

Expectations and the Transmission of Monetary Policy[†]

Vicente Jimenez-Gimpel Tomás E. Caravello
MIT MIT

October 9, 2026

[Click here for the latest version.](#)

Abstract

Monetary policy works largely through expectations: core transmission channels, including long-term yields, mortgage rates, and investment, depend on expected future policy. Markets interpret current choices as informative about the central bank’s real-time readings of the economy—readings that persist and guide future policy. When choosing a rate, the central bank takes as given how markets interpret its decisions. However, its systematic behavior shapes that interpretation, and hence how strongly mistakes in those readings affect the economy. We show that this creates a time-inconsistency problem: optimal meeting-by-meeting policy can leave market expectations excessively sensitive to policy surprises, amplifying the effects of those mistakes. We quantify this mechanism’s importance in an extended New Keynesian model estimated using macroeconomic data, Federal Reserve forecasts, natural-output uncertainty estimates, and inflation-swap responses to policy surprises. The optimal *ex ante* rule substantially reduces output-gap volatility by responding less aggressively overall, especially to estimates of natural output.

Keywords: Monetary policy, monetary transmission, expectations management, time inconsistency

JEL Codes: D83, D84, E52, E58, E61

[†]E-mails: jimenezv@mit.edu, tomasec@mit.edu. Jimenez-Gimpel is deeply grateful to Ricardo J. Caballero and Christian K. Wolf for their continued support and guidance. We are particularly grateful to Iván Werning for many helpful discussions and insightful suggestions. We also thank Daron Acemoglu, Cristina Arellano, Corina Boar, Gabriel Chodorow-Reich, Anna Cieslak, Darrell Duffie, Jesús Fernández-Villaverde, Lukas Freund, Rohan Kekre, Dmitry Mukhin, Jonathan Parker, Christina Romer, David Romer, Kunal Sangani, David Thesmar, Jean Tirole, Adrien Verdelhan, Olivier Wang, and participants at several venues for helpful comments.

1 Introduction

Central banks set policy on the basis of their real-time reading of the economy, and markets know it. Consequently, a rate decision does two things at once: it changes current borrowing costs, and it reveals information about what the central bank believes is happening. Some components of that reading are persistent: what the central bank believes today is informative about what it will believe—and do—tomorrow. Thus, a rate decision is informative about future policy and, in turn, about key macroeconomic variables such as output and inflation, even if it reveals nothing about fundamentals. Crucially, nearly every channel through which monetary policy reaches the real economy is forward-looking (e.g., long yields, mortgage rates, asset prices, investment decisions). Thus, what markets infer about future policy from a given decision is not simply a side-effect; it largely *is* monetary transmission. An immediate consequence of this is that two identical rate choices—and even identical paths of policy rates—can have very different effects, depending on what markets infer about the views behind them. Thus, central banks communicate through their decisions as well as their words, whether they intend to or not.

Federal Reserve Chair Kevin Warsh has recently argued that reducing explicit guidance would encourage markets to focus more on economic fundamentals ([Warsh, 2026](#)). Nevertheless, future policy shapes outcomes that markets care about, so they will continue to interpret the Fed’s decisions. With less explicit guidance, markets are likely to rely more heavily on the one signal the Fed can’t avoid sending: its rate decisions.

We develop a framework to study monetary policy when its transmission depends on how markets interpret policy decisions. The central bank acts on its real-time imperfect assessments of demand and the economy’s productive capacity. The private sector does not observe these assessments and must infer them from policy decisions. Since generically a policy decision reveals these assessments only partially, the systematic policy rule determines how markets revise their expectations of future policy and output—and thus how strongly the decision affects current output. Consequently, unlike in the conventional New Keynesian benchmark, the realized path of interest rates does not fully determine the effects of monetary policy.

This dependence creates a novel source of time inconsistency. A central bank setting policy meeting by meeting uses the response of market expectations to stabilize the economy, but takes their sensitivity as given. The map from policy surprises to belief revisions depends on the market’s understanding of the systematic rule, which is not affected by a single decision. The optimal *ex ante* rule provides a benchmark for the resulting inefficiency: it accounts for how systematic behavior shapes market interpretation and thus limits the

propagation of assessment errors into expectations and output. These gains arise even when future policy remains discretionary. In a richer model estimated on U.S. data, we find that this mechanism has quantitatively important stabilization costs. A rule responding only to current assessments captures almost all of the improvement under this benchmark, with less aggressive responses overall, especially to assessments of the economy’s productive capacity. Counterfactual exercises show that changing how markets interpret policy is a key source of these stabilization gains.

Framework. We study a stylized linear-quadratic economy with rigid prices in which monetary policy affects output through both the standard intertemporal-substitution channel and a *belief-revision channel*. Output is determined by the market via the zero-inflation version of a standard IS curve. The central bank’s decisions are based on private and persistent assessments of contemporaneous demand and productive capacity (natural output), the economy’s fundamentals. The market updates its beliefs about these estimates upon observing the policy rate. To isolate our mechanism, we assume the market observes the fundamentals themselves, so what it learns from a policy decision concerns only the central bank’s reading of them. A surprise tightening can mean that the central bank sees stronger demand or lower productive capacity. These assessments have different implications for future policy and output, and the systematic rule determines how markets attribute the surprise between them. Thus, the rule’s relative responses to the two assessments shape how strongly a given decision moves market expectations and, in turn, current output ([Proposition 1](#)).

The time inconsistency. We show that meeting-by-meeting policy leaves output excessively exposed to the central bank’s assessment errors. Given its current reading of the economy, the central bank uses both the direct effect of the policy rate and the response of market expectations to align expected output with its assessment of productive capacity (i.e., certainty equivalence holds). However, errors in that reading generate policy surprises, which induce a revision in the market’s expectations of future policy and output. Consequently, the same expectations channel that helps the central bank stabilize the economy amplifies its mistakes. In equilibrium, markets interpret decisions using the central bank’s systematic policy, yet the central bank takes that interpretation as given when choosing each rate ([Proposition 2](#)).

The optimal ex ante linear rule provides a natural benchmark for assessing this inefficiency. It balances stabilization conditional on the central bank’s assessments against the volatility generated when those assessments are wrong, accounting for how systematic policy shapes market interpretation ([Lemma 2](#)). Under conditions characterized below, the opti-

mal rule places less weight on natural-output assessments relative to demand, decreasing the sensitivity of market beliefs to policy surprises and thus limiting the amplification of assessment errors ([Proposition 3](#) and [Corollary 1](#)). The comparison reveals a novel source of time inconsistency that arises even when future policy remains discretionary. Once its assessments are realized, the central bank has an incentive to depart from the ex ante rule to stabilize the economy as it currently sees it, taking market interpretation as given.

Empirical and quantitative analysis. The empirical analysis suggests that the expectations response at the heart of this mechanism is economically substantial and varies across policy regimes. Using high-frequency changes in one-year OIS rates around monetary-policy events and changes in inflation swaps, we find economically significant negative responses of expected inflation to positive policy surprises. Furthermore, the estimated response per unit of policy surprise is much larger during 2010–2019 than either before the Global Financial Crisis or during 2022–2026. This pattern is consistent with the model’s core insight that systematic policy shapes how markets interpret a given rate surprise.

We then use an estimated model to assess the stabilization costs associated with this mechanism and identify how systematic policy can limit them. For this, we use a richer infinite-horizon New Keynesian model with inflation and cost-push shocks. The central bank receives private noisy signals about these fundamentals in every period, while markets learn about its assessments from the history of policy rates. We estimate the model using U.S. macroeconomic data, Federal Reserve real-time estimates, and the inflation-swaps response to policy surprises described above. At our estimates, the optimal linear commitment policy reduces the central bank’s losses by 69% relative to discretion. Commitment lowers the sensitivity of market expectations to policy surprises, limiting the amplification of the central bank’s assessment errors.

Almost all of these gains are captured by a rule that responds only to current central-bank assessments and does so much less aggressively overall, especially to assessments of natural output. The latter is highly persistent and difficult to assess in real time, so mistakes in these assessments are slow to correct. When markets attribute a policy surprise to such a mistake, they revise their expectations of policy well into the future, amplifying the effects of assessment errors on current output and inflation. A weaker response to natural output relative to demand reduces how much of a policy surprise markets attribute to these persistent mistakes, dampening the resulting fluctuations. A counterfactual exercise illustrates how much this change in interpretation matters: weakening only the response to natural-output assessments reduces the output-gap contraction on impact by 41% following an overestimate of demand, even though the realized interest-rate path is unchanged. Mar-

kets infer less pessimism about natural output from the same tightening and therefore revise their expectations of future policy less sharply.

These findings echo the case for caution in responding to noisy real-time estimates of economic slack emphasized by [Orphanides \(2001, 2003\)](#). Our mechanism provides a distinct rationale for this caution: responding less aggressively to assessments changes how markets interpret policy and limits the amplification of assessment errors, even though the central bank filters the available data optimally.

Overall, our results show that systematic policy shapes how strongly expectations respond to policy decisions, giving rise to a novel source of time inconsistency. Our quantitative analysis suggests that these costs are large, making the market’s interpretation of policy an important consideration for policy design.

Discussion and related literature. This paper is thematically related to several strands of literature. First, it relates to work on the signaling role of monetary policy. Related papers study how policy actions reveal fundamentals or policy objectives, shaping firms’ pricing and inflation expectations ([Baeriswyl and Cornand, 2010](#); [Mertens, 2016](#); [Melosi, 2017](#); [Jia, 2023](#)). Our framework assumes no superior central-bank information about fundamentals: the central bank knows its own real-time views of the economy, which markets must infer from policy decisions.¹ Relatedly, [Caballero and Simsek \(2022\)](#) study monetary policy when the central bank and markets disagree about the economic outlook. They also extend their analysis to allow the market to infer the central bank’s views from its policy decisions. There, however, systematic policy does not affect how markets attribute policy surprises to different components of the central bank’s outlook. In our model, the policy rule shapes that attribution and thus the propagation of central-bank assessment errors. This dependence is crucial for the time inconsistency at the core of our analysis.

Our paper also relates to previous work on optimal monetary policy under imperfect central-bank information (e.g., [Orphanides \(2001\)](#), [Aoki \(2003\)](#), and [Svensson and Woodford \(2003, 2004\)](#)). In these analyses, however, imperfect central-bank information does not create a persistent inference problem for markets.² In our model, by contrast, the central bank’s assessments remain private and persistent, and markets must infer them from policy

¹[Bauer and Swanson \(2023\)](#) report survey evidence that some forecasters view policy surprises as news about future policy, not about fundamentals (since they don’t believe the Fed to have superior information). Relatedly, [Bauer et al. \(2024\)](#) document revisions in perceptions of future policy following monetary policy surprises.

²In [Svensson and Woodford \(2004\)](#), contemporaneously observed endogenous variables reveal the central bank’s information relevant for the equilibrium. In [Aoki \(2003\)](#), the market does not directly observe the central bank’s noisy indicators, but the private component of its assessment is transitory and has no implications for future policy. For more extended discussions of these, see [Appendix G.1](#).

decisions. The systematic policy rule then determines how informative interest rates are and, through this inference, affects the strength of monetary transmission itself.

Relatedly, our quantitative analysis has a precedent in [Aoki \(2006\)](#), which measures the gains from commitment in a forward-looking model with imperfect central-bank information. Since its model does not feature private central bank information, the nature of their time-inconsistency is different.³ Unlike our paper, it finds that commitment gains are small for the U.S. [Lippi and Neri \(2007\)](#) also estimate a model with imperfect information using euro-area macroeconomic data, assuming that the central bank and the private sector share the same information. Information asymmetry is central to our mechanism: markets must infer the central bank’s assessments from its policy decisions. Furthermore, our estimation uses revisions to Federal Reserve real-time forecasts to discipline the precision of those assessments.

Third, the paper contributes to the literature on monetary policy and time inconsistency, originating with [Kydland and Prescott \(1977\)](#) and [Calvo \(1978\)](#). The New Keynesian literature further showed that commitment can improve stabilization even in the absence of an inflationary bias, overcoming what has come to be known as “stabilization bias” (see, e.g., [Clarida et al. \(1999\)](#); [Woodford \(1999\)](#)). Our model isolates a different role for commitment: the current policy rule shapes what markets infer about the central bank’s views of the economy and thus how errors in those views propagate. These gains arise even when subsequent policy remains discretionary, highlighting the distinct nature of time inconsistency in our setting.

This informational source of inefficiency has antecedents in the broader signal-jamming literature, where incentives to affect other agents’ beliefs via payoff-relevant actions give rise to inefficient allocations (for early exponents of this, see [Holmstrom \(1982\)](#); [Stein \(1989\)](#)). Within the realm of monetary policy, [Jia \(2023\)](#) shows that in an environment with a perfectly informed central bank, commitment can improve stabilization by changing how the policy rate shapes firms’ expectations of current inflation and hence the Phillips curve. Along different lines, [Stein and Sunderam \(2018\)](#) study how inefficient gradualism can arise if a central bank with private information about an exogenous long-run target rate has a preference for low bond-market volatility. While sharing common conceptual ground (i.e., the role of policy as a signal), the trade-off at the heart of our time inconsistency is different. Concretely, our setting features an imperfectly informed central bank pursuing standard macroeconomic stabilization. Time inconsistency arises because the policy that best stabilizes the economy given the central bank’s current views of the economy is not, in general, the rule that best limits the propagation of mistaken views through the response of market

³Furthermore, it assumes that imperfect information lasts one period only.

beliefs to policy surprises.

Roadmap. The remainder of the paper is organized as follows. [Section 2](#) introduces the physical environment and the agents’ beliefs. [Section 3](#) shows how market inference from policy decisions shapes monetary transmission and characterizes the discretionary equilibrium and the optimal commitment policy. [Section 4](#) presents the empirical evidence on the response of market expectations to policy surprises. [Section 5](#) develops and estimates the richer New Keynesian model and uses it to quantify the mechanism and the gains from commitment. [Section 6](#) concludes. The appendices provide a microfoundation of the linear model, an information-structure foundation, and a game-theoretic analysis of the main equilibrium concepts. They also contain proofs and additional results. The empirical and quantitative appendices provide further details on these analyses, together with additional exercises and robustness checks.

2 Model

This section lays out the baseline model. First, we describe the physical environment: a very simple rigid-prices economy. Then, we discuss the features of the information structure that underpin the analysis.

2.1 Environment

The model below can be obtained as an approximation of a standard nonlinear economy with households, firms and a monetary authority (see [Appendix A](#) for details).

Consider a linear economy where prices are fully rigid and thus output is demand-determined. In particular, output evolves as

$$y_t = \mathbb{E}_t^M y_{t+1} - \gamma(r_t - \rho_t), \quad (1)$$

where $\gamma > 0$ is the elasticity of intertemporal substitution, r_t the interest rate, ρ_t a demand shifter, and $\mathbb{E}_t^M$ the market’s period- t expectation.

The central bank chooses the policy stance r_t to stabilize output around its natural level y_t^n ,

$$\min_{\{r_t\}} \mathbb{E}^C \sum_{t=0}^{\infty} \beta^t (y_t - y_t^n)^2, \quad (2)$$

where $\beta \in (0, 1)$ is the discount factor and $\mathbb{E}^C$ the central bank’s ex-ante expectation. The

two fundamentals, natural output y_t^n and the demand shifter ρ_t , evolve as

$$y_1^n = \varphi_n y_0^n + \varepsilon_n, \quad y_t^n = 0 \text{ for } t \geq 2, \quad (3)$$

and

$$\rho_1 = \bar{\rho} + \varphi_\rho(\rho_0 - \bar{\rho}) + \varepsilon_\rho, \quad \rho_t = \bar{\rho} \text{ for } t \geq 2, \quad (4)$$

with $\varphi_n \in (0, 1]$, $\varphi_\rho \in [0, 1)$, and $\varphi_n \geq \varphi_\rho$: natural output is at least as persistent as demand. The $t = 1$ innovations $(\varepsilon_n, \varepsilon_\rho)$ are mean-zero, Gaussian, and mutually independent.

2.2 Beliefs

This subsection presents the relevant features of the information structure. We specify the central bank's and market's conditional beliefs directly. A detailed heterogeneous-prior foundation and discussion of what a common prior would imply is provided in [Appendix B](#).

Since it is common knowledge that $(y_t^n, \rho_t) = (0, \bar{\rho})$ at periods $t \geq 2$, we focus on periods where fundamentals are stochastic, $t \in \{0, 1\}$.

Central bank beliefs. The central bank's $t = 0$ information is summarized by a vector of imperfect assessments of the state of the economy, $(\hat{y}_0^n, \hat{\rho}_0)$. The central bank's conditional distribution of the fundamentals is given by

$$\begin{pmatrix} y_0^n \\ \rho_0 \end{pmatrix} \Big| \begin{pmatrix} \hat{y}_0^n \\ \hat{\rho}_0 \end{pmatrix} \lesssim \mathcal{N} \left(\begin{pmatrix} \hat{y}_0^n \\ \hat{\rho}_0 \end{pmatrix}, \begin{pmatrix} \sigma_{\eta,C}^2 & 0 \\ 0 & \sigma_{\varepsilon,C}^2 \end{pmatrix} \right), \quad (5)$$

so the assessments correspond to the central bank's best estimates of the fundamentals. The parameters $\sigma_{\eta,C}^2, \sigma_{\varepsilon,C}^2 > 0$ denote the corresponding (conditional) variances. For later use, we define the central bank's assessment errors as

$$\eta_0 \equiv y_0^n - \hat{y}_0^n, \quad \varepsilon_0 \equiv \rho_0 - \hat{\rho}_0.$$

We also state the central bank's ex-ante beliefs (i.e., before observing its $t = 0$ assessments), since they are relevant for the policy problem in [Section 3.4](#). These ex-ante beliefs are such that

$$\begin{pmatrix} \hat{y}_0^n \\ \hat{\rho}_0 \end{pmatrix} \lesssim \mathcal{N} \left(\begin{pmatrix} 0 \\ \bar{\rho} \end{pmatrix}, \begin{pmatrix} \sigma_{n,C}^2 & 0 \\ 0 & \sigma_{\rho,C}^2 \end{pmatrix} \right), \quad (6)$$

where $\sigma_{n,C}^2, \sigma_{\rho,C}^2 > 0$. Note that the two assessments are independent under the central bank's ex-ante beliefs.

We assume the central bank gets no additional information between its $t = 0$ and $t = 1$ policy.⁴ This implies

$$\hat{y}_1^n \equiv \mathbb{E}_1^C[y_1^n] = \varphi_n \hat{y}_0^n, \quad \hat{\rho}_1 \equiv \mathbb{E}_1^C[\rho_1] = \bar{\rho} + \varphi_\rho(\hat{\rho}_0 - \bar{\rho}), \quad (7)$$

so the parameters φ_n and φ_ρ also denote the persistence of the central bank's assessments.

Market beliefs. The market observes the value of contemporaneous fundamentals (y_t^n, ρ_t) and knows their laws of motion, but it does *not* observe the central bank's estimates. Under the market's $t = 0$ beliefs (i.e., conditional on $t = 0$ fundamentals), the estimate vector $(\hat{y}_0^n, \hat{y}_1^n, \hat{\rho}_0, \hat{\rho}_1)$ is jointly Gaussian such that $(\hat{y}_0^n, \hat{y}_1^n) \perp (\hat{\rho}_0, \hat{\rho}_1)$, consistent with the independent nature of $t = 0$ fundamentals and (7). Concretely, market beliefs regarding the central bank's $t = 0$ assessments are such that⁵

$$\begin{pmatrix} \hat{y}_0^n \\ \hat{\rho}_0 \end{pmatrix} \stackrel{M}{\sim} \mathcal{N} \left(\begin{pmatrix} \lambda_n y_0^n \\ \rho_0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & s \end{pmatrix} \right), \quad (8)$$

where $\lambda_n \in [0, 1)$ and $s > 0$. The fundamentals' $t = 1$ innovations $(\varepsilon_n, \varepsilon_\rho)$ are independent of the central bank's $t = 0$ assessments under the market's beliefs.

Note that the above entails an asymmetry, since

$$\mathbb{E}_0^M[\hat{y}_0^n] = \lambda_n y_0^n, \quad \mathbb{E}_0^M[\hat{\rho}_0] = \rho_0. \quad (9)$$

Thus, the market expects the central bank's demand estimate to be centered on realized demand and its natural-output estimate to sit at an average between realized natural output and its unconditional mean, zero.⁶

The parameter s introduced in (8) represents the market's *relative* uncertainty about the two estimates:

$$s \equiv \frac{\text{Var}_0^M(\hat{\rho}_0)}{\text{Var}_0^M(\hat{y}_0^n)} \in (0, \infty), \quad (10)$$

⁴This assumption has two elements. First, the fact that no additional *exogenous* signals are available to the central bank is a harmless simplification that allows not having to carry additional parameters that do not change the qualitative insights of the paper (this is relaxed in the quantitative model where the central bank gets signals each period). Second, the fact that the central bank does not observe output y_0 is more substantial. [Appendix B.1](#) discusses it in detail.

⁵The normalization $\text{Var}_0^M(\hat{y}_0^n) = 1$ is innocuous, since only the *relative* uncertainty about the two assessments plays a role in the model.

⁶The demand centering is a knife-edge case of the more general belief structure in [Appendix B](#). We adopt it for analytical simplicity: it delivers a unit discretionary coefficient on demand assessments. The appendix also explains why a standard proper common-prior information structure instead produces interior centerings (i.e., the loading of $\mathbb{E}_0^M[\hat{\rho}_0]$ on ρ_0 is strictly between 0 and 1). The quantitative model instead derives interior centerings from its common-prior Gaussian signal structure.

which we call the *interpretation state*. As will be apparent in the analysis, this parameter governs the signal-extraction problem a rate surprise poses: a higher s leads the market to read a given surprise more as news about the central bank's demand assessment and less as news about its natural-output assessment. Throughout, s is fixed and common knowledge.

We assume the market observes contemporaneous fundamentals to isolate learning about the central bank's assessments, which matter in their own right because they determine future policy. This assumption is stronger than needed: if the central bank observes noisy versions of the market's signals and both agents regard the market's remaining estimation errors as independent (and mean zero), the analysis carries over with market estimates replacing fundamentals. If, on the other hand, policy decisions also reveal information about fundamentals to the markets, the systematic rule shapes that inference as well, so our mechanism of interest would still be at play.

3 Equilibrium

This section shows how systematic policy shapes the market interpretation of policy decisions and why this generates a time inconsistency: a central bank that sets policy meeting-by-meeting leaves the economy too exposed to volatility driven by its assessment errors. First, we characterize the equilibria at $t \geq 1$, showing that the central bank's persistent assessments shape y_1 . Then, we turn to linear equilibria at $t = 0$ and present the key mechanism of the model: how policy decisions affect market beliefs and thus monetary transmission. We then characterize the discretionary equilibrium and the optimal ex ante commitment policy, highlighting the key difference: discretion takes market interpretation as given, whereas commitment internalizes its dependence on policy and thus reduces the propagation of assessment errors and, under conditions characterized below, the sensitivity of market beliefs.

3.1 Continuation equilibria

We characterize equilibrium by backward induction. Thus, we first solve the continuation from $t = 1$ onward. The following lemma characterizes the equilibria at $t \geq 1$.

Lemma 1 (Continuation equilibria). *For all $t \geq 2$, the central bank implements its first best, i.e., $y_t = y_t^n = 0$. At $t = 1$, it sets $r_1 = \hat{\rho}_1 - \hat{y}_1^n/\gamma$ to achieve*

$$y_1 = \hat{y}_1^n + \gamma(\rho_1 - \hat{\rho}_1). \quad (11)$$

Because all uncertainty is resolved by period 2, and the central bank has one tool (r_t) for one target (y_t) each period, its problem is straightforward and it achieves its first best for $t \geq 2$.

On the other hand, as standard in linear-quadratic environments, optimal discretionary policy satisfies certainty equivalence conditional on the central bank's estimates. Thus, at $t = 1$ the central bank simply implements its target output *on average*, i.e., it chooses policy r_1 to set $\mathbb{E}_1^C[y_1] = \hat{y}_1^n$.

Crucially, the market anticipates this when forming its expectations at $t = 0$: it knows that at $t = 1$, output will be determined by both the actual demand shifter and the central bank's $t = 1$ assessments via the latter's policy decision. Since the market does not observe those assessments and they are persistent, the market uses the $t = 0$ policy decision to learn about the central bank's contemporaneous assessments and thus about $t = 1$ output.

These revisions in expected y_1 affect current output through the IS equation, making the market's anticipation of future policy part of current monetary transmission. The next subsection develops this idea.

3.2 The role of market expectations: policy as a signal

Timing. Figure 1 lays out the sequence of events at $t = 0$. At the start of the period, the fundamentals (y_0^n, ρ_0) and the central bank's assessments $(\hat{y}_0^n, \hat{\rho}_0)$ are realized. The market observes the fundamentals but not the assessments, while the central bank observes its assessments. The central bank then sets the policy rate r_0 . Observing it, the market updates its beliefs about the central bank's assessments—and hence about next period's output—and y_0 is determined.




Figure 1: Timeline of events at $t = 0$.

Linear equilibria and policy transmission. We study *linear equilibria*, i.e., equilibria such that the central bank's policy is a linear function of its private assessments,⁷

$$r_0 = \bar{\rho} + a(\hat{\rho}_0 - \bar{\rho}) - K\hat{y}_0^n, \quad (12)$$

⁷The choice of the intercept is without loss of generality, since both in the (discretionary) equilibrium and the optimal commitment it equals $\bar{\rho}$. This is discussed in more detail in Remark 4.

for some $(a, K) \in \mathbb{R}^2$. For interpretation purposes, it is useful to reparameterize each rule with $a \neq 0$ writing $K = ab$ and thus

$$r_0 = \bar{\rho} + a(\hat{\rho}_0 - \bar{\rho} - b\hat{y}_0^n), \quad a \in \mathbb{R} \setminus \{0\}, \quad b \in \mathbb{R}. \quad (13)$$

The *scale* parameter a determines the overall aggressiveness of the policy response, while the *composition* parameter b pins down the relative weight of the policy reaction on the natural-output assessment. The case $a = 0$ is handled separately: if $K \neq 0$, the central bank responds only to its natural-output assessment, while if $K = 0$, the interest rate is completely unresponsive to the central bank's assessments.

As shown in (11) and mentioned above, $t = 1$ output is a function of the central bank's $t = 1$ assessments. Since the market anticipates this and assessments are persistent, the market learns about future output from the $t = 0$ policy decision, r_0 . The following proposition characterizes the market's inference.

Proposition 1 (Belief sensitivity and policy transmission). *For every policy rule (a, K) with $b \equiv K/a$ whenever $a \neq 0$, the market's beliefs following a policy decision r_0 are such that⁸*

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \begin{cases} \varphi_n \lambda_n y_0^n - \chi(b, s) \left(\frac{r - \mathbb{E}_0^M[r_0]}{a} \right), & \text{if } a \neq 0, \\ \varphi_n \lambda_n y_0^n - \left(\frac{\varphi_n}{K} \right) (r - \mathbb{E}_0^M[r_0]), & \text{if } a = 0, K \neq 0, \\ \varphi_n \lambda_n y_0^n, & \text{if } a = K = 0. \end{cases} \quad (14)$$

where

$$\chi(b, s) = \frac{\varphi_n b + \gamma \varphi_\rho s}{b^2 + s}. \quad (15)$$

Consequently, given a rate decision $r_0 = r$, output is given by

$$y_0 = \begin{cases} \varphi_n \lambda_n y_0^n - \gamma(r - \rho_0) - \chi(b, s) \left(\frac{r - \mathbb{E}_0^M[r_0]}{a} \right), & \text{if } a \neq 0, \\ \varphi_n \lambda_n y_0^n - \gamma(r - \rho_0) - \left(\frac{\varphi_n}{K} \right) (r - \mathbb{E}_0^M[r_0]), & \text{if } a = 0, K \neq 0, \\ \varphi_n \lambda_n y_0^n - \gamma(r - \rho_0), & \text{if } a = K = 0. \end{cases} \quad (16)$$

First, note that if $a = 0$ and $K \neq 0$, the policy decision fully reveals the central bank's natural output assessments $\hat{y}_0^n$, while providing no information about its demand assessments $\hat{\rho}_0$. In this case, the market's inference about $t = 1$ output is simply the central bank's $t = 1$ natural-output assessment, which is a scaled version of its $t = 0$ assessment. To recover that,

⁸Under the proper information structure in [Appendix B.1](#), the market's beliefs are shaped by signals and thus the expectations should be interpreted as conditional on those as well as the policy rate.

the market simply subtracts the common knowledge component of the rate, $\bar{\rho}$, and scales down the remainder using the conjectured rule. If instead $(a, K) = (0, 0)$, the policy rate is constant and conveys no information. Thus, the market’s forecast then remains equal to its prior.

In contrast, if $a \neq 0$, the market’s inference is generically imperfect on both fronts⁹, and it effectively solves a signal-extraction problem to infer the central bank’s $t = 1$ assessments from the policy surprise $r_0 - \mathbb{E}_0^M[r_0]$. Due to the Gaussian nature of all the random variables involved, the market’s inference is linear in the policy surprise, with slope $-\chi(b, s)/a$. We call $\chi(b, s)$ the *sensitivity of market beliefs*. It represents the market’s revision of expected future output in response to a policy surprise normalized by the scale a of the policy rule. Note that it is independent of the latter, since the market scales its updating to exactly offset the overall size of any given policy reaction. Put differently, the informativeness of a policy surprise is independent of the overall rate aggressiveness. It does depend, however, on the composition b and the interpretation state s . Both elements are crucial, since they shape how much the market learns about both central bank assessments, and thus determine the overall impact of policy on output. This is the central object of the paper.

To anchor intuition, consider the case where demand assessments are fully transitory, $\varphi_\rho = 0$. Then,

$$\chi(b, s) = \frac{\varphi_n b}{b^2 + s}.$$

First, note that χ is decreasing in s in the empirically relevant case where $b > 0$. This is because the more the market’s prior uncertainty is about demand-estimates, the more that it attributes a given rate surprise to news about them, which is fully transitory and thus uninformative about future output. In contrast, χ is hump-shaped in b . The reason behind this non-monotonicity is that increasing b changes both the volatility of the rate surprise and the informativeness of a given surprise. When b is large, the former effect dominates the latter, and thus the overall update is dampened. In turn, when b is small, the additional co-movement between rates and future output is stronger than the added volatility.

Because this belief sensitivity χ depends on the composition b and thus on the systematic policy rule, *the central bank’s own rule shapes the strength of its transmission*—an observation that drives everything that follows. [Figure 2](#) summarizes the channel.

Note that fully transitory demand assessments isolate the intuition but are not necessary for the mechanism. In fact, even if $\varphi_n = \varphi_\rho$, generically the rate remains an imperfect signal of future output (see [footnote 9](#) for the knife-edge exception).

⁹This holds except when $b = \varphi_n/(\gamma\varphi_\rho)$ for $\varphi_\rho > 0$. In that case, the rate is a sufficient statistic for predicting $t = 1$ output and thus the inference is perfect. Since $\varphi_n \geq \varphi_\rho$, however, that b is never attained in a discretionary equilibrium as [Proposition 2](#) shows.




Figure 2: The belief-revision channel.

On the interpretation of the model. Like any model, ours abstracts from many details to isolate our key mechanism of interest. Two of these simplifications warrant further discussion to clarify the interpretation of the results.

1. On the private-sector side, the model summarizes the forward-looking elements of the economy in households' expectations of future output. Through a standard consumption–savings decision, these expectations shape current consumption demand and, since prices are fully rigid, current output ([Appendix A](#)). Neither the belief-revision mechanism nor its endogeneity to systematic policy hinges on this particular representation. In richer and more realistic models, the relevant expectations may be formed by financial-market participants and embedded in longer-term rates, mortgage and credit conditions, and asset valuations. Revisions in these expectations would then shape the spending and investment decisions of households and firms and, through them, macroeconomic outcomes. The same belief-revision mechanism—and the dependence of its strength on systematic policy—would thus carry over to such environments.
2. On the policy side, the model collapses the rich, multidimensional set of decisions involved in conducting monetary policy into a single scalar choice: the current policy rate. The mechanism of interest, however, is not exclusive to such literal policy-rate choices. Statements, press conferences, minutes, and speeches are read by market participants as signals of the central bank's assessment and, consequently, shift the expected policy path, current financial conditions, and ultimately macroeconomic activity. In fact, a direct implication of the core feature of our framework—the fact that policy is a signal—is precisely that rate-setting and communication are not so different after all.

This becomes especially apparent in the scale-invariance of $\chi(b, s)$ above, which is essentially the formal counterpart of the phenomenon described by Stein (2014) in the context of central-bank communication: *the softer [the central bank] talks, the more the market leans in to hear better and, thus, the more the whisper gets amplified*. This, then, is a particular implication of the broader mechanism we study, namely, the endogeneity of the sensitivity of market beliefs—and hence, of policy transmission—to the systematic conduct of monetary policy.

