)d\lambda
\end{aligned}$$

We estimate the following regression in the cross-section

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

Then

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

We proceed with each term separately

$$\begin{aligned}
Cov\left(\mathbb{E}_t^i\left[y_{t+k}\right], \mathbb{E}_t^i\left[\pi_{t+k}\right]\right) &= \int_i \mathbb{E}_t^i\left[y_{t+k}\right]\mathbb{E}_t^i\left[\pi_{t+k}\right]dF_i - \int_i \mathbb{E}_t^i\left[y_{t+k}\right]dF_i \int_i \mathbb{E}_t^i\left[\pi_{t+k}\right]dF_i\\
&= \int_i\left(\int_0^\infty \mathbb{E}_t^i\left[y_{t+k}\left(\lambda\right)\right]\mathbb{E}_t^i\left[\pi_{t+k}\left(\lambda\right)\right]\mu_t\left(\lambda\right)d\lambda\right)dF_i < 0\\
Var\left(\mathbb{E}_t^i\left[\pi_{t+k}\right]\right) &= \int_i \mathbb{E}_t^i\left[\pi_{t+k}\right]^2 dF_i\\
&= \int_i\left(\int_0^\infty \mathbb{E}_t^i\left[\pi_{t+k}\left(\lambda\right)\right]^2\mu_t\left(\lambda\right)d\lambda\right)dF_i > 0
\end{aligned}$$

Then

$$\gamma_{2,t} = \int_i \alpha_t^i \left(\int_0^\infty \omega_t^i(\lambda) \frac{\mathbb{E}_t^i[y_{t+k}(\lambda)]}{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]} d\lambda \right) dF_i$$

where

$$\omega_t^i(\lambda) = \frac{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2 \mu_t}{\int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2 \mu_t d\lambda} \qquad \alpha_t^i = \frac{\int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2 \mu_t d\lambda}{\int_i \left(\int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2 \mu_t d\lambda \right) dF_i}$$

Fix i, $\frac{\mathbb{E}_t^i[y_{t+k}(\lambda)]}{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]}$ is negative, and the magnitude decreases with λ: a central bank with higher priority for inflation stability. Now let's focus on

$$\omega_t^i(\lambda) = \frac{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2 \mu_t}{\int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)\mu_t]^2 \mu_t d\lambda}$$

Suppose $\mu_t' \gtrsim_{\text{FOSD}} \mu_t$, then

$$\int_0^x \mu_t'(\lambda)\, d\lambda \leqslant \int_0^x \mu_t(\lambda)\, d\lambda.$$

Since $\mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2$ is decreasing in λ, we can integrate by parts to obtain

$$\begin{aligned}\int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2\, \mu_t'(\lambda)\, d\lambda &= \int_0^\infty \left(-\frac{d\, \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2}{d\lambda} \right) \left(\int_0^\lambda \mu_t'(s)\, ds \right) d\lambda \\ &\leqslant \int_0^\infty \left(-\frac{d\, \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2}{d\lambda} \right) \left(\int_0^\lambda \mu_t(s)\, ds \right) d\lambda = \int_0^\infty \mathbb{E}_t^i[\pi_{t+k}(\lambda)]^2\, \mu_t(\lambda)\, d\lambda.\end{aligned}$$

Define

$$\omega_t^i(\lambda) := \frac{h_t^i(\lambda)\, \mu_t(\lambda)}{\int_0^\infty h_t^i(s)\, \mu_t(s)\, ds}, \qquad \omega_t^{i'}(\lambda) := \frac{h_t^i(\lambda)\, \mu_t'(\lambda)}{\int_0^\infty h_t^i(s)\, \mu_t'(s)\, ds},$$

and the ratio

$$R^i(\lambda) := \frac{\omega_t^{i'}(\lambda)}{\omega_t^i(\lambda)} = c\, \frac{\mu_t'(\lambda)}{\mu_t(\lambda)}, \qquad c := \frac{\int_0^\infty h_t^i(s)\, \mu_t(s)\, ds}{\int_0^\infty h_t^i(s)\, \mu_t'(s)\, ds} < 1.$$

Since

$$\int_0^\infty \omega_t^{i'}(\lambda)\, d\lambda = \int_0^\infty \omega_t^i(\lambda)\, d\lambda = 1,$$

it follows that $R^i(\lambda)$ must be greater than one for some values of λ and smaller than one for others.

Now, assume $\mu_t' \gtrsim_{\text{MLR}} \mu_t$ (which is stronger than FOSD). Then $R^i(\lambda)$ is increasing in λ. Because both weights integrate to one, they must cross. Taken together, $R^i(\lambda)$ crosses exactly once (from below to above).