3.3 Discretionary equilibrium

We first focus on *discretionary* equilibria at $t = 0$, namely, such that the central bank sets $t = 0$ policy taking the market’s belief map (Proposition 1) as given. This is the outcome that follows if the central bank cannot commit to a rule before observing its $t = 0$ assessments.¹⁰

Formally, fix the *market’s conjectured* linear rule and the market belief map it induces. After observing its assessments, the central bank solves

$$\arg \min_{r \in \mathbb{R}} \mathbb{E}_0^C \left[(y_0 - y_0^n)^2 \right] \quad (17)$$

subject to (16). Crucially, the function $r \mapsto \mathbb{E}_0^M[y_1 \mid r_0 = r]$ in (16) (or, equivalently, the parameters (a, b)) is taken as given: the central bank does not internalize the fact that its policy reaction changes the coefficients (a, b) and thus policy transmission. Loss from date 1 onward is omitted because, by Lemma 1, it is independent of the current rate.

The equilibrium concept is the natural one, a fixed point in which the central bank’s optimal reaction to its assessments coincides with the market’s conjectured rule, and thus the market’s beliefs are consistent with the central bank’s actual rule. The following definition formalizes this notion.

Definition 1 (Linear discretionary equilibrium). *A policy rule $r_0(\hat{y}_0^n, \hat{\rho}_0)$ of the form (12), together with its associated state-contingent output allocation y_0 , is a linear discretionary equilibrium if: (i) the market beliefs induced by the rule satisfy (14); (ii) for every realization of the central bank’s assessments, the rate prescribed by the rule solves (17), taking the induced belief map as given; and (iii) output satisfies the IS curve (1).*

For the interested reader, Appendix C recasts the economy as a game and defines the discretion and commitment solution concepts in more transparent game-theoretic terms.

Unless stated otherwise, *discretionary equilibrium* refers to an equilibrium in which the central bank’s $t = 0$ rate-choice problem has a unique solution. This rules out a knife-edge

¹⁰The reason behind this time inconsistency will become clear in Section 3.4.

case, namely, one in which the direct and the belief-revision effects of a rate adjustment cancel out exactly and thus output is independent of the policy rate and the central bank is indifferent among all rate choices (even though only one validates this as an equilibrium). We exclude these indifference cases from the analysis since they are economically uninteresting and empirically irrelevant.

Proposition 2 (The unique linear discretionary equilibrium). *There exists a unique linear discretionary equilibrium. It is such that $b_d \in (0, (1 - \lambda_n \varphi_n) \gamma^{-1})$, and satisfies*

$$a_d = 1, \quad b_d(\gamma + (1 - \lambda_n) \chi(b_d, s)) = 1 - \lambda_n \varphi_n. \quad (18)$$

Intuitively, due to the linear-quadratic nature of the central bank's problem, certainty equivalence holds: taking as given the sensitivity of expectations to rate movements, the central bank adjusts its policy reaction to minimize average inefficient output fluctuations, considering both the direct effect of the current rate (the change in the price of consumption today relative to consumption tomorrow) and the belief-revision effect on expectations of next-period output. This (conditional) expected output gap is given by

$$\mathbb{E}_0^C[y_0 - y_0^n] = [b(\gamma a + (1 - \lambda_n) \chi(b, s)) - (1 - \lambda_n \varphi_n)] \hat{y}_0^n + \gamma(1 - a)(\hat{\rho}_0 - \bar{\rho}). \quad (19)$$

Since the belief revision effect on the demand side is mean zero, the central bank disregards that component of transmission and $a_d = 1$ follows. On the other hand, the central bank's reactions to its output target estimate do shift expectations on average, and thus it picks its policy to achieve its desired average output movements considering both components of transmission. This explains the upper bound $(1 - \lambda_n \varphi_n) \gamma^{-1}$, since that would be the desired reaction if only the standard transmission were present. Since overall transmission is stronger with the belief-revision component, a less aggressive reaction is optimal (note that the target total transmission is not one-to-one since the market prior beliefs “do part of the work”, by anticipating a future rate response of $-\lambda_n \varphi_n / \gamma$ times the natural output realization on average).

Importantly, as (19) shows, the belief-revision impact is not simply χ , since that element of transmission is discounted by what the markets already expect the policy rate to be. The natural-output component of that expected rate is $-ab\lambda_n y_0^n$, which contributes $-b\chi\lambda_n y_0^n$ to the revision in expected output. In addition, the target pass-through of natural output assessments is less than one, since a fraction $\lambda_n \varphi_n$ of it is already incorporated by the market via its ex-ante beliefs about future policy and its knowledge of the current natural output realization.

Finally, equilibrium uniqueness follows from the fact that increasing the central bank's reaction coefficient b strictly increases the output response to a given natural-output assessment: the larger direct response more than offsets any weakening of the belief revision channel. Consequently, the LHS of the second equality in (18) is strictly increasing for $b > 0$, so exactly one composition coefficient delivers the desired stabilization. Note that the strength of transmission per unit of policy-rate change *is not* in general increasing in b , since χ is hump-shaped as discussed above.

3.4 Optimal commitment policy

Now consider the problem where the central bank commits to a $t = 0$ linear policy rule *before observing its $t = 0$ assessments*. Crucially, this entails internalizing that the policy rule shapes belief sensitivity χ .

The central bank's optimal commitment policy solves

$$\arg \min_{(a,K) \in \mathbb{R}^2} \mathbb{E}^C \left[(y_0 - y_0^n)^2 \right] \quad (20)$$

subject to (12) and (16), with the market belief map induced by the chosen coefficients (a, K) . See [Appendix C](#) for a game-theoretic interpretation of the commitment problem and a discussion of the differences with the discretionary equilibrium above.

The source of time inconsistency is easiest to see in terms of what the central bank would like to promise. Before observing its assessments, the central bank knows that some of them will be wrong. Consequently, it would like to commit to a systematic mapping from assessments into rates that limits how strongly those errors are propagated through market beliefs and into output. Under the conditions developed below, the optimal rule places a lower relative weight on the natural-output assessment.

Once the assessments have been realized, however, the market's response to any rate decision is already pinned down by the systematic rule it expects. A one-time deviation cannot change that belief map. Taking this map as given, the central bank chooses the rate to set its conditional expected output gap to zero, accounting for both the direct and belief-revision channels of monetary transmission. Consequently, generically it does not follow the ex-ante rule that was chosen to shape the market interpretation.

Remark 1 (The scope of commitment). *The commitment problem above and resulting policy do not correspond to their optimal Ramsey counterparts since (i) it applies only to $t = 0$ policy, not the full path (i.e., it entails no forward guidance, policy reverts to its discretionary equilibrium at $t = 1$); and (ii) it is linear in the central bank's assessments. The former is a*

simplification that allows us to focus on the main mechanism of interest and highlights the novel nature of the time inconsistency problem that is the object of this paper (in contrast to the standard time inconsistency problem studied, for example, in [Woodford \(2003\)](#)). The latter, on the other hand, is a meaningful restriction, unlike in equivalent economies without private information, where the optimal unconstrained Ramsey policy is linear. This is because of the nonlinear nature of the market's belief updating as a function of the policy rule. However, in the spirit of the literature on simple rules, this is a natural family of policies to study. Furthermore, this allows the policy to be compared with the ones in related environments without private information. The quantitative analysis in [Section 5](#) relaxes the former restriction while keeping the latter.

Because the rule is chosen before the central bank observes its assessments, the commitment objective integrates over its ex-ante beliefs specified in [Section 2.2](#). The following lemma provides a convenient representation of the objective function of the commitment problem.¹¹

Lemma 2 (Loss function). *For every rule with $a \neq 0$, the central bank's ex ante loss is given by*

$$L(a, b; s) = (b(\gamma a + \chi(b, s)) - (1 - \lambda_n \varphi_n) - \lambda_n b \chi(b, s))^2 \sigma_{n,C}^2 + (a - 1)^2 \gamma^2 \sigma_{\rho,C}^2 \quad (21) \\ + \underbrace{(\gamma + \chi(b, s))^2 \sigma_{\varepsilon,C}^2 + (1 - \lambda_n \varphi_n + \lambda_n b \chi(b, s))^2 \sigma_{\eta,C}^2}_{\equiv \mathcal{P}(b,s)}.$$

The loss has a natural conditional mean-variance interpretation. Its first two terms capture systematic stabilization errors—that is, variation in the output gap's conditional mean across realizations of the central bank's assessments. The first measures the failure of output to track the natural-output assessment one for one. The second measures the failure to shut down the effect of demand shocks on output.

The third and fourth terms, collectively denoted $\mathcal{P}(b, s)$, capture output-gap volatility caused by assessment errors. Demand-assessment errors affect current output directly, with coefficient γ , and also generate a policy surprise that changes expected future output, with coefficient $\chi(b, s)$. The resulting total pass-through is then $\gamma + \chi(b, s)$. On the other hand, natural output assessment errors entail a one-to-one impact on the output gap that is then attenuated by the ex-ante market expectations of future policy, based on the actual current natural output. Furthermore, the market's expected rate is $\mathbb{E}_0^M[r_0] = \bar{\rho} + a(\rho_0 - \bar{\rho}) - ab\lambda_n y_0^n$.

¹¹[Lemma 2](#) uses the scale-composition parameterization $b = K/a$ and therefore applies to $a \neq 0$. Policy rules with $a = 0, K \neq 0$ and the constant rule $(a, K) = (0, 0)$ are evaluated separately in raw (a, K) coordinates in [Lemma 3](#).

Seeing a high y_0^n that the central bank doesn't, it expects the latter to cut rates by $ab\lambda_n\eta_0$ more than what it actually does. So the realized surprise rises by $ab\lambda_n\eta_0$: the central bank looks hawkish. The market reads that as news that the assessment is low, and revises expected future output down by $\chi b\lambda_n\eta_0$, which entails an output gap movement. Thus, the total pass-through of natural-output assessment errors is $1 - \lambda_n\varphi_n + \lambda_nb\chi(b, s)$.

When demand assessments are transitory, the belief revision component of the pass-through arises entirely because the market partly attributes the policy surprise to persistent natural-output news, regardless of its actual source. With persistent demand assessments, part of it also comes from the share of the surprise that the market attributes to demand-assessment news.

The following proposition presents the paper's main normative result: commitment reduces the pass-through of assessment errors into output and, under some conditions, generically lowers the sensitivity of market beliefs.¹²

Proposition 3 (Optimal Commitment Policy). *Let (a_c, K_c) be a solution of the commitment problem (20) and suppose $a_c \neq 0$.¹³ Let $b_c \equiv K_c/a_c$ and set*

$$m \equiv 1 - \frac{\varphi_n(1 + \lambda_n)}{2}.$$

Then,

- $\mathcal{P}(b_c, s) \leq \mathcal{P}(b_d, s)$;
- if either $\lambda_n = 0$ or

$$\gamma^2 s (\varphi_n - 2\varphi_\rho m) \geq \varphi_n m^2, \quad (22)$$

then

$$\chi(b_c, s) \leq \chi(b_d, s) \quad (23)$$

and the inequality is strict whenever $\lambda_n > 0$. Further, whenever $\lambda_n = 0$, the inequality is also strict except possibly when (22) holds with equality and $\varphi_n > (2 - \varphi_n)\varphi_\rho$.

¹²Note that the proposition does not assert the existence of a commitment optimum. Indeed, this does not hold in general due to the discontinuity at the constant policy rule. Intuitively, at every rule with $a \neq 0$, the market learns about the central bank's assessments from the policy decision (and this learning is in fact independent of the value of a as illustrated in Proposition 1). However, at $(a, K) = (0, 0)$, the market learns nothing. This discontinuity can create a situation in which the central bank can always improve its loss by moving closer to $a = 0$, but never reach it. Lemma 3 provides a sufficient condition for existence of a commitment optimum.

¹³An optimum with $a_c = 0$ is possible and has a different economic interpretation: if $K_c \neq 0$, the central bank responds only to its natural-output assessment, while if $K_c = 0$, the rule is constant and conveys no information. These cases are not empirically relevant and thus are not considered in the analysis from here onward.

Commitment disciplines the market's belief revisions following policy surprises. When $\lambda_n = 0$, it reduces the sensitivity of beliefs and consequently weakens the amplification of demand-assessment errors. The proposition extends this result to $\lambda_n > 0$ under an explicit sufficient condition that aligns the incentives to limit the propagation of demand and natural-output assessment errors.

At the discretionary rule, the first two terms of L are exactly zero: conditional on its assessments, the central bank sets its expected output gap to zero as detailed above. A small departure from that rule therefore creates expected stabilization losses only to second order. By contrast, changing the rule can generate a first-order reduction in output volatility caused by assessment errors $\mathcal{P}$. Commitment internalizes this margin, whereas a discretionary bank takes assessment-error volatility as independent of its current rate choice.

Intuitively, policy affects the propagation of assessment-errors by shaping both the sensitivity of market beliefs and the policy surprises generated by these errors. Demand-assessment errors propagate to output with coefficient $\gamma + \chi(b, s)$, so reducing their amplification calls for lower belief sensitivity. When $\lambda_n > 0$, natural-output assessment errors propagate to the output gap with coefficient $-[1 - \lambda_n \varphi_n + \lambda_n b \chi(b, s)]$. Concretely, this propagation depends not only on the sensitivity of market beliefs χ but also directly on the rule's composition b : a lower b makes expected policy less dependent on realized natural output, reducing the surprise generated when the central bank's assessment differs from that realization. When these two channels conflict, commitment may reduce overall assessment-error propagation while increasing belief sensitivity.

Condition (22) aligns these two mechanisms. To understand it, recall that χ is (minus) the covariance-to-variance ratio relating future output to the scale-free policy surprise. Condition (22) first requires $\varphi_n > 2\varphi_\rho m$. This makes the natural-output component strong enough relative to the persistent demand component that increasing b can raise belief sensitivity. The accompanying lower bound on $\gamma^2 s$ then ensures that the increase in covariance dominates the increase in signal variance, making χ locally increasing in b . The condition therefore places the discretionary composition on the increasing portion of $\chi(\cdot, s)$, so lowering b reduces both χ and $b\chi$. Naturally, this condition is sufficient, not necessary: the motive to lower the sensitivity of beliefs can dominate even when the two channels are not aligned. Indeed, our numerical exercise shows that commitment lowers belief sensitivity well beyond the parameterizations where the sufficient condition holds.¹⁴

¹⁴A sweep over 32,400 parameter configurations shows that $\chi(b_c, s) > \chi(b_d, s)$ in only 3.3% of cases, even though the condition above holds in only 57%. Further, for the most part whenever the ranking reverses it does so by very little: $\chi(b_c, s)/\chi(b_d, s)$ exceeds 1.01 in 1.47% of all configurations and 1.05 in only 0.66%; plus, the median ratio among the cases where $\chi(b_c, s)/\chi(b_d, s) > 1$ is 1.0068. The parameters considered are $\gamma \in \{0.3, 0.5, 1\}$, $\varphi_n \in \{0.9, 0.98, 1\}$, $\varphi_\rho \in \{0, 0.3, 0.6, 0.9\}$, $s \in \{0.02, 0.1, 0.3, 1, 3\}$, $\lambda_n \in \{0.1, 0.3, 0.6, 0.9, 0.99\}$,

How does the central bank reduce the sensitivity of market beliefs? The following corollary shows that, when assessment errors are small relative to variation in the assessments, optimal commitment implements this reduction with a lower composition.¹⁵

Corollary 1 (Local commitment composition). *Suppose (22) holds strictly. For all sufficiently small $\sigma_{\varepsilon,C}^2/\sigma_{n,C}^2$ and $\sigma_{\eta,C}^2/\sigma_{n,C}^2$, any optimal commitment rule has composition*

$$b_c < b_d.$$

The lower composition weakens belief sensitivity and reduces the pass-through of demand-assessment errors. When $\lambda_n > 0$, it also reduces the pass-through of natural-output-assessment errors through the factor $b\chi(b, s)$. Note that the result concerns the *relative weight* on natural output, not the raw natural-output coefficient ab .

When $\lambda_n > 0$, condition (22) is sufficient but not necessary for a lower composition in the small-error regime. The natural-output-error channel depends on $b\chi(b, s)$, which is locally increasing at the discretionary composition and therefore favors a reduction in b . This force may dominate even when the demand-error channel, operating via χ , favors an increase in b . When $\lambda_n = 0$, on the other hand, only the demand-error component depends on policy and hence, away from the knife-edge case where $\chi_b(b_d, s) = 0$, commitment tweaks policy in the direction that reduces χ .

Thus, the reduction in assessment-error propagation need not always be implemented through a lower composition. **Corollary 1** provides conditions under which it is. As we show in **Section 5**, this scenario is the empirically relevant one.

4 Empirical Evidence

This section presents empirical evidence that supports the key mechanisms of the model: expectations about inflation meaningfully react to monetary policy surprises and the reaction of expectations depends on the underlying policy regime. First, we describe the empirical specification. Then, we discuss the data and implementation details. Finally, we present the results and discuss potential limitations.

¹⁵ $\sigma_{\rho,C}^2 \in \{0.25, 1, 4\}$, $\sigma_{\varepsilon,C}^2 \in \{0.05, 0.2, 1\}$ and $\sigma_{\eta,C}^2 \in \{0.05, 0.2, 1, 5\}$.

Proposition 4 in **Appendix E** also shows that $a_c > 1$. This “compensating” scale increase is a consequence of the fact that commitment shapes only $t = 0$ policy: assessment-error pass-through $\mathcal{P}(b, s)$ is independent of a . In **Section 5**, commitment also changes future policy, and scale affects how inferred assessments propagate into future allocations. A simple example illustrates this: if $a \rightarrow 0$, the market does not update much about the future, not because the information content of the current rate is low (it is independent of scale as discussed above), but because in this scenario future outcomes are independent of the central bank’s assessments due to policy being unresponsive.

4.1 Empirical Specification

Let mps_t denote a revision in the expected one-year interest-rate path, and let $\Delta\mathbb{E}_t^M[x_t] \equiv \mathbb{E}_t^M[x_t] - \mathbb{E}_{t-1}^M[x_t]$ denote the corresponding revision in the market expectation of an aggregate variable of interest. Throughout, we focus on inflation due to data availability.¹⁶ Concretely, we let $x_t \in \{\pi_{t,0y \rightarrow 2y}, \pi_{t,2y \rightarrow 5y}, \pi_{t,5y \rightarrow 10y}\}$, where $\pi_{t,h_1 \rightarrow h_2}$ is the breakeven inflation at t , for annualized inflation between times $t + h_1$ and $t + h_2$. Thus, we consider changes in forward breakeven inflation rates at three maturities: 0 to 2 years, 2 to 5 years, and 5 to 10 years.

For each maturity $H \in \{0y \rightarrow 2y, 2y \rightarrow 5y, 5y \rightarrow 10y\}$, we are interested in the coefficient β_H of the following regression

$$\Delta\pi_{t,H} = \alpha_H + \beta_H mps_t + u_{t,H}, \quad (24)$$

where t denotes the date of a monetary policy event (discussed below).

Note that the empirical coefficient above is naturally related to, but different from, the market sensitivity of beliefs χ in the baseline model. The latter measures the response of expected future output to a policy surprise *normalized by the overall scale of the policy reaction*. In practice, we don't observe the central bank's fundamentals' assessments nor the coefficients of its systematic rule. Thus, we estimate the reduced-form response of inflation expectations to a one-year policy surprise. These coefficients β_H differ from χ in outcome of interest and scale, and thus the former should not be interpreted as literal estimates of the latter. They are, however, an observable counterpart of the broader belief-revision margin: how much market expectations change per unit of policy news.

4.2 Measurement and Implementation Details

This subsection discusses the source of variables and events in the specification (24) together with the sample selection.

Monetary policy surprises. We obtain both the event calendar and the policy-news measure from the U.S. Monetary Policy Event-Study Database (USMPD) introduced by Acosta et al. (2025).¹⁷ We use two types of USMPD events: *Monetary Events* and minutes releases. *Monetary Events* combine the policy statement and the subsequent press conference

¹⁶While the baseline model does not have inflation, it is conceptually straightforward to incorporate it, as shown in the quantitative model in Section 5. Furthermore, in that model, as in many others of the New Keynesian variety, market expectations of inflation and output are quite related.

¹⁷Available at <https://www.frbsf.org/research-and-insights/data-and-indicators/us-monetary-policy-event-study-database/>

(if there is one): the event window runs from 10 minutes before the statement release to 60 minutes after the press conference begins. When there is no press conference, USMPD uses the 30-minute statement window running from 10 minutes before to 20 minutes after the release. For minutes releases, USMPD uses a 30-minute window from 10 minutes before to 20 minutes after the release. For each event, we take mps_t directly from the database’s OISIY series. It is given by the intraday change in the fixed rate on a one-year U.S. dollar overnight-indexed swap over the relevant USMPD event window. Thus, mps_t captures unexpected news about the policy path conveyed by the entire FOMC event—through the policy statement and associated press conference, or through the release of minutes, *not simply* an unexpected change in the current federal funds rate. Under the expectations hypothesis, this change approximately measures the market’s revision in the expected overnight policy-rate path over the following year, up to movements in the risk premium. The coefficient β_H in (24), then, measures how expected inflation is revised per unit of policy-path news at an FOMC event. In particular, it is not intended to identify the causal effect of an exogenous Taylor-rule intercept shock.¹⁸ This is the appropriate object to study in our context, since we are interested in how the market updates its beliefs about the Fed’s assessment of the economy and thus about future macroeconomic aggregates and not in the causal effect of a one-off Fed funds shock.

Inflation swaps. We use inflation swaps to measure breakeven inflation. We obtain daily data from Bloomberg for three maturities: 2, 5 and 10 years. Letting $\pi_{t,h}$ be the swap rate at time t for maturity h , we construct the forwards as $\pi_{t,h_1 \rightarrow h_2} = (h_2\pi_{t,h_2} - h_1\pi_{t,h_1})/(h_2 - h_1)$. In order to compute the change in the expectations at the monetary event, we take the difference between the values for the day before and the day after the event.

Sample. Our sample is limited by historical inflation swap data, which starts on July 21, 2004. Thus, the first event is an FOMC meeting on August 10, 2004. Our last monetary event is an FOMC meeting on April 29, 2026. We exclude observations associated with unscheduled FOMC meetings. We also exclude events in calendar years 2020–2021 due to the ongoing Covid-19 crisis. Its very particular nature makes the market’s reaction to monetary policy surprises likely not representative. The final sample contains 318 scheduled FOMC events: 63 in the pre-GFC period (August 10, 2004 to May 21, 2008), 25 in the GFC period (June 25, 2008 to December 16, 2009), 160 in the post-GFC/pre-Covid period (January 6, 2010 to December 11, 2019), and 70 in the post-2021 period (January 5, 2022

¹⁸As discussed in [Bauer and Swanson \(2023\)](#), to quantify the dynamic causal effects of an intercept shock to a Taylor rule, we would need to take additional steps. See also [Jarociński and Karadi \(2026\)](#).




Figure 3: Estimates of (24) using the USMPD OIS1Y surprise. Markers and adjacent numbers report coefficients, whiskers show 95% HC1 confidence intervals. Horizontal axes show forward horizons in years. All panels report percentage points per percentage-point surprise. The GFC panel (June 2008–December 2009) uses y-axis limits $[-10, 5]$, versus $[-2, 1]$ elsewhere.

to April 29, 2026). It comprises 159 policy announcements (including the associated press conference when applicable) and 159 releases of FOMC minutes.

4.3 Results

Figure 3 shows the estimated coefficient in (24) for different maturities and subsamples. The figure shows two main patterns. First, point estimates are negative in every cell, but their precision differs substantially across regimes. The response is statistically different from zero at the 5% level at all three horizons during 2010–2019, at the first two horizons during 2022–2026, and at the 5–10 year horizon during the GFC. None of the pre-GFC responses is significant at that level. To get a sense of the scale, note that pooling the pre-GFC and post-2021 samples ($N = 133$) at the 0–5 year horizon, a 25-basis-point surprise implies a cumulative expected-price response of $0.25 \times (-0.2365 \times 2 - 0.3425 \times 3) \approx -0.38$ percentage points. In the 2010–2019 sample, the corresponding response is $0.25 \times (-0.8715 \times 2 - 0.8916 \times 3) \approx -1.10$, approximately three times as large.

Second, there is meaningful heterogeneity across subsamples. The pre-GFC estimates are small and not significant at the 5% level, while the post-2021 response is concentrated

in the first five years. Relative to the pre-GFC and post-2021 samples, the estimates for the 2010–2019 sample show a much larger reaction, which is relatively flat across horizons. This is consistent with the mechanism in our model. After the GFC, monetary policy changed significantly (e.g., federal funds rates became less volatile due to the binding zero lower bound and communication became an explicit tool). The cross-regime variation is in line with [Proposition 1](#): changes in systematic policy affect what markets infer from a given policy surprise. This view is also consistent with the complementary results in [Bauer et al. \(2024\)](#), who document time variation in the Fed’s perceived responsiveness to economic conditions and an increase in perceived policy inertia.

The financial-crisis period shows a distinct pattern, with particularly large point estimates at longer horizons. This subsample contains only 25 events, and the estimates are less precise than in the other periods, although the 5–10-year response is statistically different from zero. We thus interpret their magnitudes cautiously. Together with the differences across the other periods, they provide further evidence that the response of inflation compensation to policy news varies across economic and policy environments. However, it must be noted that given the exceptional financial stress during the GFC, it is particularly plausible that changes in inflation risk premia partly account for the large responses at longer horizons.

The heterogeneity is also statistically meaningful. In a pooled specification allowing for regime-specific intercepts and slopes, a Wald test rejects equality of the four regime slopes for the cumulative 0–5 year response ($p = .013$). The corresponding horizon-specific tests have p -values of .143, .101, and .002 for the 0–2, 2–5, and 5–10 year horizons, respectively.

Naturally, our explanation for the heterogeneity of the response of beliefs across subsamples is not the only one. In particular, beyond a change in the Fed’s systematic policy (and thus in the informational content of surprises), the heterogeneity could also be due to changes in the underlying macroeconomic environment, such as the persistence of shocks and the sensitivity of the risk-premium to monetary policy. Empirically disentangling these explanations is beyond the scope of the paper. Consequently, we interpret the regime pattern as suggestive evidence that the expectations response is not a policy-invariant structural constant.

4.4 Measurement Robustness and Interpretation

Alternative measure of policy news. Our baseline specification measures policy news using the change in the one-year OIS rate over the USMPD event window. [Appendix F.1](#) instead uses the fourth quarterly Eurodollar/SOFR futures contract, ED4. Although the coefficient levels differ across the two measures, the key finding is unchanged: at every

horizon, the point estimate is more negative during 2010–2019 than in the pre-GFC and post-2021 periods. Hence, our findings of interest are not specific to OIS1Y as the measure of policy-path news.

Alternative measure of inflation expectations. While our baseline uses inflation swaps to measure changes in the market’s inflation expectation, a perhaps natural alternative is using Treasury Inflation-Protected Securities (TIPS) breakevens. However, these can be affected by time-varying liquidity and other market-specific premia (Fleckenstein et al., 2014; Abrahams et al., 2016). Appendix F.2 reports the corresponding regressions using breakevens constructed from nominal Treasury and TIPS yields. The pre-GFC response is significant at the 5% level only at the 5–10 year horizon. The key finding persists: at all three horizons, the point estimate is more negative during 2010–2019 than in the pre-GFC and post-2021 periods.

Quote construction and liquidity. The baseline inflation-swap series uses Bloomberg’s LAST field. In a thin market with wide bid–ask spreads, this could partly reflect stale quotes or changes in the recorded price within the spread range instead of changes in the underlying valuation of inflation compensation. Appendix F.3 rebuilds the swap-forward revisions and runs the specification (24) twice, once using bid-side levels at every maturity and once using ask-side levels. If the baseline result was driven mostly by movements of LAST with an otherwise unchanged bid–ask spread, the estimated coefficients would change materially when the same regressions are estimated at the two sides of the market. However, reassuringly, both the bid- and ask-based estimates preserve the main results.

Inflation-risk premia. The main limitation of the exercise is that inflation swaps measure inflation compensation under the risk-neutral measure. Their response may therefore reflect changes in inflation-risk premia as well as revisions in “physical” inflation expectations. There is no straightforward way to separate these components at the frequency of monetary-policy events: TIPS breakevens contain a corresponding risk-premium component, while survey expectations are not observed in narrow event windows or at comparable horizons (besides having other limitations). Moreover, existing evidence indicates that inflation-risk premia are generally small at shorter horizons (Haubrich et al., 2012; Ajello et al., 2020). Consequently, we interpret the estimates as informative about revisions in market inflation expectations, but acknowledge the possibility of a potentially important contribution from risk premia, particularly during the GFC and at longer horizons.

Event-window timing. The policy surprise is measured over a narrow intraday USMPD window, while the inflation-swap revision runs from the last trading day before the event to the first trading day afterward. The latter may therefore contain other news arriving outside the policy window. This adds noise relative to a fully intraday measure of inflation expectations, but as far as we know there is no evident reason why such news should systematically generate the “ordering” documented above regarding strength of the response across policy subsamples.

Taken together, the evidence indicates that market inflation expectations respond meaningfully to policy-path news and that this response per unit of news varies substantially across policy regimes. The central finding survives an alternative measure of policy news, TIPS-based breakevens, and separate bid and ask swap quotes. The principal unresolved measurement issue is the possible contribution of inflation-risk premia, which cannot be cleanly separated at event frequency. Subject to that unavoidable qualification, the results provide suggestive evidence that the expectations response to monetary policy is economically substantial and depends on the policy environment.

The quantitative model below uses this evidence together with macroeconomic data and real-time Federal Reserve forecast revisions to discipline the strength of the belief-updating mechanism and assess its relevance.

5 Quantitative Analysis

This section quantifies the importance of our mechanism in a New Keynesian model estimated using U.S. data. The analysis delivers two lessons. First, how markets interpret policy decisions substantially affects the macroeconomic consequences of errors in the central bank’s real-time assessments. Second, at our estimates, these costs can be substantially reduced through more restrained systematic responses to the central bank’s assessments, and particularly a weaker response to natural output relative to demand. We first introduce the model and estimate it using U.S. macroeconomic data, Federal Reserve real-time forecasts, and inflation-swap responses to policy surprises. We then examine the gains from commitment and explain how changes in systematic policy achieve them. Finally, we assess which features of the data drive the results.

5.1 Environment

The private sector's behavior is described by a forward-looking IS curve and New Keynesian Phillips curve,

$$y_t = \mathbb{E}_t^M y_{t+1} - \gamma (i_t - \mathbb{E}_t^M \pi_{t+1} - \rho_t), \quad (25)$$

$$\pi_t = \beta \mathbb{E}_t^M \pi_{t+1} + \kappa(y_t - y_t^n) + u_t, \quad (26)$$

where y_t is output, y_t^n is natural output¹⁹, i_t is the nominal short rate, ρ_t is the demand or discount-rate shock, and u_t is a cost-push shock. As in the baseline model, γ is the elasticity of intertemporal substitution in the IS curve, and κ is the NKPC slope. We let $\bar{\rho}$ denote the unconditional mean of the natural rate and hence the steady-state nominal rate under zero trend inflation.

The central bank minimizes a standard dual-mandate loss function with an additional rate-level penalty,

$$\frac{1}{2} \mathbb{E}_0^C \sum_{t=0}^{\infty} \beta^t [x_t^2 + \vartheta_\pi \pi_t^2 + \vartheta_i (i_t - \bar{\rho})^2], \quad (27)$$

where $x_t \equiv y_t - y_t^n$ is the output gap. The fundamental state is $X_t = (y_t^n, \rho_t - \bar{\rho}, u_t)^T$. Each component follows an independent AR(1),

$$X_t = A_X X_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_\varepsilon),$$

where $A_X \equiv \text{diag}(\varphi_n, \varphi_\rho, \varphi_u)$ with $|\varphi_j| < 1$ and $\Sigma_\varepsilon \equiv \text{diag}(\sigma_n^2, \sigma_\rho^2, \sigma_u^2)$ is a positive-definite matrix.

5.2 Information and Equilibrium

The information structure is similar to the baseline model, except that the central bank's estimates are explicitly derived from noisy signals. Furthermore, the market's conditional mean and variance for the central bank's assessments are implied by the signal structure via a standard steady-state Kalman filter. Finally, all parameters of the signal structure are common knowledge and the information structure admits a standard common-prior foundation.

¹⁹We use “natural output” to denote the flexible-price output *absent markup shocks*. Since we assume the steady-state monopoly distortion is offset by a subsidy (see [Appendix A](#)), this coincides with the efficient level of output. However, it differs from the flexible-price output in the presence of cost-push shocks (indeed, under the standard microfoundation, the cost-push shock is proportional to the gap between the efficient level of output and its flexible-price counterpart). Hence, our terminology is different from that of [Blanchard and Galí \(2007\)](#) (among others), which uses “natural output” to refer to the flexible-price one.

We describe the core of the information structure below. Details can be found in [Appendix G](#).

Beliefs. The central bank does not observe the fundamental state directly. Instead, each period it observes private correctly centered noisy signals

$$s_t^C = X_t + e_t^C, \quad e_t^C \sim N(0, \Sigma_e),$$

where e_t^C is i.i.d. over time and independent of the fundamental innovations. Further, $\Sigma_e = \text{diag}((\sigma_n^e)^2, (\sigma_\rho^e)^2, (\sigma_u^e)^2)$, so each fundamental’s signal noise is independent from the others (as are fundamentals themselves). As in the baseline model, the central bank’s information consists only of these signals. In particular, it does not observe endogenous outcomes.²⁰

The market, on the other hand, observes each period’s current fundamentals X_t and the current policy rate i_t , but not the central bank’s private signal s_t^C .