Hence, if $\mu_t' \gtrsim_{\mathrm{MLR}} \mu_t$, then

$$\omega_t^{i'} \gtrsim_{\mathrm{FOSD}} \omega_t^i, \qquad \text{i.e.} \qquad \int_0^x \omega_t^{i'}(\lambda)\, d\lambda \leqslant \int_0^x \omega_t^i(\lambda)\, d\lambda.$$

Now let's focus on

$$\gamma_{2t}^i = \int_0^\infty \omega_t^i(\lambda) \frac{\mathbb{E}_t^i[y_{t+k}(\lambda)]}{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]} d\lambda = \int_0^\infty \left(1 - \int_0^\lambda \omega_t^i(x)\, d\lambda\right)\left(-\frac{d}{d\lambda}\frac{\mathbb{E}_t^i[y_{t+k}(\lambda)]}{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]}\right) d\lambda$$

Then, since $\frac{\mathbb{E}_t^i[y_{t+k}(\lambda)]}{\mathbb{E}_t^i[\pi_{t+k}(\lambda)]}$ is negative and its magnitude increasing in λ, we have that if $\mu_t' \gtrsim_{MLR} \mu_t$ then $\gamma_{2t}^{i'} < \gamma_{2t}^i$. Finally, we have

$$\gamma_{2t} = \int_i \alpha_t^i \gamma_{2t}^i dF_i = \mathbb{E}_t\left[\gamma_{2,t}^i\right] + Cov_t\left(\alpha_t^i, \gamma_{2t}^i\right)$$

where the second term is positive: forecasters who forecasts larger shocks are given more weight (larger α_t^i), and also they have a worse estimate of reputation, as the allocation becomes less dependent on the central bank's preferences and more on the economic conditions (smaller magnitude of γ_{2t}^i). However, with better reputation, shocks have a smaller pass-through on forecasts. Therefore, when $\mu_t' \gtrsim_{MLR} \mu_t$, under some general conditions, γ_{2t} becomes more negative.

D.11 Perceived Taylor Rule Coefficients

We consider two perceived Taylor rules. The first regresses forecast interest rates on forecast inflation:

$$\mathbb{E}_t^i[i_{t+k}] = \alpha_{1t}^\pi + \alpha_{2t}^\pi \mathbb{E}_t^i[\pi_{t+k}] + u_{i,t+k} \tag{D.6}$$

The second regresses forecast interest rates on forecast output:

$$\mathbb{E}_t^i[i_{t+k}] = \alpha_{1t}^y + \alpha_{2t}^y \mathbb{E}_t^i[y_{t+k}] + u_{i,t+k} \tag{D.7}$$

We establish two results relating these coefficients to reputation.

Lemma 9. *The perceived Taylor rule coefficient on inflation $\alpha_{2,t}^\pi$ is U-shaped in $\mathbb{E}_t^P[\psi^\pi]$: it is decreasing when reputation is strong (low $\mathbb{E}_t^P[\psi^\pi]$) and increasing when reputation is weak (high $\mathbb{E}_t^P[\psi^\pi]$).*

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

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

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

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

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

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

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

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

because we assumed reputation was constant between forecasters, so $\mathbb{E}_t^P\left[\psi^{\pi}\right]$ is a constant. Assume there is no heteroskedasticity in the dispersion of forecasts, so that $Var\left(\mathbb{E}_t^i\left[\varepsilon_{t+k}\right]\right)$ does not depend on the horizon k. Therefore, we have

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

Then, we have

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

And the estimand of the taylor rule coefficient is

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

The first term is increasing in $\mathbb{E}_t^P[\psi^\pi]$, whereas the second one is decreasing. Then, $\alpha_{2,t}$ is strictly convex and U-shaped.

Under heteroskedascitiy of forecasts errors, the dispersion of the forecast changes with the horizon. This adds an extra term to the estimand of $\alpha_{2,t}^\pi$ 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. ■

Lemma 10. *In the economy of Appendix C.1, where the central bank has a triple mandate, prices are fixed, and only demand shocks are present, the perceived Taylor rule coefficient on output is*

$$\alpha_{2,t}^y = \frac{1}{\mathbb{E}_t^P[\psi^y]} - \sigma \geqslant 0$$

and is increasing in reputation.

Proof. Under these assumptions, the $k-$periods ahead forecast of output gap and interest rates is

$$\mathbb{E}_t^i[y_{t+k}] = \mathbb{E}_t^P[\psi^y]\,\mathbb{E}_t^i[X_{t+k}] \tag{D.8}$$

$$\mathbb{E}_t^i[i_{t+k}] = \mathbb{E}_t^P\left[\psi^i\right]\mathbb{E}_t^i[X_{t+k}] \tag{D.9}$$

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

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

Then the coefficient becomes

$$\alpha_{2,t}^\pi = \frac{1}{\mathbb{E}_t^P[\psi^y]} - \sigma \geqslant 0$$

which is increasing in the central bank's reputation. ■

E Appendix for Quantitative Analysis

This section collects the technical details of the estimation procedure and quantitative analysis, and additional robustness exercises performed in this paper.