Timing. Within a period, fundamentals X_t and central-bank signals s_t^C are realized first. The market observes X_t , while the central bank observes s_t^C and revises its beliefs. Then, the central bank sets the policy rate i_t . Finally, the market observes i_t , revises its beliefs, and current output and inflation are determined by (25)–(26). In other words, the timing within a period is analogous to the one described in [Figure 1](#), save for the inclusion of inflation and the substitution of the central bank’s assessments by signals.

Finite-history approximation. Since the central bank’s private assessments are persistent, agents optimally use the entire history to form their beliefs: the market uses fundamentals and policy rates to infer the central bank’s current assessment, while the central bank uses signals and rates to forecast market beliefs. Thus, our economy has no finite-dimensional sufficient state. For computation, we approximate the market’s conditional expectations in (25)–(26) by projections on truncated histories and, correspondingly, shrink the central bank’s posterior state. Besides being a useful approximation, this has a natural counterpart in the U.S., since Greenbook/Tealbook estimates—the proxies for the central bank’s private information we use—are publicly released with a lag.²¹

Our baseline estimation sets the finite history length to $H = 8$ quarters. Expanding the length of the finite history to $H = 12$ leaves the headline results unchanged, although some

²⁰If the central bank could observe past endogenous outcomes, it would change its policy to alter the informativeness of signals. Since this is not the focus of the current paper and its empirical relevance is doubtful, we shut down this channel (see [Appendix B](#) for an extended discussion).

²¹See [Appendix G.1](#) for a detailed discussion on the finite-history approximation.

individual coefficients and intermediate magnitudes differ.²²

Discretion and commitment. Discretionary equilibria and commitment are defined analogously to the baseline model. For a given approximation horizon H , market forecasts are given by the (population) linear projections implied by the policy rule and the retained histories of fundamentals and rates. A discretionary equilibrium is a fixed point between this forecast rule and the central bank’s best response, which takes the forecasting rule as given. Under (linear) commitment, the central bank instead chooses ex ante a linear rule on the same finite- H information state and internalizes the forecast rule it induces.²³ [Appendix G.2](#) provides the formal definitions while details on the solution method can be found in [Appendix G.3](#).

5.3 Measurement and Estimation

We fix standard New Keynesian parameters to their values in [Giannoni and Woodford \(2005\)](#), $(\beta, \gamma, \kappa) = (0.99, 0.16, 0.024)$. We also fix $\vartheta_\pi = \vartheta_i = 1$. Equal inflation and output-gap weights are a common benchmark (plain “dual mandate”). There is less agreement on the rate-level weight, so we use one as a transparent reference value and report the result’s sensitivity to deviations from it in [Appendix H.2](#).

The remaining nine parameters—the innovation standard deviations, persistence parameters, and signal-noise ratios of the three fundamentals—are estimated using ten empirical moments taken from recent U.S. history: six macro moments (persistence and residual variance of inflation, the output gap, and the federal funds rate), one natural-output uncertainty moment $\mathcal{H}$, two Tealbook/Greenbook forecast-revision moments, and one inflation-swaps moment measuring the reaction of expected inflation to monetary-policy surprises.

Macro moments. The first set of targets contains six quarterly macro moments: AR(1) coefficients and AR(1) residual variances for inflation, the output gap, and the policy rate. The underlying series are real GDP, the PCE price index, and the effective federal funds rate from FRED (series GDPC1, PCEPI, and FEDFUNDS, respectively). We transform real GDP to log points and extract its cyclical component using the [Hamilton \(2018\)](#) filter

²²See [Appendix H.1](#) for details.

²³Throughout, we study stationary projection equilibria as defined in [Appendix G.2](#). Our focus on equilibria driven by the model’s fundamental and signal innovations is in the spirit of the minimum state variable approach of [McCallum \(1983\)](#). Establishing the unique implementation of the resulting allocation is beyond the scope of our analysis. In particular, our assumption that the central bank does not observe endogenous outcomes rules out standard direct feedback rules. More generally, as [Lubik et al. \(2023\)](#) show, imperfect information introduces additional difficulties in ruling out indeterminacy.

with parameters $h = 8$, $p = 4$. Inflation is the annualized quarter-over-quarter log change in the PCE price index, and the policy rate is measured in annual percentage points. Thus, the output gap is in percentage log-output points, while inflation and the policy rate are in annual percentage points. The sample contains 120 quarters from 1990Q1 through 2019Q4 (we omit the Covid period).

Central bank information moments. The second set of targets disciplines the precision of the central bank’s private assessments. First, to pin down the central bank’s uncertainty about natural output, we target the average one-sided filtered posterior standard deviation of potential output from Holston et al. (2023)’s model. We begin with the data vintage distributed with the official 2026Q1 HLW estimates, discard all observations after 2019Q4, and re-estimate the model’s three stages over 1961Q1–2019Q4. Holding the resulting parameter estimates fixed, we use the *one-sided* Stage-3 Kalman filter: the posterior covariance for each quarter t conditions only on observations up to t . We then average the posterior standard deviation of potential output over the 120 quarters from 1990Q1 through 2019Q4.²⁴ This gives $\mathcal{H} = 1.7578$, which we use as a proxy for central-bank uncertainty about natural output and impose exactly in the estimation.²⁵ Second, we use revisions to Tealbook (formerly Greenbook) estimates to discipline the central bank’s real-time uncertainty about inflation and real GDP growth (and thus, indirectly, about the underlying shock processes). These are prepared by the Federal Reserve Board Staff before each FOMC meeting. We retrieve these from the Federal Reserve Bank of Philadelphia’s Greenbook Data Sets.²⁶ The usable sample covers target quarters from 1996Q1 through 2019Q4 and contains 55 target quarters for each variable. In particular, for each target quarter τ , let z_τ^0 denote the current-quarter nowcast and z_τ^L the first estimate of the same quarter reported in a vintage *at least four quarters later*. We target the (centered) revision-variance share

$$R_z^{\text{data}} = \frac{\text{Var}(z_\tau^L - z_\tau^0)}{\text{Var}(z_\tau^L)}.$$

²⁴Although the state estimates are one-sided, this is not a real-time-vintage exercise: the historical observations come from the revised 2026Q1 dataset, and the parameters are estimated once through 2019Q4 (instead of recursively at each date). Unfortunately, the data in the official HLW estimates do not allow us to build forecasts that are fully based on real-time data alone.

²⁵We do so since this procedure does not yield a properly defined “sampling uncertainty” and thus it could not be weighted as the other moments. Further, it is not really a “data moment” since its construction uses a model.

²⁶Available at <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/greenbook>

This measures the size of subsequent revisions relative to the variation in the later-vintage estimates. In the model, its counterpart is the posterior-variance share

$$R_z^{\text{model}} = \frac{\text{Var}(z_\tau - \mathbb{E}_\tau^C z_\tau)}{\text{Var}(z_\tau)}.$$

Consequently, we treat the later-vintage estimate as a proxy for the ex-post realization and thus interpret the empirical revision-variance shares as proxies for the fraction of its variation that was uncertain to the central bank in real time.

Inflation swaps moment. Finally, we target a market expectations-reaction moment based on inflation swaps and monetary policy surprises as detailed in [Section 4](#).²⁷ This can be thought of as the key moment disciplining the strength of our belief-revision mechanism. Indeed, an analogous economy with imperfect but public central bank estimates would be unable to match this moment since the reaction of expectations of future allocations to policy surprises would be zero.

As described there, the event information and the measured OIS1Y changes come from USMPD while the daily inflation-swap rates come from Bloomberg. The event definitions and intraday windows are the same as in the empirical analysis.

We use the same event sample as the empirical analysis in the previous section, pooling all 318 events to estimate the inflation-swaps target. The sample contains scheduled announcements and minutes from August 10, 2004 through April 29, 2026 and excludes the entire calendar years 2020–2021. This leaves 318 events: 159 monetary-policy announcements and 159 minutes releases.

Concretely, the target is the response of the expected cumulative price change over the first five years following a monetary-policy surprise. We use the (expected) cumulative price response instead of the response of medium-term forward inflation because our model has no endogenous source of inflation persistence. Empirically, it is built from the reactions of the 0–2 year and 2–5 year inflation-swap forwards as follows

$$\beta_{1:20}^P = 2\beta_{1:8}^\pi + 3\beta_{9:20}^\pi,$$

where $\beta_{1:8}^\pi$ and $\beta_{9:20}^\pi$ denote the pooled full-sample 0–2 year and 2–5 year coefficients, using the specification in [Section 4](#) with a common intercept and slope across all events. The units are cumulative annualized percentage-point years per percentage-point policy-path surprise. In the model, $\beta_{1:20}^P$ is the analogous (population) regression of the market’s revision in expected

²⁷As acknowledged there, this treats changes in inflation-swap breakevens as a proxy for revisions in expected inflation, even though in principle they may also reflect changes in inflation risk premia.

cumulative price growth over the first five years on its revision in the expected one-year nominal policy path.²⁸ [Appendix G.3.4](#) gives its finite- H counterpart.

Estimation. We estimate the parameters by minimum distance. Since the targets combine different datasets, sampling units, and frequencies, an estimate of their joint sampling covariance matrix is not directly available. Thus, we use diagonal inverse-sampling-variance weights. On top of those, the weights on the two Tealbook moments and the inflation-swaps moment are multiplied by ten. This is a deliberate choice to over-weight the key moments associated with our model’s mechanism. In particular, absent it, the parameterization is mostly determined by the six macro moments, leaving little influence to the ones that directly discipline the mechanism at the core of our model.²⁹ [Appendix G.4](#) contains the estimation details. All model target moments, losses, and decompositions reported here and in the appendix are computed analytically.

5.4 Estimates and Model Fit

The estimated parameter values are shown in [Table 1](#).

Table 1: Baseline Estimated Parameters

Parameter	Value	Description
σ_u	1.0314	Cost-push innovation sd
σ_n	0.3855	Natural-output innovation sd
σ_ρ	0.0772	Demand innovation sd
φ_u	0.0000	Cost-push persistence
φ_n	0.9950	Natural-output persistence
φ_ρ	0.9950	Demand persistence
q_u	0.1195	Cost-push signal-noise ratio
q_n	5.6450	Natural-output signal-noise ratio
q_ρ	0.9148	Demand signal-noise ratio

Here, q_j , for $j \in \{u, n, \rho\}$ is the ratio of the signal-noise variance to the unconditional variance of the corresponding fundamental. The fitted signal structure implies steady-state posterior-variance shares of 0.0868 for the demand state and 0.2074 for natural output. Thus,

²⁸In the model, the central bank’s assessments are revealed through its rate decision alone. Empirically, policy-path news also reflects explicit communication. The target moment remains similar when minutes releases are excluded: using only the 159 policy announcements, including any associated press conference, gives $\beta_{1:20}^P = -2.092$ (paired-bootstrap standard error 0.674), while the full sample gives -2.449 .

²⁹More concretely, in a pure inverse-variance-weights estimation, the sensitivity of market expectations moments is missed by a significant margin. See [Appendix H.3](#) for the model moments resulting from such an estimation.

the central bank is estimated to have substantially more precise real-time information about demand and cost-push shocks than about natural output, consistent with the findings about the difficulty of measuring economic slack in real time (Orphanides, 2001; Orphanides and van Norden, 2002).

Natural-output and demand persistence reach the upper bound of 0.995, while cost-push persistence reaches its lower bound of zero. The high estimated persistence of natural output is consistent with the persistent movements in productive capacity (see, e.g., Holston et al. (2023), which allow permanent innovations to natural output and its trend growth). On the other hand, the very low estimated persistence of cost-push shocks has precedents in the literature (see, e.g. Sbordone (2007) and the estimates assuming rational expectations in Milani (2007); also used as a benchmark in Galí (2015)). The high demand persistence is less standard (albeit not unprecedented, see, e.g., Miao et al. (2020)). It may partly reflect the absence of endogenous propagation mechanisms such as consumption habits and investment adjustment costs.³⁰

Taken together, persistent fundamentals and noisy signals make assessment errors persistent. Following a one-time signal error, subsequent observations gradually correct the central bank’s assessment. At the estimated parameters, the effect of a natural-output signal error has a half-life of approximately sixteen quarters, while demand’s is only seven quarters (see Appendix G.3.1 for details). Since the two fundamentals have the same estimated persistence, this difference reflects the lower precision of natural-output information.

Table 2 presents the fit of the model in terms of the targeted moments. The largest miss is the AR(1) residual variance of the output gap. The two Tealbook moments and the inflation-swaps moment remain fairly close to their empirical counterparts.

³⁰This was already acknowledged by (Smets and Wouters, 2007, p. 597). They find that weakening these endogenous persistence mechanisms substantially increases the estimated persistence of demand disturbances using aggregate U.S. data.

Table 2: Targeted Moments Fit

Moment	Data	Model
AR(1): inflation	0.1321	0.2440
AR(1): output gap	0.8943	0.9219
AR(1): policy rate	0.9643	0.9121
Residual variance: inflation	2.7672	1.3397
Residual variance: output gap	1.1785	0.2946
Residual variance: policy rate	0.1650	0.2031
Natural-output uncertainty ($\mathcal{H}$)	1.7578	1.7578
Forecast-revision variance share: inflation	0.2414	0.2680
Forecast-revision variance share: real GDP growth	0.5876	0.6296
Market-expectations sensitivity $\beta_{1:20}^P$	-2.4486	-2.0117

5.5 The Role of Commitment

Commitment gains. We consider three policy regimes: discretion, commitment and a restricted-commitment regime where the central bank chooses an ex-ante linear rule that can only condition on its current assessments of the three fundamentals. Studying this simple rule serves two purposes. First, as we show below, it captures almost all of the commitment gain, making it useful for identifying which features of policy matter most for stabilization. Second, the comparison allows us to assess how much can be achieved without direct history dependence, a central feature of standard New Keynesian commitment policies. The results are summarized in [Table 3](#) below.

Table 3: Commitment Gains

Policy	Loss	Gain (% of D loss)	Share of full commitment gain
Discretion	2.29760	—	—
Commitment: simple rule	0.72762	68.33%	98.75%
Full commitment	0.70776	69.20%	100.00%

Full commitment lowers the stationary per-period loss by 69.20% and reduces the standard deviation of the output gap by 80.8% (see [Appendix G.5](#) for the full loss decomposition under each policy regime). The simple rule captures 98.75% of the full commitment gain. Thus, almost all of the loss reduction can be achieved by changing how strongly policy responds to the central bank’s current assessments. By construction, conditional on the central bank’s current assessments, the simple rule implements the same rate regardless of past assessments and rates (and, furthermore, does not incorporate a Ramsey “promise state”). Consequently, the excellent performance of this rule documented above shows that the source

of the large commitment gains is not the ability to honor past promises or to compensate for past policy mistakes caused by poor information (see, e.g., [Woodford \(2003\)](#) and [Aoki \(2006\)](#) respectively). A key question, then, is how changing the systematic response to current assessments alone delivers such a substantial improvement.

Changes in systematic policy. [Table 4](#) shows how the on-impact interest-rate response to signal innovations varies across regimes.³¹ Both commitment policies respond substantially less to each signal innovation. Furthermore, under the simple rule the absolute response to a natural-output signal innovation falls by 95.8%, compared with 78.2% for a demand-signal innovation. Thus, commitment combines an overall lower aggressiveness with a particularly sharp reduction in the response to natural output relative to demand.

Table 4: Impact Policy Responses to Signal Innovations

Signal	Discretion	Simple rule	Full commitment
Natural output	−0.0706	−0.0030	−0.0041
Demand	0.1622	0.0353	0.0208
Cost-push	0.0923	0.0027	0.0016

Notes: Interest-rate responses are in annual percentage points. Each row uses the same positive signal innovation across policy regimes, with its size set to one standard deviation of the corresponding signal noise.

As in the baseline model, a discretionary central bank takes the market’s forecasting structure as given and uses the reaction of market expectations to stabilize the economy as it perceives it (recall that its problem is linear-quadratic and certainty equivalence applies). Commitment instead internalizes how its systematic policy shapes that forecasting structure and, consequently, the transmission of assessment errors into allocations.

In this richer model, the market’s forecasting structure depends on policy via two different channels: policy determines both what markets infer about the central bank’s assessments from its decisions and how those inferred assessments translate into expectations of future policy. The first channel is the belief-revision mechanism emphasized in the baseline model. The second introduces an additional margin: commitment can change how future assessments map into future policy, a possibility deliberately excluded from the baseline analysis.³² The following illustration shows how these relationships jointly shape the propagation of assess-

³¹Under discretion and full commitment, policy depends on the central bank’s current estimates of the past values of fundamentals as well as their current values. A new signal about a persistent fundamental prompts an update of these past-value estimates too, so the on-impact policy response does not in general equal the corresponding estimate’s coefficient times the revision in that estimate.

³²In the baseline model, commitment was restricted to the $t = 0$ rule, and thus the mapping from assessments into subsequent policy was pinned down by the discretionary equilibrium for $t \geq 1$.

ment errors. The subsequent exercises establish the distinctive role of our belief-revision mechanism.

The policy reaction to a perceived demand pressure: an illustration. To make the mechanism above concrete, consider a central bank whose information suggests a positive demand pressure due to demand-signal noise.³³ A discretionary central bank naturally raises its interest rate to counteract the perceived expansionary pressure. When doing this, it anticipates that it may surprise the market and thus part of its intended stabilization will be achieved via the market revising its expectations of output and inflation downward. In doing this, it naturally takes the continuation equilibria as given (i.e., its deviation is “one off” in game-theoretic terms).

Commitment changes both the size of the policy response and the market forecasting structure through which it operates. Under the simple rule, policy responds much less to demand assessments, but its response to natural output estimates is reduced even more sharply. An unexpected tightening, thus, mostly conveys news about the central bank’s demand assessment.

These changes jointly shape how markets interpret policy and how that interpretation translates into expectations of future policy. Concretely, following a one standard deviation demand-signal error, the nominal increase falls from 0.162 under discretion to 0.035 under the simple rule, while the resulting output-gap contraction shrinks from 0.299 to 0.085. Thus, commitment prevents a large share of the assessment error’s passthrough into allocations. [Figure 4](#) shows the responses over time, together with the corresponding responses to a natural-output signal error.

Note that commitment also improves stabilization following a genuine demand innovation, which the central bank initially acknowledges only partially (and hence it also induces an assessment error). Following a one standard deviation demand innovation, the positive output-gap response on impact shrinks from 0.323 under discretion to 0.060 under the simple rule.³⁴

We next show that changing the market’s interpretation can improve stabilization without changing the realized interest rate path. Starting from the discretionary policy, we replace only the coefficients on current and past natural-output assessments with their simple-rule counterparts. Following a one-standard-deviation demand-signal error, the central bank’s

³³Note that demand disturbances (signal errors or fundamental innovations) account for 67.7% of the loss reduction from discretion to the simple rule, compared with 28.6% for natural-output disturbances and 3.7% for cost-push disturbances.

³⁴The contribution of fundamental demand innovations to the loss falls by 81.7%. [Appendix G.5.3](#) contrasts the responses to fundamental innovations and signal errors.




Figure 4: Responses to positive one-standard-deviation signal errors in natural output (top row) and demand (bottom row).

natural-output assessments remain unchanged. Since all other policy coefficients are unchanged, the two rules generate exactly the same realized nominal-rate path, as shown in Panel A of [Figure 5](#).

While the rate decision is the same, markets interpret it differently, revising their estimate of the central bank’s natural-output assessment downward much less under the counterfactual rule. On impact, this inferred decline is 0.162 under discretion and only 0.010 under the counterfactual, while the central bank’s actual natural-output assessment is unchanged (Panel B). Because natural-output assessment errors are corrected slowly, attributing a tightening to such an error has persistent implications for expected future policy. Thus, expected next-period output and inflation fall less, and the output-gap contraction on impact is 41.3% smaller (Panel C). Note that if the central bank assessments were public, changing the natural-output response would not affect the macroeconomic consequences of the demand error. [Appendix G.5.6](#) provides the details of this exercise.

These findings complement the case for caution in responding to noisy real-time information, especially estimates of economic slack, emphasized by [Orphanides \(2001, 2003\)](#). In the




Figure 5: Responses to a positive one-standard-deviation demand-signal error. The counterfactual changes only the response to natural-output assessments. Panel B shows the market’s impact revision to its estimate of the central bank’s natural-output assessment.

optimal-control benchmark of the latter paper, policy applies the full-information response coefficients to optimally filtered estimates: measurement uncertainty does not warrant further attenuation of the response to those estimates (i.e., certainty equivalence holds). In our framework, substantial stabilization gains remain available from weakening policy responses to the same optimally formed assessments. This additional restraint limits the amplification of assessment errors through the market’s interpretation of policy. The concern is particularly relevant for natural-output assessments: their low precision makes mistakes slow to correct, so policy decisions attributed to them have lasting implications for expected future policy.

The response to cost-push assessments. The distinction between a unilateral policy adjustment and a systematic rule change also helps explain the weaker response to cost-push assessments under both commitment policy regimes. Taking the market’s forecasting rule as given, a central bank facing perceived cost-push pressure can use a tightening to lower expected inflation and ease its perceived stabilization trade-off. This signal-jamming incentive (see, e.g., [Stein \(1989\)](#)) helps explain the stronger cost-push response under discretion. In equilibrium, however, markets account for the strength of the cost-push response when interpreting policy surprises. A stronger response makes an unexpected tightening less informative about persistent demand or natural-output assessments, partly offsetting the expectations benefit that motivates the unilateral deviation.

This change in interpretation also affects the propagation of errors in those other assess-

ments. At our estimates, the much stronger discretionary response delivers little additional inflation stabilization following a genuine cost-push innovation, but produces a much larger contraction following a cost-push signal error. Commitment limits these costs by choosing a weaker cost-push response jointly with the responses to the other assessments. Note, however, that keeping the discretionary cost-push coefficients while adopting the simple rule’s remaining coefficients delivers 97.9% of the simple rule’s gain over discretion. Hence, the change in the reaction to cost-push shocks is not central to the commitment gains. [Appendix G.5.4](#) illustrates this incentive numerically and examines how the benefits of attenuating the cost-push response depend on the other policy responses.

The pass-through of assessment errors. In our model, policy decisions affect relevant outcomes both directly and through shaping the market expectations of future policy. As in the theory, this creates a trade-off between stabilization conditional on the central bank’s information and the propagation of assessment errors into allocations (output and inflation). To map the quantitative results to the loss decomposition in [Lemma 2](#), note that we can write the unconditional output gap volatility in a given period t as

$$\text{Var}(x_t) = \text{Var}\left(\mathbb{E}_t^C x_t\right) + \mathbb{E}\left[\text{Var}(x_t \mid \mathcal{I}_t^C)\right].$$

The first term represents the variation in the output gap anticipated by the central bank at the time it chooses rates (i.e., after having formed its estimates of the corresponding period). The second, on the other hand, measures departures of the realized output gap from that conditional estimate, i.e., the quantitative counterpart of $\mathcal{P}$ in the baseline model.

The simple rule reduces the second component by 97.8%, while leaving variation in the central bank’s expected gap almost unchanged. It’s worth stressing that the central bank’s information is identical across regimes: what changes is the propagation of assessment errors to allocations (see [Appendix G.5](#) for the full decomposition of output-gap variance under each policy regime).

Policy surprises and expectation revisions. The differences between discretion and commitment in the demand-signal error illustration above are reflected in how policy decisions shift inflation expectations. Since markets already observe fundamentals, these policy-induced belief revisions concern the central bank’s assessments and their implications for future policy. Thus, they are an observable counterpart of the belief-revision mechanism emphasized in the theory.

We consider three statistics informative of this margin. First, the object used in the empirical analysis and one of the targeted moments in the estimation above, $\beta_{1:20}^P$. As

discussed above, it measures the revision in expected cumulative price growth over five years upon a revision of beliefs about the coming one-year policy path. Second, the standard deviation of policy surprises, a measure of the “amount” of such policy news. Finally, we consider the product of these two, which equals the standard deviation of fitted expectation revisions, a measure of the overall pass-through of policy surprises into expectations.

$$\text{SD}\left(\widehat{\Delta P}_{t,1:20}^e\right) = |\beta_{1:20}^P| \text{SD}(mps_t^M).$$

Table 5 reports these statistics across the three policy regimes.

Table 5: Policy News and Expectations Revisions

Policy	$\beta_{1:20}^P$	Policy-surprise SD	Fitted revision SD
Discretion	-2.01174	0.11439	0.23013
Commitment: simple rule	-1.26820	0.03055	0.03875
Full commitment	-0.29370	0.04796	0.01409

In the estimated economy, commitment substantially reduces the sensitivity of market expectations to policy news. The simple rule lowers the magnitude of the cumulative five-year price expectations reaction to a one-percentage-point policy-path surprise from 2.01 to 1.27, a reduction of 37%. Full commitment lowers it further to 0.29, an 85% reduction relative to discretion. Together with smaller policy surprises, these reduce the standard deviation of fitted expectation revisions by 83% and 94%, respectively.

The simple rule nevertheless captures almost all the commitment gain despite a much smaller reduction in the absolute magnitude of $\beta_{1:20}^P$ than under full commitment. Its performance reflects the combined changes in policy responses, the composition of policy news, and the future policy implied by market inferences. The sensitivity coefficient alone does not summarize these changes.

It should be noted that the desirability of these reductions depends on the balance between assessment-error amplification and the stabilization benefits of responsive expectations, which are deliberately excluded from the baseline model to isolate the mechanism of interest. First, shifting expected output and inflation allows policy to influence the economy with smaller rate deviations, decreasing the rate-level penalty given $\vartheta_i > 0$. Second, with cost-push shocks, changes in expected inflation can ease the trade-off between output-gap and inflation stabilization. Commitment balances these benefits against the propagation of assessment errors through the forecasting structure its policy rule induces.

We also compare the gains from commitment to the simple rule with the corresponding gains in two alternative economies: a full information one and one with imperfect but public

central bank information, keeping the physical environment and corresponding parameters unchanged (except, of course, the ones governing the central bank's information quality, which are moot in the full information economy). The loss reduction relative to discretion is 34.81% under full information and 34.06% under public assessments, compared with 68.33% in our benchmark economy. The main difference is the substantially worse performance of discretion when assessments are private (commitment losses are fairly similar across these economies), which highlights the acuteness of the time inconsistency problem in economies where the transmission of policy is partly via changing the beliefs of the private sector.³⁵

Together, these results establish a quantitatively important role for the mechanism emphasized in the theory. Commitment does not only reduce the rate movements generated by mistaken assessments; it also changes how a given policy decision is interpreted and transmitted into macroeconomic outcomes. The following subsection examines which features of the data determine both the magnitude of the commitment gains and the changes in policy-induced expectation revisions.

5.6 Which data features drive the results?

To assess which data features drive our key results, we move individual target moments from their point estimates to the endpoints of their 95% confidence intervals. For each perturbation, we re-estimate all nine structural parameters using the otherwise unchanged specification and recompute the discretionary equilibrium and the optimal simple and full commitment policies. [Table 6](#) reports the two perturbations that most affect the headline results.

We let G be the (full) commitment gain as a percentage of the discretionary loss and S the share of that gain captured by the simple rule. The two remaining columns report

$$R_\beta = \frac{|\beta_{1:20,C}^P|}{|\beta_{1:20,D}^P|}, \quad R_i = \frac{\text{SD}(mps_C^M)}{\text{SD}(mps_D^M)},$$

where C and D denote full commitment and discretion, respectively. Clearly, values above one indicate that full commitment increases the absolute response of expectations per unit of policy news (R_β) or the volatility of policy surprises (R_i).

³⁵See [Appendix G.5.5](#) for the losses under both policy regimes and the associated gains from commitment in each economy.

Table 6: Targeted Moments and Commitment vs. Discretion

Exercise	G	S	R_β	R_i
Baseline	69.20%	98.75%	0.146	0.419
Less-negative swaps endpoint	49.33%	71.82%	1.374	1.161
Lower GDP-revision endpoint	71.58%	95.75%	0.345	0.616

The inflation-swaps target is key for the direction in which commitment shifts the sensitivity of expectations. Moving its requested value from -2.449 to the less-negative endpoint of its confidence interval, -0.795 , moves the fitted discretionary coefficient from -2.012 to -0.793 . With weaker expectations transmission under discretion, full commitment chooses to increase both absolute sensitivity and policy-surprise volatility. This is consistent with the notion discussed above that in this model, sensitive market expectations have an upside. Although the proportional commitment gains fall, they remain substantial. Furthermore, the simple rule captures a lower but still significant share of that gain. Consequently, direct history dependence contributes more to the commitment gain than at the baseline estimates.

Importantly, in this richer model managing the belief-revision channel involves more than the one-dimensional margin of choosing how strongly market expectations respond to policy news. The policy rule also determines which assessment errors generate that news—and thus, what exactly markets infer from it—and how those inferences translate into expected future policy and allocations.

As an additional check, moving the GDP-growth revision-variance target from 0.588 to its lower confidence interval endpoint of 0.308 produces a fitted value of 0.265 (i.e., one that corresponds to a central bank with better real-time information about output). The commitment gain then slightly increases relative to the baseline. Furthermore, the share of gains the simple rule attains is essentially the same as in the baseline, and full commitment lowers both the expectations sensitivity and policy-surprise volatility.

Overall, the quantification strongly supports the relevance of the belief-revision mechanism for U.S. monetary policy. The commitment gains are economically large across the specifications considered, and almost all of the baseline gain is attained without direct dependence on lagged assessments or policy rates. Together with the sharp reductions in expectation revisions and signal-noise pass-through, this points to managing the market’s inference from policy as a central source of commitment gains (as opposed, for example, to conventional history-dependent “make-up” policy). At our baseline estimates, most of these gains are attained through more restrained responses to current assessments, particularly a weaker response to natural output relative to demand. Natural-output assessments are especially noisy and thus their errors take longer to correct, so attributing policy news to

them has persistent consequences for expected future policy.

6 Conclusion

Monetary policy works largely by shaping expectations. Thus, its effects depend on how the private sector interprets a given decision. That interpretation depends, in turn, on how the central bank systematically conducts policy. This paper studies the implications for policy design. We show that a central bank choosing its rule *ex ante* internalizes this dependence and reduces the output-gap volatility arising from the market's reaction to its imperfect real-time assessments.

Our empirical analysis provides strong support for a key premise of the mechanism: monetary policy surprises substantially shift expectations of future inflation. Furthermore, we document variation in expectations sensitivity across time subsamples, a phenomenon suggestive of a dependence on the policy regime and consistent with our model's central mechanism. In an estimated New Keynesian model incorporating this mechanism, commitment reduces the policy loss by roughly two-thirds. A rule responding only to current assessments captures almost all of this gain. Further counterfactuals show that changing how markets interpret policy accounts for an important part of the gain from commitment. At our baseline estimates, the policy that achieves this lowers the sensitivity of market expectations to policy surprises.

More broadly, the familiar image of monetary policy as a lever misses that a rate decision is also an act of communication, whether intended or not: to a large extent, market participants react based on why they think the decision was made and what that implies for the future. Effective policy design must account for this interpretation. A natural next step is to study how this implicit communication through policy decisions interacts with explicit central-bank communication: when can explaining a decision improve the inferences drawn from it, and how does that change the design of policy itself?

References

- Abrahams, M., Adrian, T., Crump, R. K., Moench, E. and Yu, R. (2016), ‘Decomposing real and nominal yield curves’, *Journal of Monetary Economics* **84**, 182–200.
- Acosta, M., Ajello, A., Bauer, M., Loria, F. and Miranda-Agrippino, S. (2025), Financial market effects of FOMC communication: Evidence from a new event-study database, Working Paper 30, Federal Reserve Bank of San Francisco.
- Ajello, A., Benzoni, L. and Chyruk, O. (2020), ‘Core and ‘crust’: Consumer prices and the term structure of interest rates’, *The Review of Financial Studies* **33**(8), 3719–3765.
- Aoki, K. (2003), ‘On the optimal monetary policy response to noisy indicators’, *Journal of Monetary Economics* **50**(3), 501–523.
- Aoki, K. (2006), ‘Optimal commitment policy under noisy information’, *Journal of Economic Dynamics and Control* **30**(1), 81–109.
- Baeriswyl, R. and Cornand, C. (2010), ‘The signaling role of policy actions’, *Journal of Monetary Economics* **57**(6), 682–695.
- Bauer, M. D., Pflueger, C. E. and Sunderam, A. (2024), ‘Perceptions about monetary policy’, *The Quarterly Journal of Economics* **139**(4), 2227–2278.
- Bauer, M. D. and Swanson, E. T. (2023), ‘An alternative explanation for the Fed Information Effect’, *American Economic Review* **113**(3), 664–700.
- Blanchard, O. and Galí, J. (2007), ‘Real Wage Rigidities and the New Keynesian Model’, *Journal of Money, Credit and Banking* **39**(s1), 35–65.
- Caballero, R. J. and Simsek, A. (2022), ‘Monetary policy with opinionated markets’, *American Economic Review* **112**(7), 2353–2392.
- Calvo, G. A. (1978), ‘On the time consistency of optimal policy in a monetary economy’, *Econometrica* **46**(6), 1411–1428.
- Clarida, R., Galí, J. and Gertler, M. (1999), ‘The science of monetary policy: A New Keynesian perspective’, *Journal of Economic Literature* **37**(4), 1661–1707.
- Dogra, K. (2024), Paradoxes and problems in the causal interpretation of equilibrium economics, Staff Report 1093, Federal Reserve Bank of New York.

- Evans, G. W. and Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Fleckenstein, M., Longstaff, F. A. and Lustig, H. (2014), ‘The TIPS–Treasury bond puzzle’, *Journal of Finance* **69**(5), 2151–2197.
- Galí, J. (2015), *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework and Its Applications*, 2nd edn, Princeton University Press.
- García-Schmidt, M. and Woodford, M. (2019), ‘Are low interest rates deflationary? A paradox of perfect-foresight analysis’, *American Economic Review* **109**(1), 86–120.
- Giannoni, M. P. and Woodford, M. (2005), Optimal inflation-targeting rules, *in* B. S. Bernanke and M. Woodford, eds, ‘The Inflation-Targeting Debate’, University of Chicago Press, Chicago, pp. 93–162.
- Gürkaynak, R. S., Sack, B. and Wright, J. H. (2007), ‘The U.S. Treasury yield curve: 1961 to the present’, *Journal of Monetary Economics* **54**(8), 2291–2304.
- Gürkaynak, R. S., Sack, B. and Wright, J. H. (2010), ‘The TIPS yield curve and inflation compensation’, *American Economic Journal: Macroeconomics* **2**(1), 70–92.
- Hamilton, J. D. (2018), ‘Why you should never use the Hodrick–Prescott filter’, *Review of Economics and Statistics* **100**(5), 831–843.
- Haubrich, J., Pennacchi, G. and Ritchken, P. (2012), ‘Inflation expectations, real rates, and risk premia: Evidence from inflation swaps’, *Review of Financial Studies* **25**(5), 1588–1629.
- Holmstrom, B. (1982), Managerial incentive problems: A dynamic perspective, *in* ‘Essays in Economics and Management in Honor of Lars Wahlbeck’.
- Holston, K., Laubach, T. and Williams, J. C. (2023), Measuring the natural rate of interest after COVID-19, Staff Report 1063, Federal Reserve Bank of New York.
- Jarociński, M. and Karadi, P. (2026), Monetary Policy Surprises with Imperfect Information.
- Jia, C. (2023), ‘The informational effect of monetary policy and the case for policy commitment’, *European Economic Review* **156**, 104468.
- Kydland, F. E. and Prescott, E. C. (1977), ‘Rules rather than discretion: The inconsistency of optimal plans’, *Journal of Political Economy* **85**(3), 473–491.

- Lippi, F. and Neri, S. (2007), ‘Information variables for monetary policy in an estimated structural model of the euro area’, *Journal of Monetary Economics* **54**(4), 1256–1270.
- Lubik, T. A., Matthes, C. and Mertens, E. (2023), ‘Indeterminacy and imperfect information’, *Review of Economic Dynamics* **49**, 37–57.
- McCallum, B. T. (1983), ‘On Non-uniqueness in Rational Expectations Models: An attempt at perspective’, *Journal of Monetary Economics* **11**, 139–168.
- Melosi, L. (2017), ‘Signalling effects of monetary policy’, *The Review of Economic Studies* **84**(2), 853–884.
- Mertens, E. (2016), ‘Managing beliefs about monetary policy under discretion’, *Journal of Money, Credit and Banking* **48**(4), 661–698.
- Miao, J., Wang, P. and Zha, T. (2020), ‘Discount shock, price–rent dynamics, and the business cycle’, *International Economic Review* **61**(3), 1229–1252.
- Milani, F. (2007), ‘Expectations, learning and macroeconomic persistence’, *Journal of Monetary Economics* **54**(7), 2065–2082.
- Morris, S. (1995), ‘The common prior assumption in economic theory’, *Economics and Philosophy* **11**(2), 227–253.
- Orphanides, A. (2001), ‘Monetary policy rules based on real-time data’, *American Economic Review* **91**(4), 964–985.
- Orphanides, A. (2003), ‘Monetary policy evaluation with noisy information’, *Journal of Monetary Economics* **50**(3), 605–631.
- Orphanides, A. and van Norden, S. (2002), ‘The unreliability of output-gap estimates in real time’, *The Review of Economics and Statistics* **84**(4), 569–583.
- Preston, B. (2005), ‘Learning about monetary policy rules when long-horizon expectations matter’, *International Journal of Central Banking* **1**(2), 81–126.
- Sbordone, A. M. (2007), ‘Inflation persistence: Alternative interpretations and policy implications’, *Journal of Monetary Economics* **54**(5), 1311–1339.
- Smets, F. and Wouters, R. (2007), ‘Shocks and frictions in US business cycles: A Bayesian DSGE approach’, *American Economic Review* **97**(3), 586–606.

- Stein, J. C. (1989), ‘Efficient capital markets, inefficient firms: A model of myopic corporate behavior’, *The Quarterly Journal of Economics* **104**(4), 655–669.
- Stein, J. C. (2014), ‘Challenges for monetary policy communication’. Remarks at the Money Marketeers of New York University, New York, New York.
- Stein, J. C. and Sunderam, A. (2018), ‘The Fed, the bond market, and gradualism in monetary policy’, *Journal of Finance* **73**(3), 1015–1060.
- Svensson, L. E. and Woodford, M. (2003), ‘Indicator variables for optimal policy’, *Journal of Monetary Economics* **50**(3), 691–720.
- Svensson, L. E. and Woodford, M. (2004), ‘Indicator variables for optimal policy under asymmetric information’, *Journal of Economic Dynamics and Control* **28**(4), 661–690.
- Townsend, R. M. (1983), ‘Forecasting the forecasts of others’, *Journal of Political Economy* **91**(4), 546–588.
- Warsh, K. (2026), ‘Transcript of chairman Warsh’s press conference’. Board of Governors of the Federal Reserve System.
- Woodford, M. (1999), ‘Optimal monetary policy inertia’, *The Manchester School* **67**(S1), 1–35.
- Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, Princeton, N.J.

Appendix

A	Microfoundation of the Baseline Model	49
A.1	Environment	49
A.2	Approximations	51
B	Information Structure	55
B.1	General structure	55
B.2	Gaussian beliefs and further assumptions in the main text	58
B.3	Common-prior restrictions	63
C	Game-Theoretic Representation and Solution Concepts	67
C.1	The game	67
C.2	Market-consistent profiles	69
C.3	Discretion	70
C.4	Commitment	71
D	Proofs	73
D.1	Proof of Lemma 1	73
D.2	Proof of Proposition 1	73
D.3	Proof of Proposition 2	75
D.4	Proof of Lemma 2	78
D.5	Proof of Proposition 3	79
D.6	Proof of Corollary 1	82
E	Additional Theoretical Results	82
F	Additional Empirical Results	87
F.1	Alternative Policy-Surprise Measure	87
F.2	TIPS-Based Inflation Compensation	88
F.3	Bid and Ask Prices	90
G	Quantitative Model Details	91
G.1	Information and Forecasting	91
G.2	Definition of Equilibria	94
G.3	Solving the Equilibrium and Computing Population Objects	96
G.4	Estimation Details	106
G.5	Additional Results	108
H	Robustness Exercises	115
H.1	History-length sensitivity	115
H.2	Rate-Level Stabilization Weight	116
H.3	Pure Inverse-Variance Weighting	117

A Microfoundation of the Baseline Model

This appendix shows that the model in the main text can be obtained as a linear approximation of a standard rigid-prices economy.

A.1 Environment

Households. Consider a representative household with preferences over consumption C_t and labor N_t given by

$$\mathbb{E}_0^M \sum_{t=0}^{\infty} \left(\prod_{j=0}^{t-1} \delta_j \right) \left[\frac{C_t^{1-\nu} - 1}{1-\nu} - \zeta_t \frac{N_t^{1+\nu}}{1+\nu} \right], \quad \nu > 0, \quad \nu > 0, \quad (28)$$

where δ_j is the one-period discount factor, the source of demand shocks. Households trade a one-period nominal bond with risk-free gross nominal return R_t set by the central bank and face the usual budget constraint

$$P_t C_t + \frac{B_{t+1}}{R_t} \leq W_t N_t + B_t + \Pi_t + T_t, \quad (29)$$

where P_t is the price of the final consumption good, W_t is the nominal wage, B_t is the household's bond holdings, Π_t denotes firm profits and T_t collects lump-sum taxes and transfers from the government. The household's Euler equation is

$$C_t^{-\nu} = \delta_t \mathbb{E}_t^M \left[R_t \frac{P_t}{P_{t+1}} C_{t+1}^{-\nu} \right], \quad (30)$$

where we denote by $\mathbb{E}_t^M$ the expectations operator corresponding to the representative household's beliefs (see [Appendix B](#) for details on the model's information structure). Importantly, the consumption decision is taken after the central bank sets the policy rate, so the expectation in (30) includes R_t in its conditioning information set.

Firms. As in the standard New Keynesian model, there is a continuum of intermediate goods producers $i \in [0, 1]$ with production technology

$$Y_t(i) = A_t N_t(i), \quad (31)$$

where A_t is an aggregate productivity parameter and $N_t(i)$ is the labor input of firm i . In addition, a representative final goods producer combines differentiated intermediate goods

according to

$$Y_t = \left(\int_0^1 Y_t(i)^{(\epsilon-1)/\epsilon} di \right)^{\epsilon/(\epsilon-1)}, \quad \epsilon > 1, \quad (32)$$

As standard, we assume an employment subsidy that exactly offsets the effect of the firms' monopolistic incentives on production and hiring is in place (paid for with lump-sum transfers).

Intermediate goods firms can reset their prices with probability $1 - \theta$ as in the standard Calvo model.

We assume the initial price distribution is symmetric and coincides with the steady-state price level (see below). Further, we consider the fully rigid-price limit $\theta \uparrow 1$. Consequently, $P_t(i) = P_t = \bar{P}$ for all i and t , i.e., there is no inflation nor price dispersion.

Since there are neither government purchases nor investment, goods-market clearing is given by $C_t = Y_t$. Furthermore, given prices are fully rigid, output is demand-determined.

As standard, we define natural output Y_t^n as the equilibrium output in an equivalent economy with fully flexible prices. Because of the employment subsidy, the flexible-price allocation is efficient and thus Y_t^n is the efficient level of output. Hence, we can find it by simply equating the marginal rate of substitution between consumption and labor to the marginal product of labor. It's straightforward to verify this gives

$$(Y_t^n)^{-\nu} = \frac{\zeta_t (N_t^n)^\nu}{A_t}, \quad (33)$$

Substituting $N_t^n = Y_t^n / A_t$ and rearranging yields

$$Y_t^n = \left(\frac{A_t^{1+\nu}}{\zeta_t} \right)^{1/(v+\nu)}.$$

In the baseline model, we assume a process directly for Y_t^n . We can think of Y_t^n as a reduced-form representation of the underlying shocks, and the law of motion for log natural output y_t^n in (3) can be generated by appropriate processes for (A_t, ζ_t) .

The market's observation of fundamentals is consistent with this microfoundation: households observe their preference shifters, firms observe productivity, and private-sector agents share a common information set. Demand and natural output are thus known to the private sector.

The (nonstochastic) steady state. A non-stochastic steady state is such that exogenous variables are constant, $A_t = \bar{A}$, $\zeta_t = \bar{\zeta}$, and $\delta_t = \bar{\delta} \equiv \beta \in (0, 1)$. Consequently, endogenous variables are constant as well.

Since prices are constant, the steady-state level of output coincides with the flexible prices one, natural output. Since by the clearing of the goods market we further have $\bar{C} = \bar{Y}$, the steady state is characterized by

$$\bar{C} = \bar{Y} = \bar{Y}^n = \left(\frac{\bar{A}^{1+\nu}}{\bar{\zeta}} \right)^{1/(\nu+\nu)}, \quad \bar{N} = \frac{\bar{Y}}{\bar{A}}, \quad (34)$$

Furthermore, evaluating the Euler [equation \(30\)](#) at constant consumption and prices gives

$$\bar{R} = \beta^{-1}. \quad (35)$$

Define $\bar{\rho} \equiv -\log \beta$. The steady-state log nominal rate is then $\log \bar{R} = \bar{\rho}$. Finally, the price level $\bar{P}$ can simply be normalized to 1.

A.2 Approximations

Demand. We now log-linearize the household Euler equation. Define consumption and output as log deviations from their steady-state values,

$$c_t \equiv \log \left(\frac{C_t}{\bar{C}} \right), \quad y_t \equiv \log \left(\frac{Y_t}{\bar{Y}} \right), \quad (36)$$

while keeping the interest rate and discount rate as simply log levels

$$r_t \equiv \log R_t, \quad \rho_t \equiv -\log \delta_t. \quad (37)$$

with log deviations from steady state

$$\tilde{r}_t \equiv r_t - \bar{\rho}, \quad \tilde{\rho}_t \equiv \rho_t - \bar{\rho}. \quad (38)$$

Since

$$\frac{C_t^{-v}}{\bar{C}^{-v}} = \exp(-vc_t), \quad \frac{\delta_t}{\beta} = \exp(-\tilde{\rho}_t), \quad \frac{R_t}{\bar{R}} = \exp(\tilde{r}_t),$$

we can divide [\(30\)](#) by $\bar{C}^{-v}$ and use $\beta\bar{R} = 1$ and $P_{t+1} = P_t$ to get

$$\exp(-vc_t) = \exp(\tilde{r}_t - \tilde{\rho}_t) \mathbb{E}_t^M [\exp(-vc_{t+1})]. \quad (39)$$

Let $\|\xi\|$ denote the size of the log deviations from the steady state. Doing a first-order Taylor expansion of [\(39\)](#) around the steady state gives

$$1 - vc_t + \mathcal{O}(\|\xi\|^2) = \left[1 + \tilde{r}_t - \tilde{\rho}_t + \mathcal{O}(\|\xi\|^2) \right] \left[1 - v \mathbb{E}_t^M c_{t+1} + \mathcal{O}(\|\xi\|^2) \right]$$

$$= 1 + \tilde{r}_t - \tilde{\rho}_t - v \mathbb{E}_t^M c_{t+1} + \mathcal{O}(\|\xi\|^2), \quad (40)$$

where the cross-terms are second order and thus included in $\mathcal{O}(\|\xi\|^2)$. Rearranging and dropping the second-order remainder yields

$$c_t = \mathbb{E}_t^M c_{t+1} - \frac{1}{v} (\tilde{r}_t - \tilde{\rho}_t). \quad (41)$$

Finally, $\tilde{r}_t - \tilde{\rho}_t = (r_t - \bar{\rho}) - (\rho_t - \bar{\rho}) = r_t - \rho_t$. Consequently, the log-linearized Euler equation is

$$c_t = \mathbb{E}_t^M c_{t+1} - \frac{1}{v} (r_t - \rho_t). \quad (42)$$

Using the equilibrium goods market clearing condition $c_t = y_t$ and defining the intertemporal elasticity of substitution $\gamma \equiv \frac{1}{v}$, the [equation \(42\)](#) becomes

$$y_t = \mathbb{E}_t^M y_{t+1} - \gamma(r_t - \rho_t), \quad (43)$$

which is exactly [\(1\)](#).

The central bank's loss function. Substituting the equilibrium condition $C_t = Y_t$, each period's household utility can be written as a function of output and labor

$$\frac{Y_t^{1-v} - 1}{1-v} - \frac{\zeta_t}{1+\nu} (N_t)^{1+\nu}. \quad (44)$$

Define the log natural output and the log output gap

$$y_t^n \equiv \log\left(\frac{Y_t^n}{\bar{Y}}\right), \quad x_t \equiv \log\left(\frac{Y_t}{Y_t^n}\right) = y_t - y_t^n. \quad (45)$$

and the natural employment $N_t^n \equiv Y_t^n / A_t$. Since $Y_t = Y_t^n \exp(x_t)$ and $N_t = N_t^n \exp(x_t)$, [\(44\)](#) can be written as

$$\tilde{u}_t(x_t) = \frac{(Y_t^n)^{1-v} \exp((1-v)x_t) - 1}{1-v} - \frac{\zeta_t (N_t^n)^{1+\nu} \exp((1+\nu)x_t)}{1+\nu}. \quad (46)$$

Differentiating [\(46\)](#) with respect to the log gap therefore yields

$$\tilde{u}'_t(x_t) = (Y_t^n)^{1-v} \exp((1-v)x_t) - \zeta_t (N_t^n)^{1+\nu} \exp((1+\nu)x_t), \quad (47)$$

$$\tilde{u}''_t(x_t) = (1-v)(Y_t^n)^{1-v} \exp((1-v)x_t) - (1+\nu)\zeta_t (N_t^n)^{1+\nu} \exp((1+\nu)x_t). \quad (48)$$

Multiplying the condition that defines natural-output [\(33\)](#) by Y_t^n and using $N_t^n = Y_t^n / A_t$

gives

$$(Y_t^n)^{1-v} = \zeta_t (N_t^n)^{1+\nu}. \quad (49)$$

Evaluating (47) and (48) at the efficient allocation $x_t = 0$ and using (49) gives

$$\tilde{u}'_t(0) = 0, \quad \tilde{u}''_t(0) = -(v + \nu)(Y_t^n)^{1-v} < 0. \quad (50)$$

Thus, a second-order Taylor expansion around $x_t = 0$ yields

$$\tilde{u}_t(x_t) = \tilde{u}_t(0) - \frac{v + \nu}{2} (Y_t^n)^{1-v} x_t^2 + \mathcal{O}(|x_t|^3). \quad (51)$$

Note that (51) is an approximation around the time-varying efficient allocation. To express it as a constant-coefficient approximation around the steady state, use that

$$(Y_t^n)^{1-v} = \bar{Y}^{1-v} \exp((1-v)y_t^n) = \bar{Y}^{1-v} [1 + (1-v)y_t^n + \mathcal{O}(\|y_t^n\|^2)]. \quad (52)$$

Since x_t^2 is already second order, the product $y_t^n x_t^2$ is third order. Hence, to second order, Y_t^n in (51) can be replaced by $\bar{Y}$ and we get

$$\tilde{u}_t(x_t) = \tilde{u}_t(0) - \frac{v + \nu}{2} (\bar{Y})^{1-v} x_t^2 + \mathcal{O}(\|y_t^n\|^3). \quad (53)$$

Having a proper expression for the instantaneous utility, it remains to address the discount factors in the household's overall intertemporal utility (28). Let

$$D_t \equiv \prod_{j=0}^{t-1} \delta_j \text{ for } t \geq 1, \quad D_0 \equiv 1. \quad (54)$$

Since $\delta_j = \exp(-\rho_j)$, $\beta = \exp(-\bar{\rho})$, and $\tilde{\rho}_j = \rho_j - \bar{\rho}$, we have

$$D_t = \beta^t \exp\left(-\sum_{j=0}^{t-1} \tilde{\rho}_j\right) = \beta^t \left[1 - \sum_{j=0}^{t-1} \tilde{\rho}_j + \mathcal{O}(\|y_t^n\|^2)\right]. \quad (55)$$

In the household intertemporal utility function, the first-order variation in D_t multiplies the second-order term x_t^2 in (53) so it is third order. Hence, to second-order we can simply replace D_t by β^t . Thus, combining (51)–(55), expected welfare under the central bank's beliefs can be written as

$$\mathcal{W}_0^C = \mathcal{W}_0^{n,C} - \frac{(v + \nu)\bar{Y}^{1-v}}{2} \mathbb{E}_0^C \sum_{t=0}^{\infty} \beta^t (y_t - y_t^n)^2 + \mathcal{O}(\|y_t^n\|^3), \quad (56)$$

where $\mathcal{W}_0^{n,C}$ is expected welfare in the equivalent flexible-price economy. Crucially, it does not depend on monetary policy.

Consequently, dropping the policy-independent term and the positive multiplicative constant in (56), to second order the central bank’s problem of maximizing its expectation of the household’s utility can be written as

$$\min_{\{r_t\}} \mathbb{E}_0^C \sum_{t=0}^{\infty} \beta^t (y_t - y_t^n)^2, \quad (57)$$

which is exactly (2).

Quantitative model. An analogous —yet more involved— derivation can be done for the quantitative model in Section 5, since the private-sector block differs from the baseline in two respects: (i) the Calvo parameter θ is left general instead of taken to the fully rigid-price limit; and (ii) the model includes a cost-push shock to the Phillips Curve (which can be microfounded in the standard way, i.e., shocks to the elasticity of substitution between varieties and thus to the optimal markup of intermediate goods firms, see, e.g., Galí (2015, Chapter 5)). Here, y_t^n still corresponds to the efficient level of output, determined by productivity and preferences over labor. The subsidy offsets the steady-state monopoly distortion, but not its stochastic fluctuations. These fluctuations generate the cost-push term u_t , while the output gap entering the central bank’s objective remains $x_t = y_t - y_t^n$.

The central bank’s loss function in the quantitative model equals the standard New Keynesian stabilization objective with an additional penalty on deviations of the policy rate from its steady-state level. Such interest-rate stabilization terms are commonly used to prevent optimal policy from implying counterfactually volatile interest rates and can be microfounded using money in utility (see Woodford (2003, Chapter 6) for a detailed discussion).

An additional potential concern is that in our model agents’ expectations are not the full information rational expectations ones, which may matter for the derivation of the Phillips Curve (see, e.g., Preston (2005) for a detailed discussion). In our environment, however, all firms share the same information, and their beliefs are model-consistent conditional expectations with respect to nested information sets. Consequently, the law of iterated expectations holds and the standard recursive derivation of the New Keynesian Phillips curve remains valid.³⁶

³⁶This applies to the exact model in which agents keep track of the full histories. However, to solve it we use a finite-dimensional state vector that truncates the histories (see Appendix G.1) and the information sets generated by the truncated state vectors are not nested in general. We assess sensitivity to this approximation by increasing the retained history from $H = 8$ to $H = 12$ at the estimated parameters and our main results

B Information Structure

This appendix gives a self-contained construction of the belief structure corresponding to the first two periods of the baseline model. It has three purposes that mirror the following subsections. First, it specifies the basic primitives: measurable space, each player's (subjective) probability measures and information at different stages. Second, it shows which of the main text's reduced-form assumptions follow from a particular case of the familiar Gaussian prior plus signals structure. Third, it states the cross-belief restrictions that a proper common prior foundation would add and thus what elements of the baseline model are inconsistent with it while showing the information structure in the quantitative model is grounded on a common prior.

B.1 General structure

Let

$$\Omega = \mathbb{R}^8, \quad \mathcal{F} = \mathcal{B}(\Omega),$$

where

$$\omega = (y_0^n, \rho_0, \hat{y}_0^n, \hat{\rho}_0, y_1^n, \rho_1, x_{n1}, x_{\rho 1}) \in \Omega.$$

Thus $(\Omega, \mathcal{F})$ is a standard Borel space that supports all random variables relevant for the model.

There are two players, the central bank (C) and the market (M). Each player is endowed with a subjective probability measure P^i over $(\Omega, \mathcal{F})$.

For convenience, partition the coordinates of ω into the $t = 0$ fundamental vector, the central bank's $t = 0$ assessment vector, the $t = 1$ fundamentals, and the $t = 1$ private signals:

$$X_0 \equiv (y_0^n, \rho_0), \quad A_0 \equiv (\hat{y}_0^n, \hat{\rho}_0), \quad X_1 \equiv (y_1^n, \rho_1), \quad Z_1 \equiv (x_{n1}, x_{\rho 1}).$$

Thus, equivalently,

$$\omega = (X_0, A_0, X_1, Z_1).$$

The primitive uncertainty therefore concerns the contemporaneous fundamentals, the central bank's $t = 0$ assessments, the future fundamentals, and the private $t = 1$ signals.

The information available to the central bank and the market at the relevant points in are, for the most part, unchanged (see [Appendix H.1](#)). By itself, this does not guarantee that the economy is close to the one with full histories, especially given that both the $H = 8$ and $H = 12$ approximations disregard the additional terms that should appear in the Phillips curve given the information sets are not nested. However, the fact that our results are robust to increasing H suggests that the additional terms are small.

time is represented by the following σ -fields:

$$\mathcal{I}_0^C = \sigma(\hat{y}_0^n, \hat{\rho}_0), \quad \mathcal{I}_{0-}^M = \sigma(y_0^n, \rho_0),$$

and

$$\mathcal{I}_1^C = \sigma(\hat{y}_0^n, \hat{\rho}_0, x_{n1}, x_{\rho 1}).$$

The σ -fields and the subjective laws of each player are common knowledge. Note that $\mathcal{I}_1^C$ contains no endogenous variable: the central bank does not observe the $t = 0$ output realization before choosing r_1 . This is substantive, since y_0 is informative about the central bank's $t = 0$ assessment errors. Thus, relaxing this assumption would have two implications. First, in the simplest case where $\lambda_n = 0$, y_0 fully reveals ρ_0 , so even if $\varphi_\rho > 0$, the market does not use the demand-assessment $\hat{\rho}_0$ to learn about the demand component of future policy (and thus, future output): it knows the central bank's assessment at the time r_1 is chosen will be $\bar{\rho} + \varphi_\rho(\rho_0 - \bar{\rho})$. This would imply that all the belief-revision channel goes through natural output, reducing the richness of the analysis (since the main mechanism reduces to the case where $\varphi_\rho = 0$).

Second, and much more importantly, if $\lambda_n > 0$, y_0 becomes an imperfect signal for the central bank that contains both information about y_0^n (via the market's ex-ante rate forecast and thus the size of the policy surprise) and ρ_0 . Crucially, the share of informativeness about each component *depends on the policy rule itself* (the composition b , in particular). This not only makes the problem much less tractable, but also introduces a new channel through which the central bank may want to distort its policy rule: it may want to make y_0 more informative about ρ_0 at the expense of informativeness about y_0^n (or vice versa). This is a distinct mechanism from the one studied in this paper, and is of doubtful empirical relevance. Thus, we assume that the central bank does not observe y_0 before choosing r_1 , and thus cannot use it to learn about its own assessment errors. Similarly, in the quantitative model it is assumed that endogenous outcomes (π_t, y_t) are never observed. Note that this does not have any implication for the extent of information about each fundamental the central bank possesses, since that can be freely shaped by changing the noise of the signals. Thus, its sole implication is precisely what it is intended to be: to avoid policy distortions motivated by changing the informativeness of endogenous outcomes.

Hence, the central bank's $t = 0$ private information is summarized by its assessment vector A_0 . At $t = 1$ the central bank additionally observes the private signal vector Z_1 . The market observes the contemporaneous fundamentals X_0 .

For any measurable $t = 0$ policy rule

$$g_0 : \mathbb{R}^2 \rightarrow \mathbb{R},$$

define the induced policy decision

$$r_0(\omega) = g_0(\hat{y}_0^n, \hat{\rho}_0).$$

Given such a rule, the market's post-decision σ -field is

$$\mathcal{I}_{0+}^{M,g_0} := \sigma(y_0^n, \rho_0, r_0),$$

and its $t = 1$ σ -field is

$$\mathcal{I}_1^{M,g_0} := \sigma(y_0^n, \rho_0, r_0, y_1^n, \rho_1).$$

For any measurable policy rule g_0 , pre- and post-decision beliefs are represented by the regular conditional probabilities³⁷

$$\pi_0^C(\omega) := P^C(\cdot \mid \mathcal{I}_0^C)(\omega), \quad \pi_{0-}^M(\omega) := P^M(\cdot \mid \mathcal{I}_{0-}^M)(\omega),$$

$$\pi_{0+}^{M,g_0}(\omega) := P^M(\cdot \mid \mathcal{I}_{0+}^{M,g_0})(\omega),$$

and $t = 1$ beliefs are represented by

$$\pi_1^C(\omega) := P^C(\cdot \mid \mathcal{I}_1^C)(\omega), \quad \pi_1^{M,g_0}(\omega) := P^M(\cdot \mid \mathcal{I}_1^{M,g_0})(\omega).$$

Equivalently, there exists a stochastic kernel K_{0+}^{M,g_0} on $(\Omega, \mathcal{F})$ given $\mathbb{R}^3$ such that

$$\pi_{0+}^{M,g_0}(\omega) = K_{0+}^{M,g_0}(y_0^n(\omega), \rho_0(\omega), r_0(\omega)) \quad P^M\text{-a.s.}$$

Note that since

$$r_0 = g_0(\hat{y}_0^n, \hat{\rho}_0),$$

the market's posterior after the policy decision depends on the policy rule g_0 : the observed action is an endogenous signal of the central bank's private assessment.

We restrict the central bank's subjective law P^C so that its $t = 0$ assessment vector A_0 coincides with the central bank's posterior mean of the contemporaneous fundamental vector

³⁷These are well-defined since $(\Omega, \mathcal{F})$ is a standard Borel space.

X_0 given its private information, that is,

$$A_0 = \mathbb{E}^C[X_0 \mid \mathcal{I}_0^C] \quad P^C\text{-a.s.},$$

or component-wise,

$$\hat{y}_0^n = \mathbb{E}^C[y_0^n \mid \mathcal{I}_0^C], \quad \hat{\rho}_0 = \mathbb{E}^C[\rho_0 \mid \mathcal{I}_0^C] \quad P^C\text{-a.s.}$$

Equivalently, defining the $t = 0$ estimation-error vector

$$e_0 \equiv X_0 - A_0,$$

we have

$$\mathbb{E}^C[e_0 \mid \mathcal{I}_0^C] = 0 \quad P^C\text{-a.s.}$$

Thus, under P^C , $(\hat{y}_0^n, \hat{\rho}_0)$ are posterior means conditional on the central bank's $t = 0$ information.³⁸

The primitive information structure is therefore given by

$$(\Omega, \mathcal{F}), \quad (\mathcal{I}_0^C, \mathcal{I}_{0-}^M, \mathcal{I}_1^C), \quad (P^C, P^M).$$

For any measurable policy rule g_0 , this structure induces the observable action r_0 , the information fields $\mathcal{I}_{0+}^{M,g_0}$ and $\mathcal{I}_1^{M,g_0}$, and the associated conditional belief kernels. These objects jointly determine the players' hierarchy of conditional beliefs.

B.2 Gaussian beliefs and further assumptions in the main text

The reduced-form model in the main text specializes the structure above to a Gaussian distribution for the fundamental, assessment, and signal blocks. In this subsection we provide the necessary background and assumptions to support the Gaussian specialization. Throughout, expectations taken under the central bank's and the market's subjective laws are denoted, respectively, by³⁹

$$\mathbb{E}^C[\cdot] \equiv \mathbb{E}^{P^C}[\cdot], \quad \mathbb{E}^M[\cdot] \equiv \mathbb{E}^{P^M}[\cdot],$$

³⁸Since the market's $t = 0$ information is given by the observed fundamentals X_0 , it would be of no use to define separate market $t = 0$ assessments.

³⁹Unlike in common-prior information structures, expectations of two agents may differ not only because of their different information sets, but also because of their different subjective probability measures. Thus, we retain the latter as a superscript even when specifying the conditioning information set.

and analogously for conditional variances and covariances.

Under P^C , natural output evolves as

$$y_0^n \sim \mathcal{N}(0, \sigma_{y_0}^2),$$

while the discount rate evolves as

$$\rho_0 \sim \mathcal{N}(\bar{\rho}, \sigma_{\rho_0}^2).$$

At date 0, define the central bank's posterior errors by

$$\eta_0 \equiv y_0^n - \hat{y}_0^n, \quad \varepsilon_0 \equiv \rho_0 - \hat{\rho}_0, \quad (58)$$

where the natural-output block $(\hat{y}_0^n, \eta_0)$ and the discount-rate block $(\hat{\rho}_0, \varepsilon_0)$ are jointly Gaussian under P^C , independent across blocks, with

$$\mathbb{E}^C[\eta_0] = \mathbb{E}^C[\varepsilon_0] = 0,$$

and

$$\text{Var}^C(\hat{y}_0^n) = \sigma_{n,C}^2, \quad \text{Var}^C(\hat{\rho}_0) = \sigma_{\rho,C}^2, \quad \text{Var}^C(\eta_0) = \sigma_{\eta,C}^2, \quad \text{Var}^C(\varepsilon_0) = \sigma_{\varepsilon,C}^2.$$

Let

$$e_0 \equiv (\eta_0, \varepsilon_0).$$

Since $\mathcal{I}_0^C = \sigma(A_0)$ and $A_0 = \mathbb{E}^C[X_0 \mid \mathcal{I}_0^C]$,

$$\mathbb{E}^C[e_0 \mid \mathcal{I}_0^C] = 0.$$

It follows that

$$\text{Cov}^C(\hat{y}_0^n, \eta_0) = 0, \quad \text{Cov}^C(\hat{\rho}_0, \varepsilon_0) = 0.$$

Under joint Gaussianity, these orthogonality restrictions are also sufficient for the conditional mean restriction above. Together with (58), these orthogonality conditions imply the variance decompositions

$$\sigma_{y_0}^2 = \sigma_{n,C}^2 + \sigma_{\eta,C}^2, \quad \sigma_{\rho_0}^2 = \sigma_{\rho,C}^2 + \sigma_{\varepsilon,C}^2.$$

Thus, the unconditional fundamental variances are implied by the assessment and posterior-error variances instead of being additional free parameters. Equivalently, the central bank's

$t = 0$ conditional law of contemporaneous fundamentals given its private assessments is

$$y_0^n \mid \mathcal{I}_0^C \sim \mathcal{N}(\hat{y}_0^n, \sigma_{\eta,C}^2), \quad \rho_0 \mid \mathcal{I}_0^C \sim \mathcal{N}(\hat{\rho}_0, \sigma_{\varepsilon,C}^2),$$

which gives (5) and (6) in the main text.