E.1 Estimation Procedure and Quantitative Simulation

Simulated method of moments. We estimate the model by SMM under the myopic policy. For a candidate vector θ, we simulate the discretized model, compute each targeted moment on every replication, and average across replications. The estimator solves

$$\hat{\theta} = \arg\min_{\theta} \ \frac{1}{2}\, \mathbf{g}(\theta)' W\, \mathbf{g}(\theta), \tag{E.1}$$

where $g_j(\theta) = \bar{m}_j(\theta) - m_j^{\text{data}}$, $\bar{m}_j(\theta)$ is the cross-replication mean of moment j, and m_j^{data} is its empirical counterpart. We set $W = I$. With a small number of targets, the off-diagonal elements of an estimated moment covariance matrix are poorly determined, and the demeaning convention described below already places the targets on comparable scales. Equal weighting therefore trades a known efficiency loss for a criterion that does not depend on a first-stage estimate.

The estimated vector is $\theta = (\lambda, \sigma_{\nu^{\varepsilon}}, \sigma_{\eta}, \tau)$: the weight on inflation in the central bank's loss, the standard deviations of the cost-push innovation and of the private sector's forecast error, and the precision of beliefs. The remaining parameters are held fixed. The persistence of the cost-push process is set to $\rho_{\varepsilon} = 0.723$, the initial location of beliefs to $\overline{\psi}_0 = 0.6$, and the New Keynesian block to standard values ($\beta = 0.99$, unit intertemporal elasticity of substitution, inverse Frisch elasticity of 5, and $\alpha = 0.25$). No feature of the initial belief distribution is estimated: with constant precision, τ governs the belief process throughout rather than its starting point.

The minimization uses a constrained interior-point algorithm (`fmincon` in MATLAB) with forward finite differences, started from a common initial guess and from three additional points drawn uniformly from the parameter bounds, retaining the best of the four. Every evaluation of the criterion reseeds the random number generator, so all candidate vectors face the same underlying shocks. These common random numbers make the criterion a deterministic and smooth function of θ, which is what permits a gradient-based method on a simulated objective; they also require the restart points to be drawn before the search rather than inside it.

Matched moments. The estimation sample is the Great Moderation, 1992-Q1 to 2019-Q4. Second moments are demeaned around fixed constants rather than around the mean of the estimation window: in the data, around long-run means computed on the full 1960–2024 sample; in the model, around each replication's own mean. The slope of the Phillips curve is swept rather than estimated. We report a grid of four specifications, $\kappa \in \{0.20, 0.40\}$ crossed with two moment sets, estimating the same four parameters in each. The small set targets $\sigma(y)$, $\sigma(\pi)$ and $\mathrm{Corr}(y, \pi)$ together with the unconditional response of inflation expectations to a cost-push shock. The large set adds the first autocorrelation of inflation and the standard deviation of γ_2, and replaces the unconditional response with the two responses conditional on credibility. Targets enter the criterion element by element: a scalar target contributes one row of $\mathbf{g}$, and a response of length L contributes L.

The responses are computed inside the same simulation that delivers the second moments. The simulated panel serves as the control arm; a treatment arm adds the shock to the realized innovations and re-runs the belief recursion alone, holding every other shock fixed. Differencing the two arms replication by replication and dividing by the size of the perturbation yields a per-unit response at a cost low enough for the responses to be evaluated inside the criterion. For the cost-push experiment the perturbation is one estimated standard deviation σ_{ν^ε}, decaying at ρ_ε thereafter, so the scale of the experiment moves with θ. Because belief updating is nonlinear in the shock, the resulting object is a difference quotient at that perturbation size rather than a derivative.

Two conventions govern the matched series. First, $\overline{\psi}_t$ is predetermined with respect to the date-t innovation, so beliefs do not move on impact. Second, the matched object is the term structure of the private sector's inflation forecast formed at the shock date, $\mathbb{E}_t^P[\pi_{t+h}]$ swept over the forecast horizon h and annualized by a factor of four. We compute the subjective forecasts as expectations over the belief distribution rather than as the forecast function evaluated at the mean belief, which retains the Jensen term generated by the nonlinearity of the reputation mapping. The empirical counterpart is available at two quarterly horizons, obtained from monthly estimates by averaging within quarters, and the model is matched at $h = 0, 1$ only. In the large moment set, the responses are split by credibility using the quartiles of γ_2 at the shock date in the control arm, comparing the most and least credible quartiles.

Simulation design and burn-in. The quantitative exercises use different simulation panels according to their computational and statistical requirements; Table E.1 summarizes their design. All stochastic processes are discretized using the Rouwenhorst method on 101-state grids, preserving the unconditional variance and persistence of the underlying AR(1) processes. For estimation, each evaluation of the SMM criterion simulates $S = 2{,}500$

independent paths, initialized at $\overline{\psi}_0$ and run for $T_{\text{burn}} + T$ quarters, with $T_{\text{burn}} = 300$ and $T = 112$ matching the length of the estimation sample. The reporting exercises are simulated only once and therefore use 10,000 replications. The welfare and delegation counterfactuals retain 500 quarters so that β^T is negligible and truncation of the discounted welfare sum does not affect comparisons across regimes. Within each panel, a fixed seed ensures that the myopic, delegated, and optimal regimes face identical shock realizations, so differences in outcomes reflect only differences in policy.

The burn-in removes dependence on initial conditions, which is particularly important for the belief state. Because the Kalman gain ω_t depends on the history of signals, paths initialized at a common prior share an early transient, causing cross-replication averages to differ from the model's stationary moments. The required burn-in therefore depends on the persistence of reputation. The estimation and welfare panels discard 300 quarters. This is sufficient for counterfactual comparisons because any remaining transient is largely common across regimes and therefore cancels when differences are computed. By contrast, the stationary-distribution exercise uses a 1,000-quarter burn-in because it estimates the *level* of long-run reputation, for which a residual transient would bias the result toward the initial belief. Across the four estimation cells, the half-lives of reputation are 61, 91, 8, and 86 quarters, so 300 quarters correspond to roughly three to five half-lives in the three more persistent cells, compared with eleven to sixteen under a 1,000-quarter burn-in.

The convergence exercise uses no burn-in by construction. It initializes reputation at a chosen value and measures the transition toward its ergodic level, so discarding the initial observations would remove the object of interest. Instead, the exogenous processes are initialized at their stationary distributions, preventing an artificial volatility ramp-up while leaving the initial belief unchanged.

Exercise	**Replications**	**Quarters kept**	**Burn-in**
SMM estimation	$2,500$	112	300
Welfare and delegation counterfactuals	$10,000$	500	300
Stationary distributions	10,000	800	1,000
Convergence transitions	10,000	800	0

Note: Great Moderation sample: 1992-Q1 to 2019-Q4 (112 obs.). Replications and quarters are per panel; the burn-in is discarded before any statistic is computed. All panels discretize the exogenous processes on 101-state Rouwenhorst grids and share a fixed seed across policy regimes. The estimation panel's length matches the sample; the others are set per exercise.

Table E.1: Simulation panels

E.2 Additional Results from the Benchmark Calibration

The benchmark calibration is the estimation cell used throughout the main text: the slope of the Phillips curve fixed at $\kappa = 0.40$, with the remaining parameters estimated on the small moment set and reported in Table 6.2. This section collects, for that calibration, additional results comparing the myopic, optimal, and delegated monetary policy regime.

Optimal Policy Regime. Table 6.3 summarizes the welfare gains from optimal policy in the stochastic steady state, but does not describe the transition through which the central bank builds reputation. Having established the stationary gains in the benchmark economy, we next use the same calibration to study the dynamics through which those gains arise. Starting from an inherited reputation, the optimal central bank responds more aggressively than under the myopic policy in order to influence private-sector beliefs, so the transition itself can involve policy allocations that differ substantially from their long-run distribution. Figure 19a illustrates another implication of this behavior by comparing the stationary distributions of reputation under the two policy regimes. Optimal policy shifts the distribution toward lower values of $\mathbb{E}^P[\psi^\pi]$, corresponding to a more hawkish reputation, while also making reputation less dispersed. Thus, optimal reputation management improves not only average reputation but also its stability.

Figure 19b shows how reputation transitions toward the stationary level associated with optimal policy from different initial conditions inherited from the myopic regime. The mean path initializes reputation at the mean of the stationary distribution under myopic policy, while the strong and weak initial-reputation paths use initial conditions corresponding to the bottom and top quintiles of that distribution, respectively. All three paths converge toward the same stationary mean under optimal policy, although both the direction and the speed of adjustment depend on the initial state. Starting from the myopic mean or from weak reputation, the central bank gradually builds reputation, whereas an initially strong reputation starts below the level sustained on average under optimal policy and converges back toward it. The adjustment is persistent but strongly front loaded: the half-life ranges from approximately six to eight quarters, while convergence takes about 23 quarters (6 years) from strong reputation, 34 quarters (8 years) from the myopic mean, and 42 quarters (10 years) from weak reputation. Thus, initial reputation matters for the length of the transition, although a substantial share of the adjustment occurs within the first several years.