The $t = 1$ variables satisfy

$$u_{n1} := y_1^n - \varphi_n y_0^n, \quad u_{\rho 1} := \rho_1 - \bar{\rho} - \varphi_\rho(\rho_0 - \bar{\rho}),$$

so $u_{n1} = \varepsilon_n$ and $u_{\rho 1} = \varepsilon_\rho$ in (3)–(4), and

$$\nu_{n1} := x_{n1} - y_1^n, \quad \nu_{\rho 1} := x_{\rho 1} - \rho_1,$$

where

$$u_{n1} \sim \mathcal{N}(0, \sigma_{un1}^2), \quad u_{\rho 1} \sim \mathcal{N}(0, \sigma_{u\rho 1}^2),$$

$$\nu_{n1} \sim \mathcal{N}(0, \sigma_{\nu n1}^2), \quad \nu_{\rho 1} \sim \mathcal{N}(0, \sigma_{\nu\rho 1}^2),$$

independently of everything else. Thus, fundamentals (y_0^n, ρ_0) evolve as Gaussian AR(1) processes with persistences φ_n and φ_ρ , respectively, and the central bank observes signals x_{n1} and $x_{\rho 1}$ at $t = 1$ that equal the fundamentals plus white Gaussian noise. Hence,

$$y_1^n \mid \mathcal{I}_0^C \sim \mathcal{N}(\varphi_n \hat{y}_0^n, \varphi_n^2 \sigma_{\eta,C}^2 + \sigma_{un1}^2),$$

$$\rho_1 \mid \mathcal{I}_0^C \sim \mathcal{N}(\bar{\rho} + \varphi_\rho(\hat{\rho}_0 - \bar{\rho}), \varphi_\rho^2 \sigma_{\varepsilon,C}^2 + \sigma_{u\rho 1}^2).$$

Consequently, the $t = 1$ posterior means are

$$\hat{y}_1^n := \mathbb{E}^C[y_1^n \mid \mathcal{I}_1^C] = (1 - \kappa_{n1})\varphi_n \hat{y}_0^n + \kappa_{n1}x_{n1},$$

$$\hat{\rho}_1 := \mathbb{E}^C[\rho_1 \mid \mathcal{I}_1^C] = \bar{\rho} + (1 - \kappa_{\rho 1})\varphi_\rho(\hat{\rho}_0 - \bar{\rho}) + \kappa_{\rho 1}(x_{\rho 1} - \bar{\rho}),$$

where

$$\kappa_{n1} \equiv \frac{\varphi_n^2 \sigma_{\eta,C}^2 + \sigma_{un1}^2}{\varphi_n^2 \sigma_{\eta,C}^2 + \sigma_{un1}^2 + \sigma_{\nu n1}^2},$$

$$\kappa_{\rho 1} \equiv \frac{\varphi_\rho^2 \sigma_{\varepsilon,C}^2 + \sigma_{u\rho 1}^2}{\varphi_\rho^2 \sigma_{\varepsilon,C}^2 + \sigma_{u\rho 1}^2 + \sigma_{\nu\rho 1}^2}$$

are the standard Kalman gains.

Under P^M , the conditional law of the central bank's $t = 0$ assessments given the market

observables (y_0^n, ρ_0) is instead assumed to follow⁴⁰

$$\begin{pmatrix} \hat{y}_0^n \\ \hat{\rho}_0 \end{pmatrix} \Big| \mathcal{I}_{0-}^M \sim \mathcal{N} \left(\begin{pmatrix} \lambda_n y_0^n \\ \bar{\rho} + \lambda_\rho (\rho_0 - \bar{\rho}) \end{pmatrix}, \begin{pmatrix} V_{n0} & 0 \\ 0 & sV_{\rho 0} \end{pmatrix} \right).$$

The parameters λ_n and λ_ρ govern how the market centers the central bank's $t = 0$ assessments on the contemporaneous fundamentals. If $\lambda_z = 1$ for some $z \in \{n, \rho\}$, the market treats the corresponding assessment as correctly centered on average. If $\lambda_z = 0$, it treats that assessment as orthogonal to the contemporaneous fundamental. These are polar cases; a standard proper common-prior signal structure generally produces coefficients strictly between zero and one, as discussed below.

The market and the central bank may not agree on the distribution of the central bank's $t = 0$ assessments conditional on the contemporaneous fundamentals. We assume, however, that under P^M , the $t = 1$ fundamental innovations $u_{n1}, u_{\rho 1}$ and signal errors $\nu_{n1}, \nu_{\rho 1}$ have the Gaussian distributions specified above, are mutually independent, and are independent of (X_0, A_0) . The market also knows that the central bank maps (A_0, Z_1) into its $t = 1$ assessments according to the filtering rules above. Hence,

$$\mathbb{E}^M[\hat{y}_0^n | \mathcal{I}_{0-}^M] = \lambda_n y_0^n,$$

$$\mathbb{E}^M[\hat{\rho}_0 | \mathcal{I}_{0-}^M] = \bar{\rho} + \lambda_\rho (\rho_0 - \bar{\rho}),$$

$$\mathbb{E}^M[\hat{y}_1^n | \mathcal{I}_{0-}^M] = ((1 - \kappa_{n1})\lambda_n + \kappa_{n1}) \varphi_n y_0^n,$$

$$\mathbb{E}^M[\hat{\rho}_1 | \mathcal{I}_{0-}^M] = \bar{\rho} + ((1 - \kappa_{\rho 1})\lambda_\rho + \kappa_{\rho 1}) \varphi_\rho (\rho_0 - \bar{\rho}).$$

Moreover, since the market's $t = 0$ conditional law is Gaussian and the $t = 1$ updating rule is linear, under P^M ,

$$\begin{pmatrix} \hat{y}_0^n \\ \hat{y}_1^n \end{pmatrix} \Big| \mathcal{I}_{0-}^M \sim \mathcal{N} \left(\begin{pmatrix} \mathbb{E}^M[\hat{y}_0^n | \mathcal{I}_{0-}^M] \\ \mathbb{E}^M[\hat{y}_1^n | \mathcal{I}_{0-}^M] \end{pmatrix}, \begin{pmatrix} V_{n0} & c_n \\ c_n & V_{n1} \end{pmatrix} \right),$$

where

$$c_n \equiv \text{Cov}^M(\hat{y}_0^n, \hat{y}_1^n | \mathcal{I}_{0-}^M) = (1 - \kappa_{n1})\varphi_n V_{n0},$$

and

$$V_{n1} \equiv \text{Var}^M(\hat{y}_1^n | \mathcal{I}_{0-}^M) = (1 - \kappa_{n1})^2 \varphi_n^2 V_{n0} + \kappa_{n1}^2 (\sigma_{un1}^2 + \sigma_{\nu n1}^2).$$

⁴⁰Since in the model the market observes fundamentals before any decision takes place, specifying their prior distribution under P^M is not necessary. However, for completeness, one can assume that it is the same as the central bank's.

Likewise,

$$\begin{pmatrix} \hat{\rho}_0 \\ \hat{\rho}_1 \end{pmatrix} \Big| \mathcal{I}_{0-}^M \sim \mathcal{N} \left(\begin{pmatrix} \mathbb{E}^M[\hat{\rho}_0 | \mathcal{I}_{0-}^M] \\ \mathbb{E}^M[\hat{\rho}_1 | \mathcal{I}_{0-}^M] \end{pmatrix}, \begin{pmatrix} sV_{\rho 0} & c_\rho(s) \\ c_\rho(s) & V_{\rho 1}(s) \end{pmatrix} \right),$$

where

$$c_\rho(s) \equiv \text{Cov}^M(\hat{\rho}_0, \hat{\rho}_1 | \mathcal{I}_{0-}^M) = (1 - \kappa_{\rho 1})\varphi_\rho s V_{\rho 0},$$

and

$$V_{\rho 1}(s) \equiv \text{Var}^M(\hat{\rho}_1 | \mathcal{I}_{0-}^M) = (1 - \kappa_{\rho 1})^2 \varphi_\rho^2 s V_{\rho 0} + \kappa_{\rho 1}^2 (\sigma_{u\rho 1}^2 + \sigma_{\nu\rho 1}^2).$$

The parameter s scales the market's perceived conditional variance of the central bank's $t = 0$ demand assessment. Given the normalization $V_{n0} = V_{\rho 0}$ imposed below, $s = 1$ is the symmetric-variance benchmark.

Further assumptions in the main text. In the main text, we take the above Gaussian information structure and further specialize it as follows.

- Normalize $V_{n0} = V_{\rho 0} = 1$.
- Assume the central bank receives no additional signal at $t = 1$. Equivalently, remove Z_1 from $\mathcal{I}_1^C$ and set $\kappa_{n1} = \kappa_{\rho 1} = 0$ in the updating formulas above.
- Set $\lambda_\rho = 1$.

The first assumption implies that the variance ratio of the market's beliefs about the central bank's demand and natural-output assessments reduces to s :

$$\frac{\text{Var}^M(\hat{\rho}_0 | \mathcal{I}_{0-}^M)}{\text{Var}^M(\hat{y}_0^n | \mathcal{I}_{0-}^M)} = \frac{V_{\rho 0}}{V_{n0}} s = s.$$

The second assumption entails that the central bank receives no new information between its $t = 0$ and $t = 1$ policy decisions. It therefore gives

$$\hat{y}_1^n = \varphi_n \hat{y}_0^n, \quad \hat{\rho}_1 = \bar{\rho} + \varphi_\rho (\hat{\rho}_0 - \bar{\rho}),$$

as assumed in (7). In particular, the $t = 1$ variances and covariances are implied by each fundamental's corresponding persistence coefficient,

$$\text{Var}^M(\hat{y}_1^n | \mathcal{I}_{0-}^M) = \varphi_n^2, \quad \text{Cov}^M(\hat{y}_0^n, \hat{y}_1^n | \mathcal{I}_{0-}^M) = \varphi_n,$$

$$\text{Var}^M(\hat{\rho}_1 | \mathcal{I}_{0-}^M) = s\varphi_\rho^2, \quad \text{Cov}^M(\hat{\rho}_0, \hat{\rho}_1 | \mathcal{I}_{0-}^M) = s\varphi_\rho.$$

The restriction is imposed for expositional simplicity. More generally, new $t = 1$ information would make the persistence of the central bank’s assessments weakly smaller than the persistence of the corresponding fundamentals, because the bank would place less weight on its $t = 0$ assessment. The quantitative analysis in [Section 5](#) relaxes this restriction since the central bank observes new signals each period.

The third assumption is a reduced-form belief specialization for analytical transparency, and its purpose is to preserve a property of the environment that makes the analysis cleaner.

Concretely, setting $\lambda_\rho = 1$ means that the market treats the central bank’s initial demand assessment as correctly centered:

$$\mathbb{E}^M[\hat{\rho}_0 \mid \mathcal{I}_{0-}^M] = \rho_0.$$

It also implies

$$\mathbb{E}^M[\hat{\rho}_1 \mid \mathcal{I}_{0-}^M] = \mathbb{E}^M[\rho_1 \mid \mathcal{I}_{0-}^M].$$

This is arguably a natural feature to assume, although exact centering is not generic under a standard finite-precision prior-and-signal structure. This centering assumption is chosen because it delivers the clean discretionary coefficient of one on the central bank’s demand assessment. The reason is that, when $\lambda_\rho = 1$, the demand component of the policy surprise $r_0 - \mathbb{E}_0^M[r_0]$ has conditional mean zero under the central bank’s beliefs. Its stochastic component is an estimation error that enters output additively, so certainty equivalence applies. The discretionary central bank therefore chooses a unit coefficient on $\hat{\rho}_0$: since the discount rate and the policy rate enter output with the same coefficient γ , the central bank wants to offset its demand assessment one for one. A central bank with commitment, by contrast, remains free to choose $a \neq 1$ because it internalizes how the policy rule affects the pass-through of assessment errors.

B.3 Common-prior restrictions

The information structure above features heterogeneous priors: P^C and P^M are different subjective probability measures. While this does not compromise either player’s Bayesian rationality, it does entail additional flexibility over standard common prior models. As discussed by [Morris \(1995\)](#), the common-prior assumption is a substantive restriction on the relation between individuals’ beliefs, and there are good reasons to doubt its validity in many economically relevant real-world contexts. Nevertheless, it remains a standard assumption in economic theory, and even more so in macroeconomics. Thus, this section details the implications a common prior foundation would have in the context of the model above.

Market's centering and uncertainty. Since the model assumes independence between blocks, the implications of a common-prior foundation are unrelated to the links between them. Thus, consider an arbitrary one in isolation. Let P be a proper shared probability measure and $x \in \mathbb{R}$ be the square-integrable and nondegenerate fundamental with prior mean zero (we assume this for simplicity, an arbitrary common knowledge prior mean does not change the argument materially). Finally, let $\mathcal{I}^C$ denote the central bank's information. By construction, the central bank's assessment of x satisfies

$$\hat{x} = \mathbb{E}_P[x \mid \mathcal{I}^C].$$

The estimate error $x - \hat{x}$ is orthogonal to every $\mathcal{I}^C$ -measurable random variable, so in particular

$$0 = \mathbb{E}_P[(x - \hat{x})\hat{x}]$$

which implies $\mathbb{E}_P[x\hat{x}] = \mathbb{E}_P[\hat{x}^2]$. Moreover, the law of iterated expectations gives

$$\mathbb{E}_P[\hat{x}] = \mathbb{E}_P[\mathbb{E}_P[x \mid \mathcal{I}^C]] = \mathbb{E}_P[x] = 0.$$

Putting the above equalities together yields

$$\text{Cov}_P(x, \hat{x}) = \text{Var}_P(\hat{x}). \tag{59}$$

Now, as in our model, suppose the market observes x and

$$\mathbb{E}_P[\hat{x} \mid x] = \lambda x \tag{60}$$

for some $\lambda \in \mathbb{R}$. Note that the fact that this holds under the common prior measure P is crucial for the argument below. In particular, it implies

$$\text{Cov}_P(x, \hat{x}) = \mathbb{E}_P[x\hat{x}] \tag{61}$$

$$= \mathbb{E}_P[x \mathbb{E}_P[\hat{x} \mid x]] \tag{62}$$

$$= \lambda \mathbb{E}_P[x^2] = \lambda \text{Var}_P(x). \tag{63}$$

Combining this with (59) gives

$$\text{Var}_P(\hat{x}) = \lambda \text{Var}_P(x), \tag{64}$$

which implies $\lambda \geq 0$. Consequently, the market's centering parameter λ , prior uncertainty about fundamentals and prior uncertainty about the central bank's assessment are linked. Furthermore, the law of total variance and (64) imply

$$\begin{aligned}\lambda \operatorname{Var}_P(x) &= \operatorname{Var}_P(\hat{x}) \\ &= \operatorname{Var}_P(\mathbb{E}_P[\hat{x} | x]) + \mathbb{E}_P[\operatorname{Var}_P(\hat{x} | x)] \\ &= \lambda^2 \operatorname{Var}_P(x) + \mathbb{E}_P[\operatorname{Var}_P(\hat{x} | x)],\end{aligned}$$

and hence

$$\mathbb{E}_P[\operatorname{Var}_P(\hat{x} | x)] = \lambda(1 - \lambda) \operatorname{Var}_P(x), \quad (65)$$

which further constrains $\lambda \in [0, 1]$. Whenever the conditional variance is constant in x , like in the Gaussian case, [equation \(65\)](#) pins it down directly as a function of λ and the prior variance. Crucially, without a common prior, (59) would hold under P^C , while (61) would hold under P^M so there would be extra degrees of freedom in the relation above.

The above implies that the endpoint values for λ have pretty stark consequences. If $\lambda = 0$, [equation \(64\)](#) implies that $\hat{x}$ is constant P -almost surely. If $\lambda = 1$, then $\operatorname{Var}_P(\hat{x}) = \operatorname{Var}_P(x)$ and thus, by orthogonality

$$\operatorname{Var}_P(x) = \operatorname{Var}_P(\hat{x}) + \operatorname{Var}_P(x - \hat{x}) \implies x = \hat{x}$$

P -almost surely. Thus, neither centering endpoint is compatible with the market having uncertainty about the central bank's assessment under a common prior environment.

While the argument above is general, below we develop it in its special $t = 0$ Gaussian prior plus signal structure version. We do so because it makes the intuition for the above constraint transparent and further illustrates that *every* interior λ can be obtained that way.

Gaussian specialization. Let the common-prior block be such that

$$x \sim \mathcal{N}(0, V), \quad z = x + \nu, \quad \nu \sim \mathcal{N}(0, R), \quad x \perp \nu,$$

where x is the fundamental, z the signal observed by the central bank, ν its noise, and $V, R > 0$. Standard Bayesian updating implies

$$\hat{x} = \mathbb{E}[x | z] = \lambda z, \quad \lambda \equiv \frac{V}{V + R} \in (0, 1).$$

If the market observes the fundamental x , its beliefs about the central bank’s assessment satisfy

$$\mathbb{E}[\hat{x} \mid x] = \lambda x, \quad \text{Var}(\hat{x} \mid x) = \lambda^2 R = \lambda(1 - \lambda)V.$$

Thus, every interior centering $\lambda \in (0, 1)$ is consistent with a common prior in the Gaussian prior and signal setting. Further, the intuition for why the endpoints are not is transparent: as $\lambda \rightarrow 0$ or $\lambda \rightarrow 1$, the market’s uncertainty regarding the central bank’s assessment approaches zero. If the central bank puts zero weight on its private information, the common prior mean is common knowledge and the market knows its assessment. Analogously, if the central bank puts full weight on its signal, the signal has no noise and equals the fundamental. Since the market knows the fundamental, it also knows the central bank’s assessment.

The baseline model, on the other hand, leaves the centering and conditional variance as free parameters and, further, assumes the ρ -centering equals 1, which is not consistent with non-zero uncertainty about the central bank’s assessment under a common prior. The baseline model is therefore not compatible with a common prior as is, but it would be so if the ρ -centering were strictly less than 1, $\lambda_n > 0$, and the above relationships between the centering and conditional variance were imposed. In fact, the quantitative model in [Section 5](#) is fully consistent with a common prior since the centering parameters are the ones implied by the rest of the information structure (in particular, they are consistent with the steady-state Kalman filter, see [Appendix G](#) for details). Concretely, a single Gaussian distribution generates fundamentals and the central bank’s signal errors, and its parameters and the signal structure are common knowledge.

An important implication of the lack of a common prior foundation for the model is that there is no natural ex-ante welfare measure, since there is more than one possible ex-ante distribution under which to compute expected values. We discuss this issue below.

Implications for “welfare” statements. In [Section 3](#), the central bank’s objective and all ex-ante loss comparisons are evaluated under the central bank’s ex-ante measure P^C . Absent a common prior, it is not straightforward to take these as actual welfare statements, even within the stylized economy under consideration and taking at face value the central bank’s loss as a proper approximation of actual utility. Thus, the claims there are stated in terms of central bank loss only. Going beyond that would require taking a stance on the ex-ante distribution, which is beyond the scope of the paper.

Importantly, this is not an issue in the richer model of [Section 5](#), the source of all quantitative welfare claims in the paper. This is because, as discussed above, that model does admit a common prior foundation and thus *there is a unique ex-ante welfare measure*,

the one used in the results those claims depend on.⁴¹

C Game-Theoretic Representation and Solution Concepts

This appendix represents the $t = 0$ interaction between the central bank and the market in [Section 2](#) as an extensive-form game. It restates the discretion and commitment solution concepts of [Section 3](#) in terms of strategies and unilateral deviations, taking the sequentially optimal continuation from $t = 1$ onward as given.

The definitions use only the primitives of [Appendix B.1](#)—the measurable space $(\Omega, \mathcal{F})$, the σ -fields representing each player’s information, and the subjective laws (P^C, P^M) . The discussion of linear rules and the pointwise version of the market’s posterior also uses the Gaussian specialization of [Appendix B.2](#). Since the parameter s is a known constant, it is suppressed from the notation.

C.1 The game

There are two players, the central bank (C) and the market (M).⁴² Recall the partition $\omega = (X_0, A_0, X_1, Z_1)$, where $X_0 = (y_0^n, \rho_0)$ are the $t = 0$ fundamentals, $A_0 = (\hat{y}_0^n, \hat{\rho}_0)$ the central bank’s $t = 0$ assessments, X_1 the $t = 1$ fundamentals and Z_1 the central bank’s $t = 1$ signals. The extensive form is as follows.⁴³

1. A state ω is realized. The central bank observes the realization of A_0 and the market observes that of X_0 .
2. The central bank chooses $r_0 \in \mathbb{R}$.
3. The market observes the realization of r_0 and chooses a forecast $f_0 \in \mathbb{R}$ of y_1 . Output at $t = 0$ is then determined by the IS curve [\(1\)](#), $y_0 = f_0 - \gamma(r_0 - \rho_0)$.

⁴¹Of course, that measure being an “appropriate” proxy for actual welfare in the economy relies on the extent to which the model is a legitimate approximation of a more standard model. On that, see [Appendix A](#), particularly the discussion of the quantitative model at the end.

⁴²For the purposes of this section, we take output to be “mechanically” determined by the IS [equation \(1\)](#) (which, as [Appendix A](#) shows, is a linear approximation of the household’s utility-maximizing behavior together with market clearing), so the market *actions* are its forecasts of next-period output. Given this and the fact that the market does not have incentives to alter its forecast even if it could affect aggregate expectations as argued in [Remark 2](#) below, we can safely think of the market as being one player and not an aggregation of a continuum of agents.

⁴³The main focus of our analysis is the $t = 0$ game since that is where the policy rate plays a signaling role that makes the game-theoretic formulation useful. Thus, we omit some notation for the $t = 1$ game and fully omit the $t \geq 2$ stages.

4. At $t = 1$, the central bank observes Z_1 , with information $\mathcal{I}_1^C$, and chooses r_1 . Then, the market observes X_1 and forecasts y_2 and output at $t = 1$ is determined by (1).
5. For $t \geq 2$ it is common knowledge that $(y_t^n, \rho_t) = (0, \bar{\rho})$ and thus a complete-information version of the stage game repeats. We omit the resulting continuation strategies below since the structure is fully standard.

A strategy for the central bank is a pair of measurable maps

$$g_0 : \mathbb{R}^2 \rightarrow \mathbb{R}, \quad g_1 : \mathbb{R}^4 \rightarrow \mathbb{R},$$

where $r_0 = g_0(\hat{y}_0^n, \hat{\rho}_0)$ and $r_1 = g_1(\hat{y}_0^n, \hat{\rho}_0, x_{n1}, x_{\rho 1})$.⁴⁴ Equivalently in [Appendix B](#) terms, r_t is $\mathcal{I}_t^C$ -measurable. A strategy for the market is a measurable map

$$\mu : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}, \quad f_0 = \mu(X_0, r_0),$$

together with its (degenerate) $t = 1$ counterpart. Note that μ is defined at *every* $r_0 \in \mathbb{R}$. The central bank's ex-ante loss and the market's $t = 0$ forecast loss are

$$\mathcal{L}^C = \mathbb{E}^C \left[\sum_{t \geq 0} \beta^t (y_t - y_t^n)^2 \right], \quad \ell_0^M = (f_0 - y_1)^2,$$

respectively. The market minimizes the conditional expectation of ℓ_0^M , while the central bank minimizes $\mathcal{L}^C$.

Remark 2 (The market's forecast does not affect its target). By [Lemma 1](#), $y_1 = \hat{y}_1^n + \gamma(\rho_1 - \hat{\rho}_1)$ depends only on the central bank's $t = 1$ assessments and the contemporaneous demand fundamental (recall we assumed the central bank does not observe y_0). Hence, y_1 is unaffected by the $t = 0$ forecast f_0 and thus the market has no incentive to distort its forecast in order to influence the object it is forecasting. This implies that (i) the market's best response is its conditional expectation of y_1 and (ii) the analysis that takes the market to be one agent is identical to the one that would arise from a continuum of atomistic forecasters.⁴⁵

Write g_1^* for the $t = 1$ strategy in [Lemma 1](#). It is the unique sequentially optimal continuation and does not depend on g_0 . Since our analysis of different solution concepts

⁴⁴In the baseline specialization, there are no $t = 1$ informative signals so g_1 depends on A_0 alone, i.e., it is a function from $\mathbb{R}^2$ to $\mathbb{R}$ (see [Appendix B](#)). We keep the general form and thus the slight abuse of notation for clarity.

⁴⁵Of course, this does matter for the derivation of the IS curve so for the purposes of the full model, the standard continuum interpretation is the appropriate one.

pertains only to the $t = 0$ game, y_1 below is the outcome induced by g_1^* unless stated otherwise.

C.2 Market-consistent profiles

As argued above, the market's optimal strategy (forecast) is always its conditional expectation of y_1 . Thus, we define the solution concepts of interest, taking that as given. The following definition formalizes this.

Definition 2. *Given a policy rule g_0 , let $\mathcal{M}(g_0)$ denote the set of optimal market strategies, i.e., μ such that for almost every (X_0, r_0) under the joint law induced by g_0 and P^M ,*

$$\mu(X_0, r_0) \in \arg \min_{f \in \mathbb{R}} \mathbb{E}^M \left[(f - y_1)^2 \mid X_0, r_0 \right].$$

The set of market-consistent profiles is $\mathcal{E} \equiv \{(g_0, \mu) : \mu \in \mathcal{M}(g_0)\}$.

Because the loss is quadratic, $\mu \in \mathcal{M}(g_0)$ if and only if

$$\mu(x, r) = \mathbb{E}^M \left[y_1 \mid X_0 = x, r_0 = r \right] \quad \text{for } P^M\text{-a.e. } (x, r). \quad (66)$$

For linear rules, the corresponding formula is given by [Proposition 1](#). Note that conditional expectations are defined only almost surely. To determine the potential deviations of the central bank and thus characterize equilibria (however defined, see below), a pointwise version needs to be specified. For the constant rule, Bayes' rule determines the forecast at $r = \bar{\rho}$ but leaves forecasts following deviations unspecified. The following definition specifies the market's forecast at every fundamental realization and policy rate, using the Gaussian updating formula for nonconstant rules and the natural “no-update” convention for the constant rule.

Definition 3. *Let $g_0^{a,K}$ denote a rule of the form [\(12\)](#). If $(a, K) \neq (0, 0)$, define $\mu^{a,K}(x, r)$ for every $(x, r) \in \mathbb{R}^2 \times \mathbb{R}$ by the affine Gaussian updating formula in [Proposition 1](#) evaluated at $X_0 = x$. Letting $x = (y^n, \rho)$, for the constant rule $a = K = 0$, define*

$$\mu^{0,0}(x, r) := \varphi_n \lambda_n y^n \quad \text{for every } (x, r) \in \mathbb{R}^2 \times \mathbb{R},$$

i.e., the prior mean of y_1 conditional on $X_0 = x$.

For nonconstant rules, the selected affine formula agrees almost surely with the market's conditional expectation. Evaluating it at every (x, r) specifies a continuous market response also at histories where that conditional expectation is not pinned down.

The constant rule is different. Bayes' rule pins down the forecast at $r = \bar{\rho}$ but says nothing about deviations. The constant-rule response is a *convention*, albeit a natural one: after a deviation from the constant rule, the market does not revise its assessment of y_1 . This is a very simple form of equilibrium refinement and is used for ruling out a constant-rate discretionary equilibrium (e.g., in [Proposition 2](#)). In the commitment case, the central bank is bound by its rule, so only the market's on-path forecast matters.

By [Lemma 1](#) and the discussion surrounding [\(17\)](#), the central bank's policy problems are separable across periods. The optimality conditions below can therefore be written using the relevant current-period loss.

C.3 Discretion

Definition 4 (Discretionary equilibrium). *A profile (g_0, g_1, μ) is a discretionary equilibrium if*

$$(i) \ (g_0, \mu) \in \mathcal{E};$$

$$(ii) \text{ for } P^C\text{-almost every realization of } A_0,$$

$$g_0(A_0) \in \arg \min_{r \in \mathbb{R}} \mathbb{E}^C \left[(y_0 - y_0^n)^2 \mid \mathcal{I}_0^C \right]$$

when the market plays μ ; and

$$(iii) \text{ for } P^C\text{-almost every realization of } (A_0, Z_1),$$

$$g_1(A_0, Z_1) \in \arg \min_{r_1 \in \mathbb{R}} \mathbb{E}^C \left[(y_1 - y_1^n)^2 \mid \mathcal{I}_1^C \right].$$

It is a linear discretionary equilibrium if in addition g_0 takes the form [\(12\)](#) and $\mu = \mu^{a,K}$ as specified in [Definition 3](#).

Condition (i) makes the market's forecast optimal at almost every history generated by the rule. For a linear equilibrium, the selected response extends that forecast to every rate. Conditions (ii) and (iii) make the central bank's action optimal for P^C -almost every realization of its information.

Crucially, the phrase in the main text *taking the belief map as given* has a precise game-theoretic meaning. The coefficients (a, b) that appear in the belief-map slope $-\chi(b, s)/a$ are features of the market's strategy μ , not of the central bank's. Condition (ii) holds μ fixed while the central bank varies r because a deviation by one player at one information set does

not alter another player's strategy. Thus, the statement in [Section 3.3](#) that the central bank *does not internalize* the effect of its reaction on (a, b) is not a behavioral restriction. It is the ordinary content of sequential rationality, and the fixed point in (i)–(ii) is the usual one when an informed player's action is observed.

Remark 3 (On linearity). *Although we study linear fixed point(s) of (i)–(ii), the central bank's optimization is unrestricted: it can choose any measurable policy rate at each point in time. For a linear policy conjecture with $a \neq 0$, [Proposition 1](#) implies that the market's strategy μ is affine in the observed rate. Under the nonzero-total-transmission convention in [Section 3.3](#), the central bank's $t = 0$ loss is a strictly convex quadratic in the rate, and its unique pointwise best response is affine in $\mathbb{E}^C[X_0 | \mathcal{I}_0^C] = A_0$. Hence, the linear policies class is closed under the (unrestricted) best-response mapping.*

Linearity does, however, restrict the class of market conjectures over which uniqueness claims hold: we do not rule out equilibria supported by nonlinear market conjectures.

Finally, when $P^C = P^M$, the discretionary-equilibrium conditions are those of a perfect Bayesian equilibrium in the corresponding common-prior signaling game, with the central bank's assessment vector A_0 as its type. The analytical model does not impose that restriction; [Appendix B.3](#) discusses the resulting scope and contrasts it with the common-prior quantitative model.

C.4 Commitment

Commitment adds a “ $t = -1$ move” to the game. Before observing A_0 , the central bank publicly chooses and commits to an interest rate rule from

$$\mathcal{G} := \{g_0^{a,K} : (a, K) \in \mathbb{R}^2\},$$

where $g_0^{a,K}$ corresponds to a generic rule of the form [\(12\)](#). The market observes the rule before forming its $t = 0$ forecast. Importantly, commitment technology binds the policy choice at $t = 0$ alone. After $t = 0$, the central bank again acts sequentially, and hence uses g_1^* .

Definition 5 (Commitment policy). *For every $g_0^{a,K} \in \mathcal{G}$, let its market response be $\mu_{g_0} \equiv \mu^{a,K}$ as in [Definition 3](#). A commitment policy, thus, is a rule $g_0 \in \mathcal{G}$ solving*

$$\min_{g_0 \in \mathcal{G}} \mathbb{E}^C \left[(y_0 - y_0^n)^2 \right], \quad r_0 = g_0(A_0), \quad f_0 = \mu_{g_0}(X_0, r_0), \quad g_1 = g_1^*.$$

Note that by definition, the continuation equilibria at $t \geq 1$ are the same under the two solution concepts.

The distinction between the two solution concepts can now be stated more precisely. In discretionary equilibria, a unilateral central bank deviation changes the chosen rate while holding the market's strategy μ fixed, as explained above. Under commitment, the central bank instead acts ex-ante choosing a publicly observed rule. The market's complete strategy specifies a response μ_g following *every possible rule* g . Thus, holding the market's strategy fixed while the central bank deviates from g to $\tilde{g}$ changes the relevant continuation response from μ_g to $\mu_{\tilde{g}}$. It is in this precise sense that the commitment bank *internalizes the effect of its rule on the market's interpretation*.

The time inconsistency arises after A_0 is observed. Absent commitment (i.e., enforcement of the pre-announced rule), the central bank would choose a rate while holding the selected map μ_{g^c} fixed. Unless g^c is itself a discretionary equilibrium, the announced rule fails condition (ii) and the central bank would deviate from what it previously announced.