Figure 20 illustrates how optimal policy affects the response of reputation to a persistent

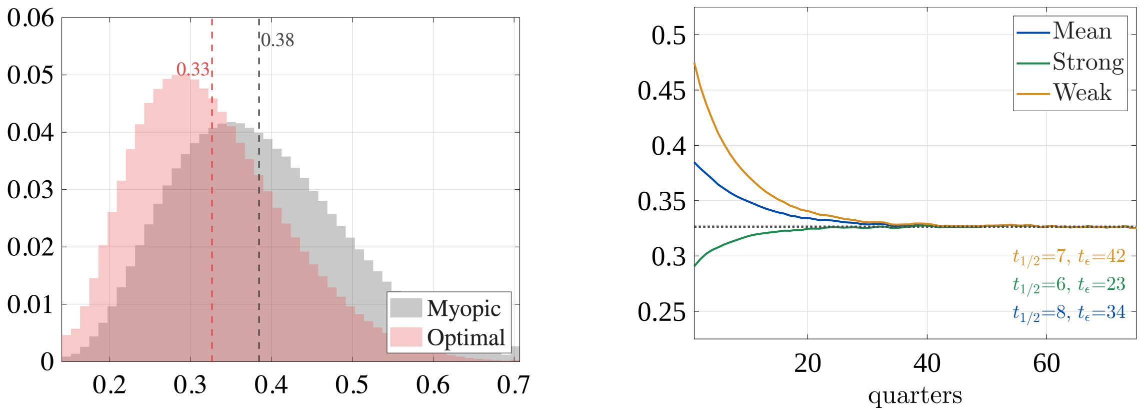

(a) Stationary reputation distribution

(b) Transition to optimal reputation

Figure 19: Reputation under myopic and optimal Policy

Notes: Panel (a) reports the stationary distribution of reputation under the myopic and optimal policy regimes. Panel (b) reports the transition of reputation after switching to the optimal policy from three initial conditions drawn from the stationary distribution under myopic policy: its mean, a strong-reputation state corresponding to the bottom quintile, and a weak-reputation state corresponding to the top quintile. The horizontal dotted line denotes the stationary mean under optimal policy.

cost-push shock. Because the model is nonlinear and the response of beliefs depends importantly on the initial level of reputation, the unconditional impulse response of reputation is estimated relatively imprecisely. We therefore report responses conditional on two initial reputation regimes, defined by the bottom and top quintiles of the stationary distribution of $\mathbb{E}^P[\psi^\pi]$, corresponding respectively to strong and weak reputation. When initial reputation is strong, a cost-push shock deteriorates reputation, but the loss is substantially smaller under optimal policy (See Figure 20a). When initial reputation is weak, the shock instead improves reputation, as the policy response provides a favorable signal about the central bank's hawkishness (See Figure 20b). Although this sign reversal may appear counterintuitive, optimal policy again dampens the response relative to the myopic regime. Thus, independently of whether a cost-push shock raises or lowers perceived hawkishness, optimal policy makes reputation less sensitive to the shock. This greater stability of beliefs is a feature of optimal reputation management.

Delegated Regimes. Figure 21 shows the transition of reputation following delegation. Panel (a) corresponds to the first exercise, in which society has preference weight λ and delegates policy to the more hawkish myopic banker with λ^*. Starting from the stationary distribution of the naive myopic regime, reputation converges toward the more hawkish level

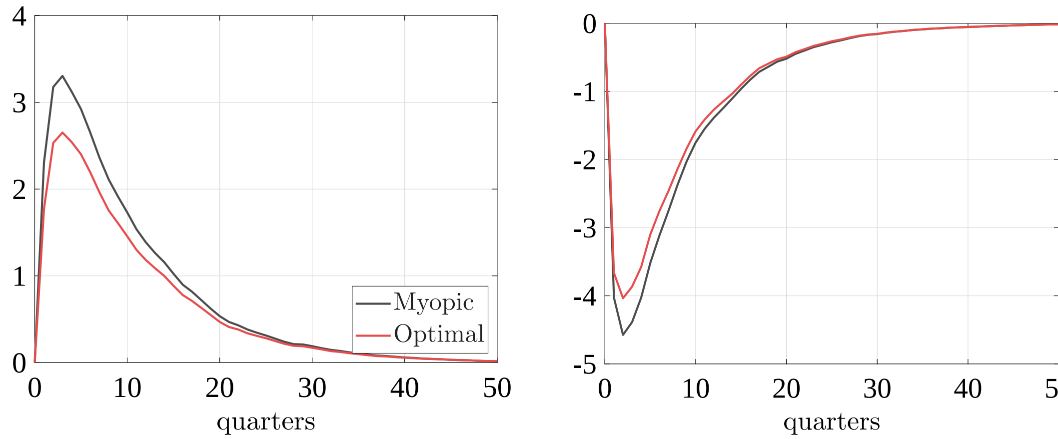

(a) Strong initial reputation **(b)** Weak initial reputation

Figure 20: Reputation responses to cost-push shocks under myopic and optimal policy