Remark 4 (On the choice of intercept). *Fixing the intercept in (12) at $\bar{\rho}$ is without loss under either solution concept. Since it is known to the market, an intercept c does not affect the policy surprise and hence does not affect belief revision. Under commitment, for fixed (a, K) , replacing $\bar{\rho}$ by c adds $\gamma^2(c - \bar{\rho})^2$ to expected loss, so the optimal intercept is $\bar{\rho}$. Under discretion, the same conclusion follows from the fixed-point condition: under the maintained unique-solution convention, matching the intercept of the central bank's pointwise best response to the conjectured intercept gives $c = \bar{\rho}$.*

The above is a partial-commitment problem. If the $t = 0$ choice also included commitment to a $t = 1$ rule

$$r_1 = \bar{\rho} + a_1(\hat{\rho}_1 - \bar{\rho}) - K_1 \hat{y}_1^n,$$

the normalized sensitivity of the market's $t = 0$ forecast would become

$$\chi_{a_1, K_1}(b, s) = \frac{\gamma K_1 \varphi_n b + \gamma a_1 \varphi_\rho s}{b^2 + s},$$

which reduces to $\chi(b, s)$ at $(a_1, K_1) = (1, 1/\gamma)$. Full commitment would therefore give the central bank an additional margin: it could use the future rule's response coefficients to affect current inference. That problem is not studied here.

Finally, note that the restriction to $\mathcal{G}$ binds under commitment: the central bank's ex-ante possibilities are *restricted* to linear rules. An unrestricted commitment problem would instead allow all measurable rules and require specifying a market-response selection for each such rule. As discussed in [Section 3.4](#), the market's updating is a nonlinear function of

the rule, so the unrestricted optimum is not necessarily linear, unlike in the corresponding economies without private information. The paper does not study this.

D Proofs

Throughout the proofs we write

$$h(b, s) := \gamma + \chi(b, s), \quad \Pi_n(b, s) := 1 - \lambda_n \varphi_n + \lambda_n b \chi(b, s),$$

$$e(a, b, s) := b(\gamma a + \chi(b, s)) - \Pi_n(b, s),$$

so that h and Π_n are the pass-through of demand- and natural-output-assessment errors respectively, e is the systematic stabilization error, and consequently the loss in [Lemma 2](#) can be written as

$$L(a, b; s) = e(a, b, s)^2 \sigma_{n,C}^2 + (a - 1)^2 \gamma^2 \sigma_{\rho,C}^2 + h(b, s)^2 \sigma_{\varepsilon,C}^2 + \Pi_n(b, s)^2 \sigma_{\eta,C}^2.$$

D.1 Proof of [Lemma 1](#)

For $t \geq 2$ there is no residual uncertainty and $(y_t^n, \rho_t) = (0, \bar{\rho})$, so the central bank attains $y_t = y_t^n = 0$ by setting $r_t = \bar{\rho}$. At $t = 1$ the problem is just slightly less straightforward: since expectations are not affected by the policy decision due to the deterministic nature of next period outcomes, minimizing $\mathbb{E}_1^C[(y_1 - y_1^n)^2]$ subject to the IS equation yields $r_1 = \hat{\rho}_1 - \hat{y}_1^n / \gamma$. Substituting into $y_1 = \mathbb{E}_1^M y_2 - \gamma(r_1 - \rho_1) = -\gamma(r_1 - \rho_1)$ (as $\mathbb{E}_1^M y_2 = 0$) yields $y_1 = \hat{y}_1^n + \gamma(\rho_1 - \hat{\rho}_1)$, i.e., $\mathbb{E}_1^C y_1 = \hat{y}_1^n$.

D.2 Proof of [Proposition 1](#)

First, consider the case where the rate is not constant, i.e., $(a, K) \neq (0, 0)$. Since r_0 and y_1 are jointly Gaussian under the market beliefs, the conditional expectation of the latter is simply the linear projection on the former. Further, by [Lemma 1](#), assessment persistence, and the mean-zero restrictions in [Section 2.2](#), $y_1 = \hat{y}_1^n + \gamma(\rho_1 - \hat{\rho}_1)$ has prior mean $\varphi_n \lambda_n y_0^n$ under the market's beliefs. Thus,

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n + \frac{\text{Cov}_0^M(y_1, r_0)}{\text{Var}_0^M(r_0)} (r - \mathbb{E}_0^M[r_0]).$$

Given $r_0 = \bar{\rho} + a(\hat{\rho}_0 - \bar{\rho}) - K\hat{y}_0^n$, and the market beliefs, we have $\text{Var}_0^M(r_0) = a^2s + K^2$. Furthermore,

$$\text{Cov}_0^M(\hat{y}_1^n, \hat{y}_0^n) = \varphi_n \text{Var}_0^M(\hat{y}_0^n) = \varphi_n.$$

Moreover, (4) and (7) imply

$$\rho_1 - \hat{\rho}_1 = \varepsilon_\rho - \varphi_\rho(\hat{\rho}_0 - \rho_0).$$

Because ρ_0 is in the market's information set and ε_ρ is independent of the central bank's $t = 0$ assessments,

$$\text{Cov}_0^M(\rho_1 - \hat{\rho}_1, \hat{\rho}_0) = -\varphi_\rho \text{Var}_0^M(\hat{\rho}_0) = -\varphi_\rho s.$$

The remaining cross-block covariances vanish by independence. Therefore,

$$\begin{aligned} \text{Cov}_0^M(y_1, r_0) &= \text{Cov}_0^M(\hat{y}_1^n + \gamma(\rho_1 - \hat{\rho}_1), a\hat{\rho}_0 - K\hat{y}_0^n) \\ &= a \text{Cov}_0^M(\hat{y}_1^n, \hat{\rho}_0) - K \text{Cov}_0^M(\hat{y}_1^n, \hat{y}_0^n) \\ &\quad + \gamma a \text{Cov}_0^M(\rho_1 - \hat{\rho}_1, \hat{\rho}_0) - K\gamma \text{Cov}_0^M(\rho_1 - \hat{\rho}_1, \hat{y}_0^n) \\ &= -K\varphi_n - \gamma a \varphi_\rho s. \end{aligned}$$

Thus,

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n - \left(\frac{K\varphi_n + \gamma a \varphi_\rho s}{a^2 s + K^2} \right) (r - \mathbb{E}_0^M[r_0]).$$

For $a \neq 0$, substitute $K = ab$ to obtain

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n - \chi(b, s) \left(\frac{r - \mathbb{E}_0^M[r_0]}{a} \right).$$

For $a = 0$ and $K \neq 0$, note that the rate r fully reveals $\hat{y}_0^n$ and nothing about $\hat{\rho}_0$. Thus, the market simply scales the rate to recover $\hat{y}_0^n$ and we get

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n - \left(\frac{\varphi_n}{K} \right) (r - \mathbb{E}_0^M[r_0]).$$

Finally, if $(a, K) = (0, 0)$, the rate is constant and thus trivially the market does not learn anything from it, so the conditional expectation is simply the prior mean

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n.$$

The expressions for output in each case follow directly from the IS equation and the above conditional expectations.

D.3 Proof of Proposition 2

We first rule out $a = 0$ in raw (a, K) coordinates. If $K = 0$, the conjectured rate is constant and thus the market makes no update whatsoever upon observing the central bank's choice. Then, the market's prior expectation of next period's output is $\varphi_n \lambda_n y_0^n$, given its belief that the central bank's contemporaneous assessment has a mean $\lambda_n y_0^n$. Hence, the central bank's expected output gap following an arbitrary rate r is

$$\varphi_n \lambda_n \hat{y}_0^n - \gamma(r - \hat{\rho}_0) - \hat{y}_0^n.$$

More precisely, using the definitions of the central bank's assessments, its conditional expected loss is

$$\mathbb{E}_0^C[(y_0 - y_0^n)^2] = [-\gamma(r - \hat{\rho}_0) - (1 - \varphi_n \lambda_n) \hat{y}_0^n]^2 + \mathbb{E}_0^C[(\gamma(\rho_0 - \hat{\rho}_0) - (1 - \varphi_n \lambda_n)(y_0^n - \hat{y}_0^n))^2].$$

The second term is independent of r . Hence its unique minimizer satisfies

$$r_0^* - \bar{\rho} = (\hat{\rho}_0 - \bar{\rho}) - \left(\frac{1 - \varphi_n \lambda_n}{\gamma} \right) \hat{y}_0^n,$$

which evidently does not coincide with the constant conjecture. Thus, $(a, K) = (0, 0)$ cannot be a discretionary equilibrium.

Now suppose $a = 0$ and $K \neq 0$. The conjectured rule is $r_0 = \bar{\rho} - K \hat{y}_0^n$, so $\mathbb{E}_0^M[r_0] = \bar{\rho} - K \lambda_n y_0^n$. The projection in Proposition 1 gives the belief map following an arbitrary current rate r as

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n - \frac{\varphi_n}{K}(r - \bar{\rho} + K \lambda_n y_0^n) = -\frac{\varphi_n}{K}(r - \bar{\rho}). \quad (67)$$

Given this, the central bank's expected output gap is

$$-\left(\gamma + \frac{\varphi_n}{K}\right)(r - \bar{\rho}) + \gamma(\hat{\rho}_0 - \bar{\rho}) - \hat{y}_0^n.$$

In particular, its conditional expected loss is

$$\begin{aligned} \mathbb{E}_0^C[(y_0 - y_0^n)^2] &= \left[-\left(\gamma + \frac{\varphi_n}{K}\right)(r - \bar{\rho}) + \gamma(\hat{\rho}_0 - \bar{\rho}) - \hat{y}_0^n \right]^2 \\ &\quad + \mathbb{E}_0^C[(\gamma(\rho_0 - \hat{\rho}_0) - (y_0^n - \hat{y}_0^n))^2]. \end{aligned}$$

The second term is independent of r . If $\gamma + \varphi_n/K \neq 0$, the rate-choice problem therefore

has the unique solution

$$r_0^* - \bar{\rho} = \frac{\gamma}{\gamma + \varphi_n/K}(\hat{\rho}_0 - \bar{\rho}) - \left(\frac{1}{\gamma + \varphi_n/K} \right) \hat{y}_0^n.$$

Its coefficient on $\hat{\rho}_0 - \bar{\rho}$ is nonzero since $\gamma > 0$, while the conjectured rule's coefficient is zero, so the best response cannot coincide with the conjecture. If instead $\gamma + \varphi_n/K = 0$, equivalently $K = -\varphi_n/\gamma$, expected loss is independent of r and the rate choice is not unique. Thus no discretionary equilibrium satisfying the unique-solution convention can have $a = 0$. Consequently, it can be written in scale-composition form (a, b) and we will do so from here onward.

Fix a conjectured rule (a, b) and write the total effect of a change in the policy rate (keeping the induced belief map fixed) as

$$H := \gamma + \frac{\chi(b, s)}{a}.$$

The market's prior expected policy rate is given by

$$\begin{aligned} \mathbb{E}_0^M[r_0] &= \bar{\rho} + a(\mathbb{E}_0^M[\hat{\rho}_0] - \bar{\rho}) - ab\mathbb{E}_0^M[\hat{y}_0^n] \\ &= \bar{\rho} + a(\rho_0 - \bar{\rho}) - ab\lambda_n y_0^n, \end{aligned}$$

where the second equality uses $\mathbb{E}_0^M[\hat{\rho}_0] = \rho_0$ and $\mathbb{E}_0^M[\hat{y}_0^n] = \lambda_n y_0^n$ from (9). Thus, the belief map induced by the conjecture entails the following expected output after observing an arbitrary current rate r :

$$\mathbb{E}_0^M[y_1 \mid r_0 = r] = \varphi_n \lambda_n y_0^n - \frac{\chi(b, s)}{a} (r - \bar{\rho} - a(\rho_0 - \bar{\rho}) + ab\lambda_n y_0^n).$$

Substituting this into the IS curve gives

$$\begin{aligned} y_0 - y_0^n &= -\gamma(r - \rho_0) - \frac{\chi(b, s)}{a} (r - \bar{\rho} - a(\rho_0 - \bar{\rho}) + ab\lambda_n y_0^n) - (1 - \varphi_n \lambda_n) y_0^n \\ &= -H(r - \bar{\rho}) + (\gamma + \chi(b, s))(\rho_0 - \bar{\rho}) - b\chi(b, s)\lambda_n y_0^n - (1 - \varphi_n \lambda_n) y_0^n. \end{aligned}$$

By definition of the central bank's assessments, its conditional expected loss can be written as

$$\begin{aligned} \mathbb{E}_0^C[(y_0 - y_0^n)^2] &= [-H(r - \bar{\rho}) + (\gamma + \chi(b, s))(\hat{\rho}_0 - \bar{\rho}) - b\chi(b, s)\lambda_n \hat{y}_0^n - (1 - \varphi_n \lambda_n) \hat{y}_0^n]^2 \\ &\quad + \mathbb{E}_0^C \left[((\gamma + \chi(b, s))(\rho_0 - \hat{\rho}_0) - (1 + b\chi(b, s)\lambda_n - \varphi_n \lambda_n)(y_0^n - \hat{y}_0^n))^2 \right]. \end{aligned}$$

The second term is independent of r . The unique-solution convention gives $H \neq 0$, so the central bank's best response is

$$r_0^* - \bar{\rho} = \frac{\gamma + \chi(b, s)}{H}(\hat{\rho}_0 - \bar{\rho}) - \left(\frac{1 + b\chi(b, s)\lambda_n - \varphi_n\lambda_n}{H} \right) \hat{y}_0^n.$$

Note that the policy-dependent part of the central bank's objective is the square of an affine function with slope $-H$. Thus, the best response above is the unique global minimizer (strict convexity holds given $H \neq 0$).

Equilibrium requires this response to coincide with $r_0 - \bar{\rho} = a(\hat{\rho}_0 - \bar{\rho}) - ab\hat{y}_0^n$, and hence

$$a = \frac{\gamma + \chi(b, s)}{H}, \quad ab = \frac{1 + b\chi(b, s)\lambda_n - \varphi_n\lambda_n}{H}.$$

Substituting the definition of H into the first equality yields $a\gamma + \chi(b, s) = \gamma + \chi(b, s)$, and therefore $a_d = 1$. Rearranging, the second equality then becomes

$$\underbrace{b_d(\gamma + (1 - \lambda_n)\chi(b_d, s))}_{\equiv G(b_d; s)} = 1 - \varphi_n\lambda_n$$

which is (18).

It remains to establish existence and uniqueness. Expanding G yields

$$G(b; s) = \gamma b + \frac{\varphi_n b^2 + \gamma \varphi_\rho b s}{b^2 + s}(1 - \lambda_n).$$

For $b < 0$, its second term satisfies

$$\frac{\varphi_n b^2 + \gamma \varphi_\rho b s}{b^2 + s}(1 - \lambda_n) \leq \frac{\varphi_n b^2}{b^2 + s}(1 - \lambda_n) < \varphi_n(1 - \lambda_n) \leq 1 - \lambda_n,$$

while $\gamma b < 0$. Hence, $G(b; s) < 1 - \lambda_n \leq 1 - \varphi_n\lambda_n$ for all $b < 0$ and thus there can't be an equilibrium with $b < 0$. Furthermore, G is clearly continuous in b , and

$$\frac{\partial G}{\partial b} = \frac{\gamma b^4 + \gamma(2 - (1 - \lambda_n)\varphi_\rho)b^2 s + \gamma(1 + (1 - \lambda_n)\varphi_\rho)s^2 + 2(1 - \lambda_n)\varphi_n b s}{(b^2 + s)^2} > 0$$

for $b > 0$. Thus, G is strictly increasing for $b > 0$ and it equals zero at $b = 0$. Finally, let $\bar{b} = (1 - \lambda_n\varphi_n)\gamma^{-1}$. Note that

$$G(\bar{b}; s) = \gamma\bar{b} + (1 - \lambda_n)\bar{b}\chi(\bar{b}, s) = (1 - \lambda_n\varphi_n) + (1 - \lambda_n)\bar{b}\chi(\bar{b}, s) > 1 - \lambda_n\varphi_n,$$

since $\chi(\bar{b}, s) > 0$, $\bar{b} > 0$ and $\lambda_n < 1$. Hence, $G(\cdot; s)$ crosses $1 - \varphi_n\lambda_n$ exactly once, at some

$b_d \in (0, \bar{b})$. Thus, the discretionary equilibrium is unique and satisfies $b_d \in (0, \bar{b})$ as stated.

D.4 Proof of Lemma 2

Define the estimation errors as

$$\eta \equiv y_0^n - \hat{y}_0^n, \quad \varepsilon \equiv \rho_0 - \hat{\rho}_0.$$

Using the policy rule (13) and the market-centering assumptions in (9), the market's expected rate is

$$\mathbb{E}_0^M[r_0] = \bar{\rho} + a(\rho_0 - \bar{\rho}) - ab\lambda_n y_0^n.$$

The normalized policy surprise that enters the belief map is therefore

$$\begin{aligned} \frac{r_0 - \mathbb{E}_0^M[r_0]}{a} &= \hat{\rho}_0 - \rho_0 - b(\hat{y}_0^n - \lambda_n y_0^n) \\ &= -\varepsilon + b\eta - b(1 - \lambda_n)y_0^n. \end{aligned}$$

Likewise, substituting $\rho_0 = \hat{\rho}_0 + \varepsilon$ into the direct policy component gives

$$r_0 - \rho_0 = (a - 1)(\hat{\rho}_0 - \bar{\rho}) - ab\hat{y}_0^n - \varepsilon.$$

Hence, using (16) and $y_0^n = \hat{y}_0^n + \eta$, we get

$$\begin{aligned} y_0 - y_0^n &= -\gamma[(a - 1)(\hat{\rho}_0 - \bar{\rho}) - ab\hat{y}_0^n - \varepsilon] + \varphi_n \lambda_n y_0^n - \chi(b, s)(-\varepsilon + b\eta - b(1 - \lambda_n)y_0^n) - y_0^n \\ &= -\gamma[(a - 1)(\hat{\rho}_0 - \bar{\rho}) - ab\hat{y}_0^n - \varepsilon] + (\varphi_n \lambda_n - 1)(\hat{y}_0^n + \eta) \\ &\quad - \chi(b, s)(-\varepsilon + b\eta - b(1 - \lambda_n)(\hat{y}_0^n + \eta)) \\ &= (b(\gamma a + (1 - \lambda_n)\chi(b, s)) - (1 - \varphi_n \lambda_n))\hat{y}_0^n - \gamma(a - 1)(\hat{\rho}_0 - \bar{\rho}) \\ &\quad + (\gamma + \chi(b, s))\varepsilon + (\varphi_n \lambda_n - 1 - b\lambda_n \chi(b, s))\eta. \end{aligned}$$

Under the central-bank ex-ante restrictions, the two estimates and their estimation errors are orthogonal across and within blocks, so every cross term vanishes. Consequently,

$$\begin{aligned} \mathbb{E}^C[(y_0 - y_0^n)^2] &= (b(\gamma a + (1 - \lambda_n)\chi(b, s)) - (1 - \varphi_n \lambda_n))^2 \sigma_{n,C}^2 + \gamma^2(a - 1)^2 \sigma_{\rho,C}^2 \\ &\quad + (\gamma + \chi(b, s))^2 \sigma_{\varepsilon,C}^2 + (\varphi_n \lambda_n - 1 - b\lambda_n \chi(b, s))^2 \sigma_{\eta,C}^2, \end{aligned}$$

which, after rearrangement, becomes (21).

D.5 Proof of Proposition 3

Suppose that $a_c \neq 0$ and set $b_c = K_c/a_c$. All four terms in (21) are nonnegative. The regular discretionary rule $(1, b_d)$ is feasible and makes the first two terms zero. Hence, since all the terms in the loss are non-negative

$$\mathcal{P}(b_c, s) \leq L(a_c, b_c; s) \leq L(1, b_d; s) = \mathcal{P}(b_d, s).$$

Now, if $\lambda_n = 0$,

$$\mathcal{P}(b, s) = (\gamma + \chi(b, s))^2 \sigma_{\varepsilon, C}^2 + \sigma_{\eta, C}^2.$$

Thus, cancelling out the constants and taking square roots in the inequality above gives

$$\gamma + \chi(b_c, s) \leq |\gamma + \chi(b_c, s)| \leq \gamma + \chi(b_d, s),$$

where the final term is positive because $b_d > 0$ and thus $\chi(b_d, s) > 0$. Subtracting γ yields $\chi(b_c, s) \leq \chi(b_d, s)$.

For strictness, consider a change in b around the discretionary rule while holding $a = 1$. Let $h_d := h(b_d, s) > 0$. Differentiating (21) with respect to b and using $e(1, b_d, s) = 0$ gives

$$\begin{aligned} L_b(1, b_d; s) &= 2e(1, b_d, s) (h_d + b_d \chi_b(b_d, s)) \sigma_{\varepsilon, C}^2 + 2h_d \chi_b(b_d, s) \sigma_{\varepsilon, C}^2 \\ &= 2h_d \chi_b(b_d, s) \sigma_{\varepsilon, C}^2. \end{aligned}$$

If $\chi_b(b_d, s) \neq 0$, this derivative is nonzero. Moving b a sufficiently small distance in the direction opposite its sign, while keeping $a = 1$, therefore produces a feasible rule with loss strictly below $L(1, b_d; s)$. Since (a_c, b_c) is a global minimizer,

$$\mathcal{P}(b_c, s) = \sigma_{\varepsilon, C}^2 (\gamma + \chi(b_c, s))^2 + \sigma_{\eta, C}^2 \leq L(a_c, b_c; s) < L(1, b_d; s) = \mathcal{P}(b_d, s) = \sigma_{\varepsilon, C}^2 h_d^2 + \sigma_{\eta, C}^2.$$

It follows that $|\gamma + \chi(b_c, s)| < h_d$, and hence $\chi(b_c, s) < \chi(b_d, s)$.

It remains to show that this derivative cannot vanish away from the stated knife-edge. Differentiating with respect to b gives

$$\chi_b(b, s) = \frac{\varphi_n(s - b^2) - 2\gamma\varphi_\rho sb}{(b^2 + s)^2}.$$

Because $(b_d^2 + s)^2 > 0$, $\chi_b(b_d, s) = 0$ is equivalent to

$$\varphi_n(s - b_d^2) = 2\gamma\varphi_\rho s b_d. \tag{68}$$

Multiplying the numerator of $\chi(b_d, s)$ by $2b_d$ and using (68) gives

$$\begin{aligned} 2b_d(\varphi_n b_d + \gamma \varphi_\rho s) &= 2\varphi_n b_d^2 + 2\gamma \varphi_\rho s b_d \\ &= 2\varphi_n b_d^2 + \varphi_n(s - b_d^2) \\ &= \varphi_n(b_d^2 + s). \end{aligned}$$

Dividing by $2b_d(b_d^2 + s) > 0$ therefore gives

$$\chi(b_d, s) = \frac{\varphi_n}{2b_d}.$$

The discretionary condition $b_d(\gamma + \chi(b_d, s)) = 1$ then becomes

$$\gamma b_d + \frac{\varphi_n}{2} = 1,$$

so

$$b_d = \frac{1 - \varphi_n/2}{\gamma}. \quad (69)$$

Next, rearrange (68) as

$$s(\varphi_n - 2\gamma \varphi_\rho b_d) = \varphi_n b_d^2.$$

Substituting (69) yields

$$\begin{aligned} s \left[\varphi_n - 2\gamma \varphi_\rho \left(1 - \frac{\varphi_n}{2} \right) \right] &= \frac{\varphi_n}{\gamma^2} \left(1 - \frac{\varphi_n}{2} \right)^2, \\ \gamma^2 s [\varphi_n - (2 - \varphi_n)\varphi_\rho] &= \varphi_n \left(1 - \frac{\varphi_n}{2} \right)^2. \end{aligned}$$

The right-hand side is strictly positive. Thus this equality can hold only if $\varphi_n > (2 - \varphi_n)\varphi_\rho$, in which case it is equivalent to

$$\gamma^2 s = \frac{(1 - \varphi_n/2)^2 \varphi_n}{\varphi_n - (2 - \varphi_n)\varphi_\rho}.$$

Thus, we have shown that $\chi_b(b_d, s) = 0$ implies the two primitive conditions in the proposition. The contrapositive, thus, implies that away from that knife-edge, $\chi_b(b_d, s) \neq 0$. Hence, the preceding argument applies and the strict inequality $\chi(b_c, s) < \chi(b_d, s)$ holds.

Suppose now that $\lambda_n > 0$ and (22) holds. From the expression for χ_b above, its sign is that of $\varphi_n(s - b^2) - 2\gamma \varphi_\rho s b$, which is strictly decreasing in b and positive at $b = 0$. Hence, there is a unique $b^* > 0$ at which $\chi_b = 0$ and $\chi(\cdot, s)$ is hump-shaped with maximum at b^* (i.e., strictly increasing on $(0, b^*)$ and strictly decreasing thereafter). Further, the computation above

showing $\chi(b_d, s) = \varphi_n/(2b_d)$ used only that the numerator of χ_b vanishes at that point. Applying it at b^* , where the numerator vanishes by definition, gives $\chi(b^*, s) = \varphi_n/(2b^*)$. Evaluating that numerator at $b = m/\gamma$ and multiplying by γ^2 gives $\gamma^2 s(\varphi_n - 2\varphi_\rho m) - \varphi_n m^2$, so (22) states exactly that it is nonnegative there, which is equivalent to $\gamma b^* \geq m$. Now, note that

$$\frac{\partial}{\partial b} [b\chi(b, s)] = \frac{s[2\varphi_n b + \gamma\varphi_\rho(s - b^2)]}{(b^2 + s)^2},$$

and, by construction, for $0 < b \leq b^*$ the numerator of χ_b is nonnegative, so $\varphi_n(s - b^2) \geq 2\gamma\varphi_\rho s b \geq 0$ and hence $s \geq b^2$, making the derivative above strictly positive. The left-hand side of the discretionary condition (18) $\gamma b + (1 - \lambda_n)b\chi(b, s)$ is therefore strictly increasing on $(0, b^*]$; it vanishes at $b = 0$ and by the above, at b^* equals

$$\gamma b^* + (1 - \lambda_n)\frac{\varphi_n}{2} \geq m + \frac{(1 - \lambda_n)\varphi_n}{2} = 1 - \lambda_n\varphi_n.$$

Thus, it attains the RHS of (18) $1 - \lambda_n\varphi_n$ on $(0, b^*]$, and since the discretionary equilibrium is unique (Proposition 2), $b_d \leq b^*$.

Now, note that if $b \leq 0$,

$$\chi(b, s) - \chi(0, s) = \frac{b(\varphi_n - \gamma\varphi_\rho b)}{b^2 + s} \leq 0.$$

Since $0 < b_d \leq b^*$, the strict monotonicity of χ on $[0, b^*]$ implies $\chi(0, s) < \chi(b_d, s)$. Thus, if $b_c \leq 0$, $\chi(b_c, s) < \chi(b_d, s)$ holds, so we can focus on the case $b_c > 0$ for the rest of the proof.

Suppose now $\chi(b_c, s) > \chi(b_d, s)$. If $b_c \leq b_d$ then $0 < b_c \leq b_d \leq b^*$, where $\chi(\cdot, s)$ is increasing: a contradiction. Hence, $b_c > b_d$, and both factors of $b\chi(b, s)$ are then larger at b_c , so $b_c\chi(b_c, s) > b_d\chi(b_d, s)$. Since the two terms in $\mathcal{P}$ are positive and increasing in $\chi(b, s)$ and in $b\chi(b, s)$ respectively, it follows that $\mathcal{P}(b_c, s) > \mathcal{P}(b_d, s)$: another contradiction. Consequently, $\chi(b_c, s) \leq \chi(b_d, s)$.

For strictness in the $\lambda_n > 0$ case, note that at $a = 1$ the first term in the loss function (21) is $e(1, b, s)^2 \sigma_{n,C}^2$, which equals zero at b_d . Thus, differentiating (21) in b at $(1, b_d)$ leaves

$$L_b(1, b_d) = 2(\gamma + \chi(b_d, s))\chi_b(b_d, s)\sigma_{\varepsilon,C}^2 + 2\lambda_n(1 - \lambda_n\varphi_n + \lambda_nb_d\chi(b_d, s))\left.\frac{\partial[b\chi(b, s)]}{\partial b}\right|_{b_d}\sigma_{\eta,C}^2 > 0,$$

because by the arguments above, $b_d \leq b^*$ makes $\chi_b(b_d, s) \geq 0$ and the last derivative strictly positive. Thus, lowering b marginally while keeping $a = 1$ has loss strictly below $L(1, b_d; s)$, so $\mathcal{P}(b_c, s) \leq L(a_c, b_c; s) < \mathcal{P}(b_d, s)$ and in particular $b_c \neq b_d$. Now suppose $\chi(b_c, s) = \chi(b_d, s)$. By strict monotonicity on $(0, b^*]$, $b_c < b_d \leq b^*$ can't hold, so $b_c > b_d$. In that case,

$b_c\chi(b_c, s) > b_d\chi(b_d, s)$ and $\mathcal{P}(b_c, s) > \mathcal{P}(b_d, s)$: a contradiction. Hence, $\chi(b_c, s) < \chi(b_d, s)$.

D.6 Proof of Corollary 1

The result follows directly from the coefficient characterization in Proposition 4 in Appendix E.

E Additional Theoretical Results

Lemma 3 (Existence of a commitment optimum). *Suppose*

$$\mathcal{P}(b_d, s) < \gamma^2 \sigma_{\rho, C}^2.$$

Then, there exists a solution to the commitment problem, (a_c, K_c) . Moreover, every solution satisfies $a_c \neq 0$.

Proof. For every $(a, K) \neq (0, 0)$, define

$$\zeta(a, K; s) := \frac{\varphi_n K + \gamma \varphi_\rho s a}{a^2 s + K^2},$$

i.e., minus the slope of the market's output expectations with respect to the policy rate surprise (in the (a, b) representation of policy, $\zeta = \chi/a$). Direct substitution gives the commitment loss in raw coordinates:

$$\begin{aligned} \tilde{L}(a, K; s) = & [\gamma K + (1 - \lambda_n) K \zeta(a, K; s) - (1 - \lambda_n \varphi_n)]^2 \sigma_{n, C}^2 + \gamma^2 (a - 1)^2 \sigma_{\rho, C}^2 \\ & + [\gamma + a \zeta(a, K; s)]^2 \sigma_{\varepsilon, C}^2 + [1 - \lambda_n \varphi_n + \lambda_n K \zeta(a, K; s)]^2 \sigma_{\eta, C}^2, \end{aligned}$$

which reduces to $L(a, b; s)$ in Lemma 2 when $a \neq 0$ and we set $b = K/a$. Instead, when $a = 0$ and $K \neq 0$, it gives

$$\tilde{L}(0, K; s) = [\gamma K - (1 - \varphi_n)]^2 \sigma_{n, C}^2 + \gamma^2 \sigma_{\rho, C}^2 + \gamma^2 \sigma_{\varepsilon, C}^2 + \sigma_{\eta, C}^2.$$

The preceding formula is defined only for $(a, K) \neq (0, 0)$. It cannot generally be extended continuously to the origin: although $a\zeta(a, K; s)$ and $K\zeta(a, K; s)$ remain bounded as $(a, K) \rightarrow (0, 0)$, their limits depend on the direction of approach.

At the exact constant rule $(a, K) = (0, 0)$, the rate conveys no information, and hence the market retains its prior output forecast. Consequently, the associated loss is

$$L_{\text{const}} = (1 - \lambda_n \varphi_n)^2 (\sigma_{n, C}^2 + \sigma_{\eta, C}^2) + \gamma^2 (\sigma_{\rho, C}^2 + \sigma_{\varepsilon, C}^2).$$

Thus, the central bank's objective function with commitment is

$$\mathcal{L}(a, K; s) = \begin{cases} \tilde{L}(a, K; s), & (a, K) \neq (0, 0), \\ L_{\text{const}}, & (a, K) = (0, 0). \end{cases} \quad (70)$$

The function $\mathcal{L}$ is continuous on $\mathbb{R}^2 \setminus \{(0, 0)\}$, but not necessarily at the origin. We prove the existence of a solution to the central bank's commitment problem by showing that the sublevel set generated by the feasible discretionary rule is compact and bounded away from this discontinuity point.

Let

$$S := \{(a, K) \in \mathbb{R}^2 : \mathcal{L}(a, K; s) \leq \mathcal{P}(b_d, s)\}$$

be the aforementioned sublevel set. The discretionary rule $(a, K) = (1, b_d)$ is feasible and satisfies

$$\mathcal{L}(1, b_d; s) = \mathcal{P}(b_d, s),$$

so S is nonempty.

Now we show that S is bounded. For every $(a, K) \neq (0, 0)$,

$$K\zeta(a, K; s) = \frac{\varphi_n K^2 + \gamma \varphi_\rho s a K}{a^2 s + K^2}.$$

Using $2|aK|\sqrt{s} \leq a^2 s + K^2$, it follows that

$$|K\zeta(a, K; s)| \leq \varphi_n + \frac{\gamma \varphi_\rho \sqrt{s}}{2}.$$

Consequently, for some finite constant C (independent of (a, K)),

$$|\gamma K + (1 - \lambda_n)K\zeta(a, K; s) - (1 - \lambda_n \varphi_n)| \geq \gamma|K| - C.$$

Thus, the first term of $\mathcal{L}(a, K; s)$ diverges as $|K| \rightarrow \infty$. Moreover,

$$\gamma^2(a - 1)^2 \sigma_{\rho, C}^2 \rightarrow \infty \quad \text{as } |a| \rightarrow \infty.$$

Then,

$$\mathcal{L}(a, K; s) \rightarrow \infty \quad \text{as } \|(a, K)\| \rightarrow \infty,$$

so S is bounded.