Notes: The figure reports impulse responses of reputation, $\mathbb{E}^P[\psi^\pi]$, to a persistent cost-push shock under the myopic and optimal policy regimes. Due to nonlinearity, responses are computed conditional on two regions of the stationary reputation distribution. Strong reputation corresponds to initial values in the bottom quintile, while weak reputation corresponds to the top quintile.

sustained by the delegate from all three initial conditions. As under the optimal policy, the adjustment is strongly front loaded: half lives range from 6 to 8 quarters, while full convergence takes between 23 and 42 quarters depending on inherited reputation. Panel (b) reports the analogous transition for the inverse exercise, in which society has preference weight λ^{**} and delegates policy to the observed myopic banker with preference λ. Delegation again shifts reputation toward a substantially more hawkish long run level. The adjustment is even faster when initial reputation is already close to the delegated steady state, with half lives ranging from 3 to 7 quarters and convergence times from 5 to 42 quarters. Overall, both exercises imply similar dynamics: appointing a banker whose inflation stabilization weight is roughly twice that of society produces a sizable and relatively rapid improvement in reputation, while inherited reputation mainly determines the length of the remaining transition.

E.3 Robustness: Counterfactuals across Estimation Cells

The main text presents the dynamic counterfactuals at the $\kappa = 0.40$ calibration estimated with the small moment set. This section reports the same exercises for all four estimation cells and separates the conclusions that survive the change in parameterization from the magnitudes that do not. Estimated parameter values for each cell are reported in Table 6.2.

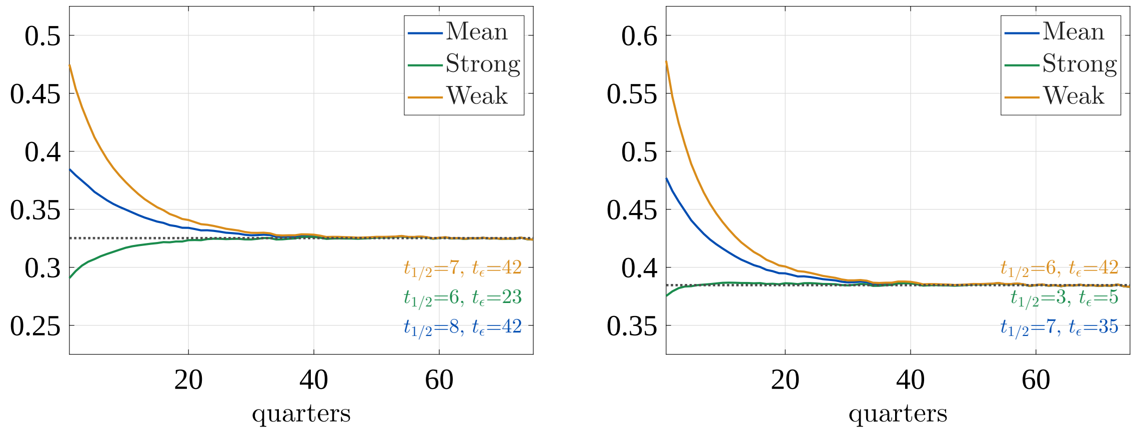

(a) Delegation from λ to λ^* **(b)** Delegation from λ^{**} to λ

Figure 21: Reputation Dynamics under Delegation

Notes: Panel (a) reports reputation dynamics for delegation from λ to λ^*; Panel (b) reports the inverse exercise from λ^{**} to λ. Initial conditions are the mean, strong, and weak reputation states under the corresponding naive myopic regime. The dotted line denotes the delegated steady state.

Estimation cells. The four cells do not fit the data equally well. At the estimated parameters the SMM criterion equals 0.21 and 0.45 for the two $\kappa = 0.20$ cells and 0.08 for $\kappa = 0.40$ with the large moment set, against 1.9×10^{-8} for $\kappa = 0.40$ with the small set, which reproduces every target to three decimals. That exact fit reflects identification rather than validation. The small moment set contains three second moments together with the response of inflation expectations at horizons $h = 0$ and $h = 1$. Because the model implies $\mathbb{E}_t^P[\pi_{t+h}] = \rho_\varepsilon^h \mathbb{E}_t^P[\pi_t]$, and the empirical responses satisfy the same ratio by construction of ρ_ε, the horizon-one response does not provide an independent restriction. The small specification therefore contains four effective moment conditions for four estimated parameters. This does not, however, guarantee an exact fit: conditional on a fixed κ, the empirical moment vector need not belong to the set of moment combinations attainable by the model. At $\kappa = 0.40$, the calibration reproduces all four effective targets essentially exactly, whereas at $\kappa = 0.20$ the minimized criterion remains positive. The large-moment specifications are overidentified.

The cells differ in every estimated parameter, but one difference dominates the counterfactuals. The precision of beliefs is $\tau \approx 200$, 294 and 307 in three cells and 12.2 in the $\kappa = 0.40$ small-moment cell. Because the Kalman gain is decreasing in τ, that cell describes an economy in which the private sector revises its assessment of the central bank an order of magnitude faster than in the others. A second consequence is visible in the level of reputation. Under the myopic policy the pass-through of shocks into expectations is constant at $1/(1+\lambda\kappa^2)$, so average reputation should equal that value, and it does in three cells to within 0.006. In the $\kappa = 0.40$ small-moment cell the simulated average is 0.384 against a full-information value of 0.336. When beliefs are diffuse, the average of the mapping from the belief location to expected pass-through no longer coincides with that mapping evaluated at the average location, and the wedge is a measure of how much belief uncertainty the estimated precision leaves in the economy.

Results. Table E.2 sets the myopic and optimal regimes side by side for each cell. Three features hold in every one of them. First, the optimal central bank ends with better credibility than the myopic one: average $\mathbb{E}^P[\psi^\pi]$ falls from 0.55 to 0.42, from 0.46 to 0.36, from 0.38 to 0.33 and from 0.35 to 0.27. Since $\mathbb{E}^P[\psi^\pi]$ is the expected pass-through of shocks into inflation expectations, a lower value is a more anchored private sector. Second, the optimal policy purchases that anchoring by accepting more output-gap volatility in exchange for less inflation volatility: $\sigma(\pi)$ falls by roughly 6 to 13 percent while $\sigma(y)$ rises by roughly 3 to 5 percent, and $\sigma(i)$ moves by at most 1.2 percent in either direction, so the reallocation occurs between the two objectives rather than through a more aggressive interest rate path. Third, the welfare gain over the myopic benchmark is positive in every cell. The direction of each of these comparisons is therefore a property of the model, not of a parameterization.

Volatilities, reputation, and welfare	$\kappa = 0.20$		$\kappa = 0.40$	
	Small	Large	Small	Large
λ (estimated)	19.76	29.28	12.35	11.43
$\sigma(y)$	2.14 / 2.23	2.23 / 2.29	1.78 / 1.86	1.78 / 1.83
$\sigma(\pi)$	1.87 / 1.62	1.68 / 1.54	2.27 / 2.14	2.27 / 2.12
$\sigma(i)$	7.40 / 7.35	7.27 / 7.25	5.99 / 6.02	5.46 / 5.39
Reputation, $E^P[\psi^\pi]$	0.55 / 0.42	0.46 / 0.36	0.38 / 0.33	0.35 / 0.27
Cost-push pass-through to $E_t^P[\pi_{t+h}]$, $h=0$	0.94 / 0.64	0.72 / 0.53	0.67 / 0.54	0.50 / 0.37
Welfare gain over myopic (%)	7.47	5.14	2.11	4.26

Note: Great Moderation sample: 1992-Q1 to 2019-Q4 (112 obs.). Each column is one estimation cell at its estimated parameters. Each entry is myopic / optimal. Welfare is the discounted sum and gains are in percent of the benchmark.

Table E.2: Optimal vs. myopic central bank, by estimation cell

The magnitudes differ substantially. In the gain row of Table E.2 the welfare gain ranges from 2.1 to 7.5 percent, and average reputation is more hawkish when $\kappa = 0.40$ under both policy regimes. The gain is larger for $\kappa = 0.20$ under both moment sets, and one factor contributing to this pattern is the lower estimated value of λ in the steep-Phillips-curve cells. The welfare benefit of using reputation management to reduce inflation volatility is then smaller relative to the additional output-gap volatility it generates. Read the other way, the value of actively managing reputation is greater when inflation stabilization is intrinsically more costly. When the Phillips curve is flatter, improving reputation delivers a larger benefit because it lowers the future output cost of stabilizing inflation; as monetary policy becomes more effective on impact, this intertemporal smoothing margin becomes less valuable and the gain from departing from the myopic policy falls.[50] It is worth noting that the main text reports the smallest of the four gains, so the quantitative claims there are conservative relative to the rest of the grid.

Table E.3 reports the two delegation exercises. The delegated bank behaves with a weight λ^* above society's, chosen to matching the optimal policy's welfare. That the two deliver the same welfare is therefore true by construction and carries no information. What the exercise does establish is that the delegate reproduces the optimal allocation, which is not imposed: comparing the delegate entries in the upper subrows of Table E.3 with the optimal column of Table E.2, its output-gap and inflation volatilities agree to within half a percent and its average reputation to within 1.1 percent. A myopic bank with a distorted objective therefore replicates a forward-looking bank that internalizes the effect of its choices on beliefs, rather than merely tying with it in the welfare metric. The required degree of conservatism is also stable: λ^*/λ lies between 1.53 and 1.79 in all four cells, and the inverse exercise, which reads the estimated weight as the product of a delegation already in place, returns λ/λ^{**} between 1.63 and 1.94. Both ratios are noticeably higher under the small moment set than under the large one and are almost invariant to κ, so the size of the delegation wedge is disciplined more by which moments are targeted than by the slope of the Phillips curve. One exception is worth stating: in the $\kappa = 0.40$ small-moment cell the welfare frontier peaks 0.045 percentage points short of the optimum, so λ^* there is the welfare-maximizing delegate rather than an exact match, and delegation recovers almost but not all of the gain from optimal policy.