We next show that S is bounded away from $a = 0$. By assumption,

$$\mathcal{P}(b_d, s) < \gamma^2 \sigma_{\rho, C}^2.$$

Thus, we can choose $\delta \in (0, 1)$ such that

$$\gamma^2(1 - \delta)^2 \sigma_{\rho, C}^2 > \mathcal{P}(b_d, s).$$

Then, for every *nonconstant* rule with $|a| < \delta$,

$$\mathcal{P}(b_d, s) < \gamma^2(a - 1)^2 \sigma_{\rho, C}^2 \leq \mathcal{L}(a, K; s)$$

Further, at the constant rule,

$$L_{\text{const}} \geq \gamma^2 \sigma_{\rho, C}^2 > \mathcal{P}(b_d, s).$$

Then,

$$S \subseteq \{(a, K) : |a| \geq \delta\}.$$

On the closed set $\{(a, K) : |a| \geq \delta\}$ and hence on S , the loss $\mathcal{L}$ reduces to $\tilde{L}$. Thus, since $\tilde{L}$ and $\mathcal{P}(b_d, s)$ are continuous, S is closed. Since we established above that it is also bounded and it is a subset of $\mathbb{R}^2$, it is compact.

By the Weierstrass theorem, $\mathcal{L}$ attains its minimum on S . By construction, any policy outside S has loss strictly greater than $\mathcal{P}(b_d, s)$, while S contains the *feasible* discretionary rule with loss $\mathcal{P}(b_d, s)$. Therefore, a minimizer of $\mathcal{L}$ on S is a global solution of the commitment problem.

Finally, every rule with $a = 0$, including the constant rule, has loss at least

$$\gamma^2 \sigma_{\rho, C}^2 > \mathcal{P}(b_d, s),$$

and the global minimum is at most $\mathcal{P}(b_d, s)$. Hence, every solution satisfies $a_c \neq 0$. $\square$

Proposition 4 (Local implementation). *Suppose (22) holds strictly. For all sufficiently small $\sigma_{\varepsilon, C}^2/\sigma_{n, C}^2$ and $\sigma_{\eta, C}^2/\sigma_{n, C}^2$, any optimal commitment rule satisfies*

$$a_c > 1, \quad 0 < b_c < b_d.$$

Proof. Note that L is homogeneous of degree 1 in the variance vector, so all that matters is

the ratio of the variances. Thus, define the normalized error variances

$$q_\varepsilon := \frac{\sigma_{\varepsilon,C}^2}{\sigma_{n,C}^2}, \quad q_\eta := \frac{\sigma_{\eta,C}^2}{\sigma_{n,C}^2}, \quad q := (q_\varepsilon, q_\eta).$$

We consider sufficiently small positive pairs (q_ε, q_η) , holding all other parameters fixed. Further, let $\Delta_\rho := \gamma^2 \sigma_{\rho,C}^2 / \sigma_{n,C}^2$ so that

$$L(a, b; s) = \left(e(a, b, s)^2 + \Delta_\rho (a - 1)^2 + q_\varepsilon h(b, s)^2 + q_\eta \Pi_n(b, s)^2 \right) \sigma_{n,C}^2.$$

Also let

$$G(b; s) := b(\gamma + (1 - \lambda_n)\chi(b, s)).$$

The discretionary coefficient is the unique positive solution of $G(b_d; s) = 1 - \lambda_n \varphi_n$.

We first note that a global commitment optimum exists for all sufficiently small positive pairs q . Under the variances parameterization above, the sufficient condition in [Lemma 3](#),

$$\mathcal{P}(b_d, s) < \gamma^2 \sigma_{\rho,C}^2,$$

is equivalent, after dividing by $\sigma_{n,C}^2$, to

$$q_\varepsilon h(b_d, s)^2 + q_\eta \Pi_n(b_d, s)^2 < \Delta_\rho. \tag{71}$$

Since $h(b_d, s)$, $\Pi_n(b_d, s)$, and $\Delta_\rho > 0$ are fixed, this condition holds for all sufficiently small positive pairs (q_ε, q_η) . [Lemma 3](#) then implies that the commitment problem has a global minimizer and that every global minimizer has $a \neq 0$.

Now let's start at $q = (0, 0)$. Since $\Delta_\rho > 0$, the second term in the objective is zero only if $a = 1$. Conditional on $a = 1$, the first term is zero only if $G(b; s) = 1 - \lambda_n \varphi_n$, and hence only if $b = b_d$. Further, note that in raw coordinates (a, K) , every policy with $a = 0$, including the constant rule, has (normalized) loss at least Δ_ρ when $q = (0, 0)$, because the central bank does not respond to its demand assessment. Thus, $(1, b_d)$ is the unique global optimum when $q = (0, 0)$.

The first-order conditions of the commitment problem can be written as

$$0 = \gamma b e(a, b, s) + \Delta_\rho (a - 1), \tag{72}$$

$$\begin{aligned} 0 &= (\gamma a + (1 - \lambda_n)(\chi(b, s) + b\chi_b(b, s))) e(a, b, s) \\ &\quad + q_\varepsilon h(b, s)\chi_b(b, s) + q_\eta \lambda_n \Pi_n(b, s)(\chi(b, s) + b\chi_b(b, s)). \end{aligned} \tag{73}$$

We next show that $(1, b_d)$ entails uniformly lower losses than all policies outside any neighborhood of it at $q = (0, 0)$. Let

$$L_0(a, b; s) = e(a, b, s)^2 + \Delta_\rho(a - 1)^2.$$

We claim that, for every neighborhood U of $(1, b_d)$, there exists $c_U > 0$ such that

$$L_0(a, b; s) \geq c_U \quad \text{for all } (a, b) \notin U.$$

To see this, suppose instead that there were a sequence outside U with $L_0(a_j, b_j; s) \rightarrow 0$. Policies with $a_j = 0$ are excluded because their loss is at least Δ_ρ . Otherwise, nonnegativity of the two terms in L_0 implies

$$a_j \rightarrow 1, \quad e(a_j, b_j, s) \rightarrow 0.$$

Moreover, b_j must be bounded. If $|b_j| \rightarrow \infty$, then $\chi(b_j, s) = O(1/b_j)$, and hence

$$e(a_j, b_j, s) = \gamma a_j b_j + O(1),$$

which cannot converge to zero. Now take any convergent subsequence $b_{j_k} \rightarrow \bar{b}$. Since $a_{j_k} \rightarrow 1$ and $e(a_{j_k}, b_{j_k}, s) \rightarrow 0$, continuity of e implies

$$0 = e(1, \bar{b}, s),$$

i.e., $G(\bar{b}; s) = 1 - \lambda_n \varphi_n$. Since G attains $1 - \lambda_n \varphi_n$ only at b_d , we must have $\bar{b} = b_d$. Hence, $b_j \rightarrow b_d$. Thus, $(a_j, b_j) \rightarrow (1, b_d)$, contradicting the requirement that the sequence remain outside U .

Since (22) holds strictly, the proof of [Proposition 3](#) gives $\gamma b^* > m$, where b^* is such that $\chi_b(b^*, s) = 0$, and also $b^* \chi(b^*, s) = \varphi_n/2$. Hence

$$G(b^*; s) = \gamma b^* + (1 - \lambda_n) \frac{\varphi_n}{2} > m + (1 - \lambda_n) \frac{\varphi_n}{2} = 1 - \lambda_n \varphi_n = G(b_d; s),$$

so $b_d < b^*$ by strict monotonicity of $G(\cdot; s)$. Hence $\chi_b(b_d, s) > 0$ and, as also shown there, $\partial [b\chi(b, s)] / \partial b > 0$ on $(0, b^*]$ and in particular at b_d .

The above plus a standard continuity argument implies that we can choose a sufficiently small neighborhood U of $(1, b_d)$, with closure in the positive quadrant, such that

$$\chi_b(b, s) > 0, \quad \chi(b, s) + b\chi_b(b, s) > 0$$

for all $(a, b) \in U$. Applying the argument above to this neighborhood gives $c_U > 0$ such that outside U ,

$$L_q(a, b; s) \geq L_0(a, b; s) \geq c_U,$$

where $L_q = L/\sigma_{n,C}^2$ denotes the normalized loss and the leftmost inequality follows from the fact that the error terms are nonnegative. On the other hand, the feasible policy $(1, b_d)$ has loss

$$L_q(1, b_d; s) = q_\varepsilon h(b_d, s)^2 + q_\eta \Pi_n(b_d, s)^2.$$

Hence, if

$$q_\varepsilon h(b_d, s)^2 + q_\eta \Pi_n(b_d, s)^2 < c_U, \tag{74}$$

every global minimizer must lie inside U : the feasible discretionary rule has loss below c_U , whereas every rule outside U with $a \neq 0$ has loss at least c_U . Note that both (74) and the earlier existence condition (71) hold for all sufficiently small pairs of error-variance ratios.

Consequently, every optimum has $a_c > 0$ and $b_c > 0$. The loss is differentiable there, so each optimum satisfies the first-order conditions (72)–(73). In (73), evaluated at an optimum, the coefficient multiplying $e(a_c, b_c, s)$ is positive, and the remaining term is strictly positive since $q_\varepsilon h(b_c, s)\chi_b(b_c, s) > 0$, while the natural-output-error term is nonnegative. Thus, $e(a_c, b_c, s) < 0$. Then, (72) implies

$$a_c - 1 = -\frac{\gamma b_c}{\Delta_p} e(a_c, b_c, s) > 0.$$

Finally, using the definitions of e and G ,

$$G(b_c; s) - (1 - \lambda_n \varphi_n) = e(a_c, b_c, s) - \gamma b_c (a_c - 1) < 0.$$

Since G is strictly increasing for $b > 0$ and $G(b_d; s) = 1 - \lambda_n \varphi_n > 0$, it follows that $b_c < b_d$. $\square$

F Additional Empirical Results

F.1 Alternative Policy-Surprise Measure

Figure 6 below shows the results of the specification (24) using the fourth quarterly Eurodollar/SOFR futures contract, ED4, on the sample that has both OIS1Y and ED4 data. While OIS1Y approximately measures the expected average overnight rate over the following year, ED4 measures a three-month rate roughly one year ahead. Thus, the latter places weight near the end of the one-year horizon instead of averaging expected rates over

the entire path. The coefficient levels differ, but the key findings are unchanged: for every horizon and either surprise measure, the 2010–2019 estimate is more negative than the corresponding pre-GFC and post-2021 estimates.




Figure 6: Estimates of (24) using the USMPD ED4 surprise. Markers and adjacent numbers report coefficients; whiskers show 95% HC1 confidence intervals. Horizontal axes show forward horizons in years. The common OIS1Y/ED4 sample matches Figure 3: 318 events, with 63, 25, 160, and 70 by chronological panel. All panels report percentage points per percentage-point surprise. The GFC panel (June 2008–December 2009) uses y-axis limits $[-10, 5]$, versus $[-2, 1]$ elsewhere.

F.2 TIPS-Based Inflation Compensation

The baseline exercise proxies the market’s inflation expectations using inflation swaps. To check that the results are not driven by the specifics of that market or by the construction of the swap series, we run the specification in (24) using Treasury Inflation-Protected Securities (TIPS) instead. Concretely, we use the Federal Reserve’s fitted inflation-compensation series, built from nominal Treasury and TIPS yield curves as detailed in [Gürkaynak et al. \(2007\)](#) and [Gürkaynak et al. \(2010\)](#). These adjust for the fact that the coupon and maturity structures of nominal Treasuries and TIPS differ, and thus simply subtracting yields would not be a proper apples-to-apples comparison.

Let $b_{t,h}^{\text{TIPS}}$ denote the continuously compounded zero-coupon inflation-compensation rate at maturity h . Using the two-, five-, and ten-year rates, we build forward inflation-compensation

rates over the same horizons used in the baseline analysis:

$$b_{t,0-2}^{\text{TIPS}} = b_{t,2}^{\text{TIPS}}, \quad b_{t,2-5}^{\text{TIPS}} = \frac{5b_{t,5}^{\text{TIPS}} - 2b_{t,2}^{\text{TIPS}}}{3}, \quad b_{t,5-10}^{\text{TIPS}} = \frac{10b_{t,10}^{\text{TIPS}} - 5b_{t,5}^{\text{TIPS}}}{5}.$$

Just as in the baseline analysis, for each policy event, the revision is the change from the last trading day before the event to the first trading day after it. Because the availability of the TIPS series differs from that of the baseline outcomes, we use the events for which the TIPS-based outcome and the policy surprise are both available. This leaves 301 events: 57 in the pre-GFC period, 25 during the GFC, 154 during 2010–2019, and 65 during 2022–2026.

Figure 7 reports the results. The central regime comparison is preserved. During 2010–2019, the point estimates are more negative at every horizon than in the pre-GFC and post-2021 subsamples. The GFC point estimates are more negative than those for 2010–2019 at all three horizons, with wider confidence intervals.




Figure 7: Estimates of (24) using OIS1Y and TIPS-based breakeven inflation. Markers and adjacent numbers report coefficients; whiskers show 95% HC1 confidence intervals. Horizontal axes show forward horizons in years. The sample comprises 301 available events from the event universe in Figure 3, with 57, 25, 154, and 65 by chronological panel. All panels report percentage points per percentage-point surprise. The GFC panel (June 2008–December 2009) uses y-axis limits $[-10, 5]$, versus $[-2, 1]$ elsewhere.

Note that just as with the baseline measure, TIPS-based inflation compensation can also reflect inflation risk premia, on top of a liquidity premium since regular nominal treasuries

are more liquid than TIPS. However, the fact that we recover the main results using a different set of securities makes an explanation based on inflation-swap measurement issues considerably less plausible.

F.3 Bid and Ask Prices

Figures 8 and 9 present results for inflation swaps using the bid and ask prices, respectively. The main regime pattern is unchanged.




Figure 8: Estimates of (24) using OIS1Y and bid-side inflation-swap quotes. Markers and adjacent numbers report coefficients; whiskers show 95% HC1 confidence intervals. Horizontal axes show forward horizons in years. The sample matches Figure 3: 318 events, with 63, 25, 160, and 70 by chronological panel. All panels report percentage points per percentage-point surprise. The GFC panel (June 2008–December 2009) uses y-axis limits $[-10, 5]$, versus $[-2, 1]$ elsewhere.




Figure 9: Estimates of (24) using OIS1Y and ask-side inflation-swap quotes. Markers and adjacent numbers report coefficients; whiskers show 95% HC1 confidence intervals. Horizontal axes show forward horizons in years. The sample matches Figure 3: 318 events, with 63, 25, 160, and 70 by chronological panel. All panels report percentage points per percentage-point surprise. The GFC panel (June 2008–December 2009) uses y-axis limits $[-10, 5]$, versus $[-2, 1]$ elsewhere.

G Quantitative Model Details

Throughout this appendix we work with the interest rate centered around its long-run mean, $\tilde{i}_t \equiv i_t - \bar{\rho}$.

G.1 Information and Forecasting

Let $W_t = (y_t, \pi_t)'$ denote the vector of endogenous variables the market forecasts one period ahead. Under the within-period timing described in Section 5, the market's information immediately before and after the current policy decision is, respectively

$$\mathcal{I}_t^{M,-} \equiv \sigma \left(\{X_\tau\}_{\tau \leq t}, \{\tilde{i}_\tau\}_{\tau \leq t-1} \right), \quad (75)$$

$$\mathcal{I}_t^M \equiv \sigma \left(\{X_\tau, \tilde{i}_\tau\}_{\tau \leq t} \right). \quad (76)$$

In linear equilibria, the market forecast of next period’s output and inflation is

$$\mathbb{E}[W_{t+1} \mid \mathcal{I}_t^M] = \sum_{\ell=0}^{\infty} F_{\ell} X_{t-\ell} + \sum_{\ell=0}^{\infty} G_{\ell} \tilde{i}_{t-\ell}, \quad (77)$$

where F_{ℓ} is 2×3 and G_{ℓ} is 2×1 .

Note that unlike the similar setting in [Aoki \(2003\)](#), the market uses the entire history of fundamentals and policy rates to predict future endogenous outcomes.⁴⁶ Consequently, in a discretionary equilibrium, policy will depend on the central bank’s current expectations of the entire history of fundamentals. The reason behind this difference is that [Aoki \(2003\)](#) assumes that the central bank perfectly observes past realizations of fundamentals. Hence, there is no persistent information asymmetry the market has to keep track of: the only uncertainty the market faces regarding policy comes from the contemporaneous realization of the central bank’s signal, whose best predictor from their point of view is the current fundamentals alone, since the signal error is independent across time. This amounts to a significant simplification. Note further that for our purposes that assumption would not work, since in that case the market does not extract any information from current rates that is useful to predict future rates (conditional on it already knowing contemporaneous fundamentals) and hence the mechanism we study is absent.⁴⁷

Relation to [Svensson and Woodford \(2004\)](#). Our “rate-only” representation differs from the indicator-variable approach in [Svensson and Woodford \(2004\)](#). They also study optimal monetary policy with imperfect central-bank information in a linear forward-looking economy, but their state-space representation allows current endogenous variables to act as contemporaneous indicators of the central bank’s estimate. In their notation, forward-looking variables can be written as

$$z_t = G_1 X_t + (G - G_1) \widehat{X}_{t|t}^C, \quad z_t \equiv (y_t, \pi_t)'. \quad (78)$$

If z_t is in the private sector’s date- t information set, then these variables reveal linear combinations of $\widehat{X}_{t|t}^C$ beyond what is revealed by the policy rate alone. The infinite distributed-lag representation is then bypassed: the equilibrium can be solved with finite state-space objects,

⁴⁶[Aoki \(2003\)](#)’s physical environment is the same as ours, except it does not have cost-push shocks so the divine coincidence holds. Furthermore, in their model the central bank’s signals are noisy versions of endogenous objects (inflation and output) instead of fundamental shocks.

⁴⁷Note that this would be the case even if the signal’s errors weren’t independent across time. At $t + 1$, observing X_t allows the central bank to recover $e_t = s_t - X_t$ and subtract the component of e_{t+1} predictable from past errors from its new signal. Hence, its $t + 1$ assessment depends only on the new signal-error innovation, about which i_t contains no information. Consequently, the previous policy rate does not help the market predict future policy (conditional on public fundamentals).

Kalman gains, and covariance matrices.

Note that this is a rational-expectations fixed point, but it is not the “rate-only” information structure studied here. Concretely, in our model the only way information about the central bank’s signal errors gets to the private sector is via observed policy decisions (processed using the knowledge of fundamentals that the market directly observes by assumption). Naturally, this is crucial for our mechanism. In this sense, our “rate-only” information structure enforces a causal information constraint that is absent from the formulation in [Svensson and Woodford \(2004\)](#) (this is related to the much broader critique of equilibrium models in [Dogra \(2024\)](#)).

Finite-history approximation. As described in the main text above, we use a finite approximation of the state in the market’s forecasting operator (77). Concretely, we replace the market’s projection on the full history of fundamentals and past rates with a projection on the current plus the first H lags of fundamentals and policy rates. The central bank’s state is also truncated accordingly.

Delayed publication of central-bank estimates provides a natural reason why sufficiently old policy rates may eventually contain little additional information (thus making the finite H approximation reasonable): once the underlying assessment becomes public, earlier rates and fundamentals do not contain additional information. In the U.S., Greenbook/Tealbook forecasts are released to the public with a five-year lag⁴⁸. See [Townsend \(1983\)](#) for a related interpretation of finite-dimensional approximations to infinite-dimensional state spaces in rational-expectations models.

Since all relevant random variables are Gaussian, the finite-history market forecasts of future allocations are given by a linear projection of future output and inflation on the relevant variables up to horizon H , i.e.,

$$M_t \equiv \text{Proj}(W_{t+1} \mid Z_t^H) = \Theta Z_t^H, \quad (79)$$

where $Z_t^H \equiv (X'_t, X'_{t-1}, \dots, X'_{t-H}, \tilde{i}_t, \tilde{i}_{t-1}, \dots, \tilde{i}_{t-H})'$ and Θ contains the coefficients of the finite-history population projection

$$\Theta \in \arg \min_{\tilde{\Theta}} \mathbb{E} \left[\|W_{t+1} - \tilde{\Theta} Z_t^H\|^2 \right]. \quad (80)$$

Note that since older omitted histories are correlated with the retained regressors, Θ does not in general equal a direct truncation of the infinite-history coefficient sequences $\{F_\ell, G_\ell\}_{\ell \geq 0}$.

⁴⁸<https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/greenbook>

From here onward, M_t and Θ denote the finite-history projection and its associated coefficient matrix (since H is fixed, we drop the superscript H for notational convenience).

In the case of the central bank, its finite-history state retains its current posterior estimates of the $X_t, \dots, X_{t-H}$, together with H lagged policy rates. Formally, denoting the finite history of fundamentals by

$$\mathcal{X}_t^H \equiv (X_t', X_{t-1}', \dots, X_{t-H}')', \quad (81)$$

the central bank state is given by

$$\mathcal{S}_t^{C,H} \equiv \begin{pmatrix} \widehat{\mathcal{X}}_{t|t}^{C,H} \\ \tilde{i}_{t-1} \\ \vdots \\ \tilde{i}_{t-H} \end{pmatrix}, \quad (82)$$

where

$$\widehat{\mathcal{X}}_{\tau|t}^{C,H} \equiv \mathbb{E}[\mathcal{X}_\tau^H \mid \mathcal{I}_t^C], \quad \mathcal{I}_t^C \equiv \sigma(\{s_\tau^C, \tilde{i}_\tau\}_{\tau \leq t}). \quad (83)$$

G.2 Definition of Equilibria

Discretionary equilibria. As in the simpler model, a discretionary central bank takes the market's forecasting rule Θ as given. Note that for a given market forecast $M_t = (m_{y,t}, m_{\pi,t})'$, current allocations are pinned down by the IS curve and NKPC

$$\begin{aligned} x_t &= -y_t^n + \gamma(\rho_t - \bar{\rho}) - \gamma\tilde{i}_t + m_{y,t} + \gamma m_{\pi,t}, \\ \pi_t &= u_t + \kappa x_t + \beta m_{\pi,t}. \end{aligned} \quad (84)$$

Since the market forecast is a linear function of the private sector state, allocations also are. Thus, the discretionary central bank faces a linear-quadratic problem. Defining

$$R_t(i) \equiv \begin{pmatrix} x_t(i) \\ \pi_t(i) \\ i - \bar{\rho} \end{pmatrix}, \quad \mathcal{Q} \equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & \vartheta_\pi & 0 \\ 0 & 0 & \vartheta_i \end{pmatrix},$$

the discretionary problem can then be written recursively as

$$V(\mathcal{S}_t^{C,H}) = \min_i \mathbb{E}_t^C \left[\frac{1}{2} R_t(i)' \mathcal{Q} R_t(i) + \beta V(\mathcal{S}_{t+1}^{C,H}(i)) \right]. \quad (85)$$

The continuation value matters because the current policy rate becomes part of future public histories. If market forecasts condition on lagged rates, then today's i_t enters future Z_{t+j}^H through the lagged-rate block, which affects future market expectations and hence future policy tradeoffs. Thus, even under discretion, optimal policy features an intertemporal element (unlike in the standard New Keynesian model). Since the problem is linear quadratic, there exists an equilibrium where the optimal policy function is of the form

$$\tilde{i}_t = K^D(\Theta)\mathcal{S}_t^{C,H} \quad (86)$$

where, as the notation highlights, policy depends on the market forecasting structure Θ . This is because the central bank understands that the transmission of policy to macro outcomes depends on how the market interprets policy rates. Concretely, it is aware that expectations M_t are partly shaped by the contemporaneous policy instrument i_t .

Definition 6. *A stationary linear finite- H -history discretionary equilibrium is a set of stochastic processes $\{y_t, \pi_t, i_t, M_t\}_{t=0}^\infty$ and coefficient matrices (K^D, Θ) such that*

- **(Forecast Consistency)** *Given the policy rule (86), $M_t = \Theta Z_t^H$, where Θ solves the optimal forecasting problem (80) under the stationary distribution induced by the equilibrium.*
- **(Policy Optimality)** *Given the market's forecasting matrix Θ , $\tilde{i}_t = K^D \mathcal{S}_t^{C,H}$ solves (85).*
- **(Private-Sector Equilibrium)** *Given $M_t = (m_{y,t}, m_{\pi,t})'$, allocations satisfy (84), together with $y_t = x_t + y_t^n$.*

A game-theoretic definition analogous to the one in [Appendix C](#) for the baseline model applies here (and in the commitment case below).

Optimal linear commitment policy. A central bank with commitment chooses a policy rule ex-ante and thus acknowledges that the market's forecasting depends on its chosen rule. We restrict the space of policies to linear rules using the same variables as the discretionary rule.⁴⁹ Then, a commitment rule is a vector K^C such that

$$\tilde{i}_t = K^C \mathcal{S}_t^{C,H} \quad (87)$$

⁴⁹Thus, in general this does not correspond to the unrestricted Ramsey allocation. This is done for simplicity since computing the Ramsey allocation would be much harder than in the standard full information benchmark due to the non-linearity of the market's forecasting as a function of policy. Thus, this might be seen as providing a lower bound on the commitment gains from the unrestricted Ramsey problem.

where each K^C implies a stationary distribution of the state with corresponding market forecasting rule $\Theta(K^C)$. Let

$$\ell_t(K^C) \equiv \frac{1}{2} \left(x_t^2 + \vartheta_\pi \pi_t^2 + \vartheta_i \tilde{i}_t^2 \right)$$

denote per-period loss under that stationary distribution. The restricted commitment problem solved can then be written as⁵⁰

$$K^{C,*} \in \arg \min_{K^C} \mathbb{E}[\ell_t(K^C)], \quad (88)$$

where the latter expectation is the unconditional one corresponding to the induced stationary distribution.

Definition 7. *A stationary linear finite- H -history commitment equilibrium consists of stochastic processes $\{y_t, \pi_t, i_t, M_t\}_{t=0}^\infty$ and coefficient matrices (K^C, Θ) such that:*

- **(Forecast Consistency)** *Given the policy rule, $M_t = \Theta Z_t^H$, where Θ solves the market's optimal forecasting problem (80) under the stationary distribution induced by the equilibrium.*
- **(Policy Optimality)** *The policy rate satisfies $\tilde{i}_t = K^C \mathcal{S}_t^{C,H}$, where K^C solves (88).*
- **(Private-Sector Equilibrium)** *Given $M_t = (m_{y,t}, m_{\pi,t})'$, allocations satisfy (84), together with $y_t = x_t + y_t^n$.*

This formulation makes the difference with the discretionary equilibrium explicit. In the discretionary problem, the current policymaker treats the market's expectation formation process Θ as given when choosing i_t . Under commitment, the central bank chooses the systematic rule K^C and internalizes that changing the rule changes the informativeness of policy rates and therefore the equilibrium forecast matrix $\Theta(K^C)$.

The current-assessment-only commitment equilibrium is analogous to the above except it restricts to zero all coefficients on lagged assessments and lagged policy rates.

G.3 Solving the Equilibrium and Computing Population Objects

G.3.1 State Dynamics and Forecast Projections

We detail the central bank's filtering problem, the state dynamics, and population projections for a candidate policy rule K and forecast matrix Θ . These are used to solve both

⁵⁰This is equivalent to the normalized discounted loss $(1 - \beta) \mathbb{E} \sum_{j=0}^\infty \beta^j \ell_{t+j}$ if the economy starts from that stationary distribution.

discretionary and commitment equilibria.

Kalman filter. Recall the fundamental law of motion is given by

$$X_t = A_X X_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim (0, \Sigma_\varepsilon),$$

and the central bank's signal equation by

$$s_t^C = X_t + e_t^C, \quad e_t^C \sim (0, \Sigma_e),$$

where A_X , Σ_ε , and Σ_e are 3×3 primitives of the model pinning down the persistence of the fundamentals, the fundamental shock covariance, and the signal noise covariance, respectively.

The finite-history state in (81) has a law of motion given by

$$\mathcal{X}_t^H = \mathcal{A}_H \mathcal{X}_{t-1}^H + \mathcal{E}_t^H, \quad (89)$$

where

$$\mathcal{A}_H = \begin{pmatrix} A_X & 0 & 0 & \cdots & 0 \\ I & 0 & 0 & \cdots & 0 \\ 0 & I & 0 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & I & 0 \end{pmatrix}, \quad \mathcal{E}_t^H = (\varepsilon_t', 0', \dots, 0')'.$$

Further, the central bank's signal evolution can be written as

$$s_t^C = \mathcal{A}_{s,H} \mathcal{X}_t^H + e_t^C, \quad \mathcal{A}_{s,H} = \begin{pmatrix} I_3 & 0_{3 \times 3H} \end{pmatrix}. \quad (90)$$

The central bank's filtering problem on this augmented system delivers the central bank assessments (83).

Persistence of assessment errors. The independent fundamental and signal processes allow the filter for each contemporaneous fundamental to be written separately. For component $j \in \{n, \rho, u\}$, let ℓ_j denote its steady-state Kalman gain. The evolution of the contemporaneous assessment of fundamental j is then

$$\widehat{X}_{j,t|t} = (1 - \ell_j) \varphi_j \widehat{X}_{j,t-1|t-1} + \ell_j s_{j,t}^C.$$

Consequently, starting from the steady state, a one-time signal error of size $e_{j,0}^C$ generates a sequence of assessment errors given by

$$\widehat{X}_{j,h|h} = \ell_j [(1 - \ell_j)\varphi_j]^h e_{j,0}^C, \quad h \geq 0.$$

The decay rate of the induced assessment error is $(1 - \ell_j)\varphi_j$. At the baseline estimates, it is approximately 0.95845 for natural output and 0.90054 for demand, implying half-lives of 16.3 and 6.6 quarters, respectively. Both fundamental persistences are 0.995, so the difference reflects the smaller Kalman gain for natural output, which follows from the fact that its corresponding signal is particularly noisy.

State vector and stationary covariance. Rewrite the law of motion of the lag-augmented fundamentals' state (81) as

$$\mathcal{X}_t^H = \mathcal{A}_H \mathcal{X}_{t-1}^H + \mathcal{B}_H \varepsilon_t,$$

where

$$\mathcal{B}_H = \begin{pmatrix} I_3 \\ 0_{3H \times 3} \end{pmatrix}$$

and thus $\mathcal{B}_H \varepsilon_t = \mathcal{E}_t^H$ loads the innovation into the current fundamental block. Recall, further, the law of motion of the signal (90) and let $\mathcal{K}$ denote the steady-state Kalman gain for this lag-augmented system. Then, the central bank's posterior mean defined in (83) follows the stationary updating rule

$$\widehat{\mathcal{X}}_{t|t}^{C,H} = \mathcal{K} \mathcal{A}_{s,H} \mathcal{A}_H \mathcal{X}_{t-1}^H + (I_{3(H+1)} - \mathcal{K} \mathcal{A}_{s,H}) \mathcal{A}_H \widehat{\mathcal{X}}_{t-1|t-1}^{C,H} + \mathcal{K} \mathcal{A}_{s,H} \mathcal{B}_H \varepsilon_t + \mathcal{K} e_t^C.$$

Partitioning a generic fixed policy rule $K = \begin{pmatrix} K_{\mathcal{X}} & K_i \end{pmatrix}$, the policy rate can be written as

$$\tilde{i}_t = K_{\mathcal{X}} \widehat{\mathcal{X}}_{t|t}^{C,H} + K_i \begin{pmatrix} \tilde{i}_{t-1} \\ \vdots \\ \tilde{i}_{t-H} \end{pmatrix}. \quad (91)$$

To describe the joint evolution of fundamentals, central-bank assessments, and lagged policy rates, define the state vector

$$\mathcal{S}_t^H \equiv \begin{pmatrix} \mathcal{X}_t^H \\ \widehat{\mathcal{X}}_{t|t}^{C,H} \\ \tilde{i}_{t-1} \\ \vdots \\ \tilde{i}_{t-H} \end{pmatrix}. \quad (92)$$

Note that this is different from the central-bank state $\mathcal{S}_t^{C,H}$ since it includes the true lagged fundamentals $\mathcal{X}_t^H$. The state follows a linear law of motion

$$\mathcal{S}_t^H = A_S \mathcal{S}_{t-1}^H + B_S \eta_t, \quad \eta_t \equiv \begin{pmatrix} \varepsilon_t \\ e_t^C \end{pmatrix}, \quad \Sigma_\eta \equiv \begin{pmatrix} \Sigma_\varepsilon & 0 \\ 0 & \Sigma_e \end{pmatrix} \quad (93)$$

where

$$A_S = \begin{pmatrix} \mathcal{A}_H & 0 & 0 \\ \mathcal{K}\mathcal{A}_{s,H}\mathcal{A}_H & (I - \mathcal{K}\mathcal{A}_{s,H})\mathcal{A}_H & 0 \\ 0 & K_{\mathcal{X}} & K_i \\ 0 & 0 & (I_{H-1} \ 0) \end{pmatrix}, \quad B_S = \begin{pmatrix} \mathcal{B}_H & 0 \\ \mathcal{K}\mathcal{A}_{s,H}\mathcal{B}_H & \mathcal{K} \\ 0_{H \times 3} & 0_{H \times 3} \end{pmatrix},$$

where we omit the dimensions of the zero and identity matrices to keep the notation concise. The first two block rows of A_S correspond to fundamentals and the central bank's assessments. The third includes the policy rate chosen in the previous period, and the fourth shifts the remaining rate lags. The rate-lag rows of B_S are zero because those rates were chosen before the current innovations.