The convergence panel of Table E.3 reports how long a delegated bank takes to bring average reputation to its new long-run level, and this is the dimension on which the cells disagree most. Starting from the source regime's mean, the half-life is 66, 99, 8 and 92

[50] The two readings cannot be separated within this grid. Estimated λ is systematically lower in the $\kappa = 0.40$ cells (12.3 and 11.4, against 19.8 and 29.3), so the slope of the Phillips curve and the weight on inflation move together across the four columns, and the welfare gain reflects the joint effect of the estimated preference and learning parameters rather than λ or κ alone.

	$\kappa = 0.20$		$\kappa = 0.40$	
	Small	Large	Small	Large
Preference weights				
λ (estimated)	19.76	29.28	12.35	11.43
λ^* (delegate for social-weighted λ)	33.76	44.92	22.08	17.59
λ^{**} (social weight for delegated λ)	10.23	17.73	6.35	6.99
λ^*/λ	1.71	1.53	1.79	1.54
λ/λ^{**}	1.93	1.65	1.94	1.64
Volatilities, reputation, and welfare				
$\sigma(y)$	2.14 / 2.23	2.23 / 2.30	1.78 / 1.87	1.78 / 1.83
	2.01 / 2.14	2.13 / 2.23	1.65 / 1.78	1.71 / 1.78
$\sigma(\pi)$	1.87 / 1.62	1.68 / 1.53	2.27 / 2.13	2.27 / 2.12
	2.37 / 1.87	1.97 / 1.68	2.62 / 2.27	2.59 / 2.27
$\sigma(i)$	7.40 / 7.35	7.27 / 7.26	5.99 / 6.03	5.46 / 5.39
	7.49 / 7.40	7.32 / 7.27	5.96 / 5.99	5.58 / 5.46
Reputation, $E^P[\psi^\pi]$	0.55 / 0.42	0.46 / 0.36	0.38 / 0.32	0.35 / 0.27
	0.70 / 0.55	0.58 / 0.46	0.48 / 0.38	0.47 / 0.35
Welfare gain over naive (%)	7.47	5.14	2.06	4.26
	10.93	8.15	4.03	7.32
Convergence of $E^P[\psi^\pi]$ (quarters)				
half-life, from mean	66	99	8	92
	47	78	7	74
half-life, from strong	71	106	6	97
	51	83	3	79
half-life, from weak	62	94	7	88
	43	73	6	71
ϵ-crossing, from mean	479	656	42	591
	358	555	35	531
ϵ-crossing, from strong	474	610	23	534
	355	534	5	508
ϵ-crossing, from weak	510	672	42	627
	395	591	42	551

Note: Great Moderation sample: 1992-Q1 to 2019-Q4 (112 obs.). Each column is one estimation cell at its estimated parameters. Each entry has two subrows: society λ above, society λ^{**} below, and within a subrow reads naive / delegate. Welfare is the discounted sum and gains are in percent of the benchmark. Convergence indicators are in quarters and refer to the delegated bank; definitions in the text.

Table E.3: Delegation: naive vs. delegated central bank, by estimation cell

quarters for the delegate appointed by a society of weight λ, and 47, 78, 7 and 74 quarters for the delegate implied by λ^{**}. The ordering tracks the estimated precision of beliefs: the

$\kappa = 0.40$ small-moment cell, at $\tau = 12.2$, adjusts an order of magnitude faster than the three cells with τ near 200 or above. Within a cell the two delegates differ by at most forty percent, against a spread of more than ten to one across cells, so the speed at which a reputation is rebuilt is governed by how precisely the private sector holds its prior rather than by the weight the bank is given. The policy choice determines the level of reputation the economy converges to; the learning technology determines how long it takes to get there.

Two caveats follow. The starting point matters little: the same panel reports half-lives from a strong and from a weak initial reputation, which shift the numbers by less than ten percent in the three slow cells and leave the ordering across cells unchanged. Full convergence, however, is slow in the high-precision cells. The ϵ-crossing rows, which report the horizon at which the mean path settles within a fixed tolerance of its long-run level, exceed a century in all three of these cells. The half-life is therefore the more informative measure of adjustment speed, and the benchmark cell's eight-quarter half-life should be interpreted as an unusually fast case rather than as a generic prediction of the model.