Define the unconditional covariance matrix of the state as

$$\Sigma_S \equiv \mathbb{E}[\mathcal{S}_t^H (\mathcal{S}_t^H)'], \quad (94)$$

which exists provided the matrix A_S is stable. Using the law of motion above, it is straightforward to show that the covariance matrix satisfies

$$\Sigma_S = A_S \Sigma_S A_S' + B_S \Sigma_\eta B_S'. \quad (95)$$

Forecast consistency. The market's optimal forecasting condition (80) can be written as the orthogonality condition

$$\mathbb{E} \left[(W_{t+1} - \Theta Z_t^H) (Z_t^H)' \right] = 0. \quad (96)$$

Since Z_t^H is a linear function of $\mathcal{S}_t^H$, we can write

$$Z_t^H = C_Z \mathcal{S}_t^H$$

for some matrix C_Z that selects the relevant components of $\mathcal{S}_t^H$. Similarly, current allocations

are linear in $\mathcal{S}_t^H$ for a fixed Θ and policy rule K , so we can write

$$W_t = (y_t, \pi_t)' = C_W(\Theta, K)\mathcal{S}_t^H,$$

where the rows of C_W are defined by substituting (79), $y_t = x_t + y_t^n$ and the policy rule into (84).

Since η_{t+1} is orthogonal to $\mathcal{S}_t^H$,

$$\mathbb{E}[W_{t+1}(Z_t^H)'] = C_W A_S \Sigma_S C_Z', \quad \mathbb{E}[Z_t^H(Z_t^H)'] = C_Z \Sigma_S C_Z'.$$

Thus, for a given vector of parameters $\theta \in \mathbb{R}^9$ and policy rule K , forecast consistency is equivalent to the fixed-point condition

$$\Theta = \underbrace{C_W(\Theta, K) A_S \Sigma_S C_Z' (C_Z \Sigma_S C_Z')^{-1}}_{\equiv \mathcal{F}_\theta(\Theta, K)}. \quad (97)$$

Given the policy rule K , for any conjectured forecast matrix Θ , $\mathcal{F}_\theta(\Theta, K)$ returns the optimal projection coefficient implied by the resulting allocation process. Forecasts are consistent when those coefficients equal the conjecture.

G.3.2 Discretionary Equilibrium

Let d_X be the dimension of X_t . The law of motion for the central bank's state (82) can be written as

$$\mathcal{S}_{t+1}^{C,H} = A_C \mathcal{S}_t^{C,H} + B_C \tilde{i}_t + \xi_{t+1}^C, \quad \mathbb{E}_t^C[\xi_{t+1}^C] = 0, \quad (98)$$

where

$$A_C = \begin{pmatrix} \mathcal{A}_H & 0_{d_X(H+1) \times H} \\ 0_{1 \times d_X(H+1)} & 0_{1 \times H} \\ 0_{(H-1) \times d_X(H+1)} & [I_{H-1} \quad 0_{(H-1) \times 1}] \end{pmatrix}, \quad B_C = \begin{pmatrix} 0_{d_X(H+1) \times 1} \\ 1 \\ 0_{(H-1) \times 1} \end{pmatrix}.$$

This representation uses the standard optimal forecasting implication that the expected next posterior mean is the corresponding predicted state.

Policy optimality. Consider a fixed Θ and define the central bank's forecast of the market forecast given a current demeaned rate $\tilde{i}$ as

$$\begin{pmatrix} \widehat{m}_{y,t}(\tilde{i}) \\ \widehat{m}_{\pi,t}(\tilde{i}) \end{pmatrix} \equiv \Theta \begin{pmatrix} \widehat{\mathcal{X}}_{t|t}^{C,H} \\ \tilde{i} \\ \tilde{i}_{t-1} \\ \vdots \\ \tilde{i}_{t-H} \end{pmatrix}$$

Define the selectors

$$\iota_n = (1, 0, 0)', \quad \iota_\rho = (0, 1, 0)', \quad \iota_u = (0, 0, 1)'.$$

Using (84), the central bank's expected allocations can be written as

$$\widehat{x}_t(\tilde{i}; \Theta) = -\iota'_n \widehat{X}_{t|t}^C + \gamma \iota'_\rho \widehat{X}_{t|t}^C - \gamma \tilde{i} + \widehat{m}_{y,t}(\tilde{i}) + \gamma \widehat{m}_{\pi,t}(\tilde{i}), \quad (99)$$

$$\widehat{\pi}_t(\tilde{i}; \Theta) = \iota'_u \widehat{X}_{t|t}^C + \kappa \widehat{x}_t(\tilde{i}; \Theta) + \beta \widehat{m}_{\pi,t}(\tilde{i}). \quad (100)$$

Splitting the forecast matrix according to the order of the regressors in Z_t^H , we get

$$\Theta =: \begin{pmatrix} \Theta_{\mathcal{X}} & \theta_0 & \Theta_- \end{pmatrix},$$

where $\Theta_{\mathcal{X}}$ multiplies $\mathcal{X}_t^H$, θ_0 multiplies the contemporaneous rate $\tilde{i}_t$, and Θ_- multiplies $(\tilde{i}_{t-1}, \dots, \tilde{i}_{t-H})'$. Let

$$\Theta_S \equiv \begin{pmatrix} \Theta_{\mathcal{X}} & \Theta_- \end{pmatrix}$$

be the part multiplying the central-bank state $\mathcal{S}_t^{C,H}$. Let d_n, d_ρ, d_u select from $\mathcal{S}_t^{C,H}$ the current posterior means of y_t^n , $\rho_t - \bar{\rho}$, and u_t , respectively. Finally define row vectors $a' = (1, \gamma)$ and $b' = (0, 1)$. Then, we can write

$$\widehat{R}_t(\tilde{i}; \Theta) \equiv \begin{pmatrix} \widehat{x}_t(\tilde{i}; \Theta) \\ \widehat{\pi}_t(\tilde{i}; \Theta) \\ \tilde{i} \end{pmatrix} = C_R(\Theta) \mathcal{S}_t^{C,H} + D_R(\Theta) \tilde{i}. \quad (101)$$

where

$$C_R(\Theta) \equiv \begin{pmatrix} C_x(\Theta) \\ C_\pi(\Theta) \\ 0_{1 \times \dim(\mathcal{S}_t^{C,H})} \end{pmatrix}, \quad D_R(\Theta) \equiv \begin{pmatrix} D_x(\Theta) \\ D_\pi(\Theta) \\ 1 \end{pmatrix}.$$

and

$$C_x(\Theta) \equiv -d'_n + \gamma d'_\rho + a'\Theta_S, \quad D_x(\Theta) \equiv -\gamma + a'\theta_0, \quad (102)$$

$$C_\pi(\Theta) \equiv d'_u + \kappa C_x(\Theta) + \beta b'\Theta_S, \quad D_\pi(\Theta) \equiv \kappa D_x(\Theta) + \beta b'\theta_0. \quad (103)$$

Thus, the central bank's contemporaneous expected loss can be written as $\frac{1}{2}\hat{R}'_t \mathcal{Q} \hat{R}_t$ plus a term involving posterior variances. Since the variance term is independent of the *current rate* (but not of the forecast matrix Θ !), it is irrelevant for the discretionary problem.

Since the problem is linear-quadratic and we are looking for linear equilibria, we guess the continuation value is quadratic,

$$V_\Theta(\mathcal{S}) = \frac{1}{2}\mathcal{S}'P_\Theta\mathcal{S} + \text{constant}$$

for some P_Θ to be determined. Consequently, minimizing the current expected loss plus the discounted continuation value gives

$$\tilde{i}_t = K^D(\Theta)\mathcal{S}_t^{C,H}, \quad K^D(\Theta) = -(\mathbf{R}_\Theta + \beta B'_C P_\Theta B_C)^{-1}(\mathbf{N}'_\Theta + \beta B'_C P_\Theta A_C), \quad (104)$$

where

$$\mathbf{Q}_\Theta \equiv C'_R \mathcal{Q} C_R, \quad \mathbf{N}_\Theta \equiv C'_R \mathcal{Q} D_R, \quad \mathbf{R}_\Theta \equiv D'_R \mathcal{Q} D_R.$$

Substituting the policy into the Bellman Equation, then, pins down P_Θ as the solution to

$$P_\Theta = \mathbf{Q}_\Theta + \beta A'_C P_\Theta A_C - (\mathbf{N}_\Theta + \beta A'_C P_\Theta B_C)(\mathbf{R}_\Theta + \beta B'_C P_\Theta B_C)^{-1}(\mathbf{N}'_\Theta + \beta B'_C P_\Theta A_C). \quad (105)$$

Since we are looking for a stationary equilibrium, we further require that the optimal policy (104), after substituting the solution to (105), induces stationary state dynamics. Concretely, all the eigenvalues of the transition matrix $A_C + B_C K^D(\Theta)$ (obtained by substituting the policy rule into (98)) must have magnitudes strictly below one. This ensures that the effects of initial conditions vanish and, given finite innovation variances, a finite stationary covariance exists.

Thus, as per Definition 6 above, equations (104), (105), (96) and (84) determine the discretionary equilibria.

Numerical implementation and equilibrium selection. For an arbitrary candidate parameter vector $\theta \in \mathbb{R}^9$, the discretionary equilibrium fixed point can be written as

$$\Theta = \underbrace{\mathcal{F}_\theta(\Theta, K_\theta^D(\Theta))}_{\equiv \mathcal{T}_\theta(\Theta)}, \quad (106)$$

where $K_\theta^D(\Theta)$ is given by (104) and $\mathcal{F}_\theta$ is defined in (97). Consequently, $\mathcal{T}_\theta$ incorporates both the forecast-consistency condition and the central bank's best response.

Thus, each evaluation of (106) entails computing the central bank's best response above, solving for the implied state stationary covariance (95) and using it to compute the population moments needed to update the forecast coefficients via $\mathcal{F}_\theta$.

The baseline implementation solves the fixed point (106) using Anderson-accelerated fixed-point iteration with depth five and relaxation 0.5. An additional nonlinear root finder is available as fallback and for final polishing. Accepted solutions must satisfy the above forecast-consistency and filtering conditions at tight numerical tolerances. We also require stationary state dynamics as stated above and a (numerically) non-singular covariance matrix of the market's forecasting regressors, ensuring finite population moments and uniquely determined forecast coefficients.

For some parameter vectors, the discretionary problem has more than one stationary fixed point. Among those, we restrict attention to locally E-stable equilibria, using the adaptive-learning criterion discussed by Evans and Honkapohja (2001).⁵¹ Intuitively, as agents update their forecasting coefficients using observed outcomes, their revised expectations also feed back into policy and allocations. The criterion requires this feedback to correct small initial errors in the forecasting coefficients instead of reinforcing them. Formally, this entails a constraint on the eigenvalues of the Jacobian of the mapping $\mathcal{T}_\theta$ at the fixed point Θ^* ,

$$\max_{\lambda \in \sigma(J_\theta)} \text{Re}(\lambda) < 1, \quad J_\theta \equiv D_\Theta \mathcal{T}_\theta(\Theta^*).$$

G.3.3 Optimal Commitment Policy

Naturally, the commitment problem uses the same central-bank filter, information state, allocation equations, and population-projection condition as the discretionary problem. The difference is that the central bank chooses the coefficients K^C directly in (88). Thus, for each candidate rule, we compute the market forecasts it induces and evaluate the resulting loss.

⁵¹While in principle there may exist multiple locally E-stable equilibria, that did not happen in any of our numerical exercises so the “refinement” allowed us to recover a unique allocation.

Forecast consistency. Fix a candidate policy K^C . Combining this rule with the fundamental processes and the Kalman filter determines the state's law of motion (93). We compute its stationary covariance as in (95) above. Like for the discretionary equilibrium, we require the covariance matrix of the market's forecasting regressors to be numerically non-singular.

Note that given a policy rule, these state dynamics and their covariances do not depend on the conjectured forecast matrix Θ . Allocations, however, naturally do depend on Θ linearly via (84). Consequently, the population-projection update is affine in the forecast coefficients

$$\text{vec } \mathcal{F}_\theta(\Theta, K^C) = a(K^C) + B(K^C) \text{vec } \Theta.$$

for some matrix $B(K^C)$ and vector $a(K^C)$. Given a candidate policy rule, then, the fixed-point condition (97) can be written as a linear system in $\text{vec } \Theta$,

$$[I - B(K^C)] \text{vec } \Theta(K^C) = a(K^C). \quad (107)$$

We solve this directly.

Policy optimization. We minimize the central bank's objective (88) over policy coefficients, recomputing the stationary covariance and consistent market forecasts (i.e., the solutions to (107)) for each candidate rule. The simple commitment rule is obtained by restricting the coefficients on lagged assessments and lagged rates to zero. Note that the central bank's objective (88) is non-convex in general. Thus, we can only guarantee that our numerical solution is a local minimum. To address this, we use multiple starting points in the numerical optimization and check local optimality of the solution.

G.3.4 Model-implied Moments

This section derives the model counterparts to the data moments used in the estimation. Once the stationary equilibrium has been found using the procedure above, the same state-space representation allows computing these moments analytically. Concretely, let v_t be any scalar, for example π_t , x_t , or $\tilde{i}_t$. Since we can write

$$v_t = c_v \mathcal{S}_t^H,$$

for some c_v , we can compute its second moments directly as

$$\text{Var}(v_t) = c_v \Sigma_{\mathcal{S}} c_v', \quad \text{Cov}(v_t, v_{t-1}) = c_v A_{\mathcal{S}} \Sigma_{\mathcal{S}} c_v'.$$

Macro moments. Since all variables are measured as deviations from steady state, the AR(1) coefficient can be computed by projecting v_t on v_{t-1} alone. Thus, it can be computed as

$$\rho_v = \frac{\text{Cov}(v_t, v_{t-1})}{\text{Var}(v_t)} = \frac{c_v A_{\mathcal{S}} \Sigma_{\mathcal{S}} c'_v}{c_v \Sigma_{\mathcal{S}} c'_v}. \quad (108)$$

Similarly, the corresponding AR(1) residual variance can be computed as

$$\sigma_{v, \text{AR1}}^2 = \text{Var}(v_t) - \rho_v^2 \text{Var}(v_{t-1}) = c_v \Sigma_{\mathcal{S}} c'_v - \frac{(c_v A_{\mathcal{S}} \Sigma_{\mathcal{S}} c'_v)^2}{c_v \Sigma_{\mathcal{S}} c'_v}. \quad (109)$$

Applying these formulas, we obtain the six macro moments

$$\rho_{\pi}, \rho_x, \rho_i, \quad \sigma_{\pi, \text{AR1}}^2, \sigma_{x, \text{AR1}}^2, \sigma_{i, \text{AR1}}^2.$$

Signal noise and natural-output uncertainty. For each fundamental $j \in \{n, \rho, u\}$, let

$$V_j \equiv \frac{\sigma_j^2}{1 - \varphi_j^2}, \quad q_j \equiv \frac{(\sigma_j^e)^2}{V_j}.$$

Thus, V_j is fundamental j 's unconditional variance and q_j is the ratio of the signal-noise variance and the unconditional variance of the corresponding fundamental. Let $P_{t|t}^{C,H}$ denote the stationary covariance matrix from the central-bank filter (after observing the corresponding contemporaneous signal) and let $d_n^{\mathcal{X}}$ select its current natural-output component. Again letting $\theta \in \mathbb{R}^9$ denote an arbitrary vector of primitives, the model counterpart of the HLW target is

$$\mathcal{H}^{\text{model}}(\theta) = \sqrt{(d_n^{\mathcal{X}})' P_{t|t}^{C,H}(\theta) d_n^{\mathcal{X}}}. \quad (110)$$

The data target is the average of the 120 one-sided, post-signal posterior standard deviations from the re-estimated HLW filter over 1990Q1–2019Q4.

Greenbook/Tealbook moments. For a given variable z_t , its model posterior-variance share is obtained by projecting it onto the central bank's state $\mathcal{S}_t^{C,H}$ defined in (82),

$$R_z^{\text{model}} = \frac{\text{Var}(z_t - \text{Proj}(z_t | \mathcal{S}_t^{C,H}))}{\text{Var}(z_t)} = 1 - \frac{\text{Cov}(z_t, \mathcal{S}_t^{C,H}) \text{Var}(\mathcal{S}_t^{C,H})^{-1} \text{Cov}(\mathcal{S}_t^{C,H}, z_t)}{\text{Var}(z_t)}. \quad (111)$$

We compute (111) for inflation π_t and real GDP growth, $y_t - y_{t-1}$. The annualization factor applied to GDP growth cancels from the variance ratio. Because a conditional-expectation error has mean zero, its posterior variance equals its mean squared error. Thus, the posterior error-variance share coincides with the posterior MSE share. The empirical counterpart is the

variance of later-vintage-minus-nowcast revisions divided by the variance of the later-vintage estimate.

Inflation-swaps moment. Let $Z_t^{H,-}$ denote the pre-policy market state, obtained by removing the contemporaneous policy rate from Z_t^H . For any variable a , define its finite-history revision upon seeing the policy rate

$$\Delta_{t,H}^M a \equiv \text{Proj}(a \mid Z_t^H) - \text{Proj}(a \mid Z_t^{H,-}).$$

The model revisions can then be written as

$$\Delta P_{t,1:20}^e \equiv \Delta_{t,H}^M \left(\frac{1}{4} \sum_{h=1}^{20} \pi_{t+h} \right), \quad (112)$$

$$mps_t^M \equiv \Delta_{t,H}^M \left(\frac{1}{4} \sum_{h=0}^3 i_{t+h} \right), \quad (113)$$

and the relevant sensitivity coefficient is

$$\beta_{1:20}^P \equiv \frac{\text{Cov}(\Delta P_{t,1:20}^e, mps_t^M)}{\text{Var}(mps_t^M)}.$$

Note that both revisions are linear functions of the state $\mathcal{S}_t^H$. Consequently, the covariance and variance defining $\beta_{1:20}^P$ can be evaluated analytically using the stationary state covariance. We omit the explicit matrix expressions for the sake of brevity.

G.4 Estimation Details

Recall we denote by $\theta \in \mathbb{R}^9$ an arbitrary vector of parameters. Let j index our relevant moments and let m_j^{data} and $m_j(\theta)$ denote the corresponding data moment and its model counterpart. Further, let $g_j(\theta) \equiv m_j(\theta) - m_j^{\text{data}}$, and let v_j be the estimated sampling variance of the corresponding data moment. Our minimum-distance objective function is

$$Q(\theta) = \sum_{j \in \mathcal{M}} \frac{g_j(\theta)^2}{v_j} + 10 \sum_{j \in \mathcal{G}} \frac{g_j(\theta)^2}{v_j} + 10 \frac{g_{\beta^P}(\theta)^2}{v_{\beta^P}} \quad (114)$$

where in addition we impose $\mathcal{H}^{\text{model}}(\theta) = \mathcal{H}^{\text{data}}$ and $\mathcal{M}$ contains the six macro moments and $\mathcal{G}$ the two Greenbook/Tealbook moments.

Sampling variances. The sampling variances are estimated by bootstrap. For each macro series, we use 5,000 draws from a univariate AR(1) sieve bootstrap, resampling centered

AR(1) residuals and removing the first 200 quarters' observations from the sample. For the output-gap moments, the bootstrap is conducted using the Hamilton-filtered series extracted from real GDP as introduced in the main text. For each Greenbook/Tealbook moment, we use 5,000 target/quarter resamples of the paired revision and later-vintage estimate and recompute the forecast-revision variance share. For the inflation-swaps moment, we use 10,000 bootstrap draws that resample event dates. In each draw, we re-estimate the 0-2 year and 2-5 year response coefficients using the same resampled events and recompute the combined target $\beta_{1:20}^P = 2\beta_{1:8}^\pi + 3\beta_{9:20}^\pi$.

Table 7: Minimum-Distance Targets and Weights

Moment	Data	Model	Sampling variance	Multiplier	Q contribution
AR(1) (π)	0.1321	0.2440	0.007482	1	1.6737
AR(1) (x)	0.8943	0.9219	0.002562	1	0.2989
AR(1) (i)	0.9643	0.9121	0.001644	1	1.6579
Residual variance (π)	2.7672	1.3397	1.308162	1	1.5577
Residual variance (x)	1.1785	0.2946	0.038263	1	20.4230
Residual variance (i)	0.1650	0.2031	0.000884	1	1.6455
Natural-output uncertainty ($\mathcal{H}$)	1.7578	1.7578	–	imposed	–
Revision-variance share (π)	0.2414	0.2680	0.005532	10	1.2765
Revision-variance share (GDP growth)	0.5876	0.6296	0.044610	10	0.3968
$\beta_{1:20}^P$	–2.4486	–2.0117	0.820098	10	2.3271
				Total	31.2568

Notes: “Multiplier” is the target-priority weight applied on top of inverse variance weighting. The natural-output uncertainty target is imposed exactly and does not contribute to Q .

The estimated parameters are $(\sigma_n, \sigma_\rho, \sigma_u, \varphi_n, \varphi_\rho, \varphi_u, q_n, q_\rho, q_u)$. We enforce the following bounds

$$10^{-4} \leq \sigma_j \leq 20, \quad 0 \leq \varphi_j \leq 0.995, \quad 10^{-4} \leq q_n, q_u \leq 100, \quad 10^{-4} \leq q_\rho \leq 1000.$$

Imposing $\mathcal{H}$ exactly leaves eight free coordinates.

For every candidate parameter vector, the estimator solves the discretionary equilibrium using the procedure described in the previous section. The reported estimate has $Q = 31.2568$, and the following parameter constraints are binding: $\varphi_n = \varphi_\rho = 0.995$ and $\varphi_u = 0$.

G.5 Additional Results

G.5.1 Loss Decomposition

Table 8 below reports the decomposition of total loss into its three components for each of the relevant policy regimes.

Table 8: Loss decomposition in three policy regimes

Policy	Output gap	Inflation	Rate	Total
Discretion	0.98153	0.71223	0.60385	2.29760
Commitment: simple rule	0.05027	0.60789	0.06947	0.72762
Full commitment	0.03623	0.58282	0.08870	0.70776

G.5.2 Conditional Variances and Commitment Gains

Table 9 decomposes output-gap and inflation variance into variation in the central bank’s conditional mean and variation around that mean. Concretely, for $z_t \in \{x_t, \pi_t\}$,

$$\text{Var}(z_t) = \text{Var}(\mathbb{E}_t^C z_t) + \mathbb{E}[\text{Var}(z_t | \mathcal{I}_t^C)].$$

Table 9: Conditional Variances and the Sources of Commitment Gains

	Discretion	Simple rule	Full commitment
<i>A. Output-gap variance</i>			
Variation in the CB’s expected gap	0.05280	0.05786	0.05866
Variation around the CB’s expected gap	1.91026	0.04267	0.01380
Total	1.96306	0.10054	0.07246
<i>B. Inflation variance</i>			
Variation in the CB’s expected inflation	1.04268	1.09726	1.04749
Variation around the CB’s expected inflation	0.38177	0.11852	0.11816
Total	1.42446	1.21578	1.16564
<i>C. Contributions to the gain over discretion (%)</i>			
Combined conditional-error reduction	–	67.86	67.93
Output gap	–	59.48	59.64
Inflation	–	8.38	8.29
Conditional-mean allocation loss reduction	–	–1.90	–0.34
Rate-variance reduction	–	34.04	32.40
Total gain	–	100.00	100.00

Notes: Panels A and B report stationary variances. Panel C applies the baseline loss weights, $\vartheta_\pi = \vartheta_i = 1$, and the factor 1/2 in the objective. Negative entries indicate an increase in that loss component.

The reduction in loss attributable to variation around the central bank's expected output gap and inflation amounts to 1.06542 under the simple rule and 1.08004 under full commitment, accounting for approximately 68% of both commitment gains. Smaller rate deviations account for the rest, with a small increase in the loss from higher conditional average stabilization. Thus, the lion's share of the gains from commitment is about realized allocations relative to what the central bank anticipates.

G.5.3 Demand Innovations and Demand-Signal Errors

Table 10 documents the impact of demand-related shocks supporting the brief discussion in the main text. Each shock corresponds to a positive one standard deviation increase in the corresponding primitive (demand innovation and demand-signal error innovation).

Table 10: Demand Disturbances under Discretion and the Simple Rule

	Demand-signal error		Demand innovation	
	Discretion	Simple rule	Discretion	Simple rule
<i>A. Impact responses</i>				
Policy rate	0.16217	0.03529	0.01693	0.00369
Output gap	-0.29860	-0.08532	0.32309	0.06047
Inflation	-0.07694	-0.01912	0.10884	-0.02026
Expected next-period inflation	-0.07048	-0.01724	0.10210	-0.02193
Ex-ante real rate	0.23265	0.05253	-0.08517	0.02562
<i>B. Stationary contribution of each innovation</i>				
Output-gap variance	0.59373	0.03898	0.72678	0.05804
Loss	0.34845	0.02376	0.90465	0.16570

Notes: The real rate is $i_t - m_{\pi,t}$, where $m_{\pi,t}$ is the market's finite-history forecast of next-period inflation. Note that the shocks are a one-standard deviation increase in the corresponding primitive, and the standard deviations are different for the two shocks. This explains the different policy rate response magnitudes (the reaction coefficient is of course the same for both since the central bank does not see them separately but only their sum).

For the demand-signal error, the smaller rate hike under the simple rule entails a smaller fall in expected inflation and a much smaller contraction. Although signal errors are independent over time, their effects on the central bank's assessments persist as subsequent signals correct the initial mistake only gradually. Concretely, the effect of an initial demand-signal error decays at approximately 0.901 per quarter, compared with 0.958 for a natural-output signal error (see [Appendix G.3.1](#) for details). The resulting transitory tightening produces the conventional contractionary and disinflationary effects. The market's belief revision under discretionary policy significantly amplifies the contractionary effect of the initial error under discretion (see [Appendix G.5.6](#) for an exercise that shows this).

Instead, a genuine demand innovation persistently raises the natural real rate and, as the central bank’s subsequent signals arrive, its demand assessments. To understand the resulting inflation response, note that since a demand innovation leaves natural output constant, the IS equation implies

$$i_t - \mathbb{E}_t^M \pi_{t+1} = \rho_t + \frac{\mathbb{E}_t^M x_{t+1} - x_t}{\gamma}. \quad (115)$$

During the first periods following the shock, subsequent signals gradually correct the central bank’s initial underestimation of the demand innovation, while demand itself remains elevated due to its high persistence. After that “learning phase” is (essentially) over, under discretion the nominal-rate increase is approximately 1.32 times the remaining natural-rate increase and the equilibrium features positive expected inflation. Under the simple commitment rule, that ratio is approximately 0.50 and the equilibrium features negative expected inflation. In both cases, the resulting real rate remains close to the natural rate (given by ρ_t , since natural output growth is zero in this exercise) and thus, together with an essentially constant output gap, is consistent with (115) above and the remaining equilibrium conditions.⁵²

Thus, despite its more restrained nominal tightening, the simple commitment rule induces a higher real rate and a substantially smaller initial output expansion. Beyond the effects on impact, the stationary loss contribution of fundamental demand innovations falls by 81.7%.

G.5.4 Cost-Push Assessments and the Incentive to Tighten

The reduction in the response to cost-push assessments under commitment illustrates the distinction between the incentive to change today’s rate and the consequences of changing the systematic rule. We compare the effects of genuine cost-push innovations with signal errors, discuss the incentives a discretionary central bank faces, and finally consider the interaction with the other policy responses.

Fundamental innovations and signal errors. Table 11 reports the on-impact responses to positive one-standard-deviation innovations in the cost-push fundamental and its signal noise (starting from the steady state and setting all other innovations to zero as standard). Since the market observes the fundamental innovation, its associated policy response is anticipated and generates no announcement surprise. Under the simple rule, the zero persistence of the cost-push fundamental and the absence of direct policy history dependence also imply

⁵²While at odds with the standard view that higher nominal rates cool down the economy, this Fisherian phenomenon is a feature of standard New Keynesian models with rational expectations (see [García-Schmidt and Woodford \(2019\)](#) for a discussion of this and an alternative approach). Our analysis inherits this feature from the forward-looking New Keynesian block and the rational expectations equilibrium concept.

no response of expected next-period inflation. Under discretion, history dependence generates a small positive response. Despite its much larger rate increase, discretion reduces current inflation only slightly more than the simple rule.

Table 11: Cost-Push Disturbances under Discretion and the Simple Rule

	Fundamental innovation		Signal error	
	Discretion	Simple rule	Discretion	Simple rule
Policy rate	0.26692	0.00794	0.09228	0.00275
Output gap	−0.03453	−0.00127	−0.16992	−0.00664
Inflation	1.03115	1.03136	−0.04378	−0.00149
Expected next-period inflation	0.00059	0.00000	−0.04011	−0.00134

Notes: Impact responses to one-standard-deviation innovations ($\sigma_u = 1.03139$ and $\sigma_u^e = 0.35659$ respectively). Rates and inflation are in annual percentage points, the output gap is in percentage log-output points.

By contrast, a cost-push signal error produces an unexpected rate increase. The market does not observe the error and interprets the tightening partly as evidence of the central bank’s assessments of natural output and demand. These inferences generate a substantial fall in expected inflation and a contraction under discretion. The simple rule sharply attenuates the policy response and the resulting movements in expectations and allocations.

The incentive for a unilateral deviation. Facing perceived cost-push pressure, the central bank would like to lower the market’s inflation expectations to alleviate the trade-off between output and inflation stabilization. Under the simple commitment rule, an unexpectedly large tightening is interpreted mainly as evidence of a stronger demand assessment, and partly of a lower natural-output assessment. Since these are fairly persistent, they are expected by the market to continue to influence policy in subsequent periods, so it expects tighter future policy and revises expected output and inflation downward. Lower expected inflation helps tame current inflation without requiring such a large output gap contraction. Thus, the central bank has an incentive to tighten more today, even though it intends to return to the prescribed rule thereafter.

For example, suppose the central bank receives a positive cost-push signal of size σ_u . The simple rule prescribes a rate increase of 0.00794. Taking the market’s forecasting coefficients as given and expecting to return to that rule in subsequent periods, the central bank would instead choose an increase of 0.07563. This deviation lowers its conditional expectation of current inflation from 0.92078 to 0.88411, while its expected output gap falls from −0.00319 to −0.16683. Thus, its (conditional) expected loss improves after accounting for the inflation, output-gap, and rate losses and the effects on continuation losses.

The responses are chosen jointly. As the counterfactual analysis above illustrates, each policy response affects the propagation of errors beyond its corresponding assessment. To examine these interactions, Table 12 considers changes to the cost-push coefficients starting from both discretion and the simple rule, keeping all other coefficients fixed. Market forecasts are recomputed for each rule.

Table 12: Losses under Cost-Push Coefficient Replacements

Other coefficients	Cost-push coefficients	
	Discretion	Simple rule
Discretion	2.29760	2.34873
Simple rule	0.76067	0.72762

Notes: Stationary per-period losses at the baseline parameters. The cost-push block contains the coefficients on current and past cost-push assessments. All remaining assessment coefficients and lagged-rate coefficients are held at the policy specified in the first column.

With the other coefficients fixed at discretion, replacing the cost-push block lowers the loss attributable to cost-push disturbances from 0.58997 to 0.53195. However, it raises the combined loss attributable to natural-output and demand disturbances by 0.10915, more than offsetting that improvement. As in the exercise above, their corresponding realized rate paths are unchanged so these negative consequences are due to the change in the market’s interpretation of those rates. On the other hand, with the other coefficients fixed at the simple commitment rule, the same cost-push replacement lowers the cost-push contribution from 0.57731 to 0.53191, but raises the combined contribution of the other disturbances by only 0.01235. Hence, total loss falls. Thus, the restrained cost-push response is desirable as part of the jointly chosen rule: reducing that response alone from discretion does not necessarily improve stabilization.

Note that a rule combining the simple rule’s remaining coefficients with the discretionary cost-push block has a loss of 0.76067, compared with 0.72762 under the simple rule and 2.29760 under discretion. Thus, it achieves 97.9% of the gain from discretion to the simple rule despite retaining the stronger cost-push response.

G.5.5 Commitment Gains across Information Structures

We consider two environments that share the physical structure with our quantitative model but differ in their information structure: full information (i.e., contemporaneous fundamental shocks are common knowledge) and imperfect but public central bank information (i.e., the central bank has imperfect signals about fundamentals that are contemporaneously

observed by the market). We keep the baseline parameters at their baseline estimated values, except for the parameters governing the central bank’s information quality, which in the full information economy are such that the central bank observes the fundamentals perfectly.

In each economy, we compare discretion with the analogue of our simple commitment rule: a linear rule using only the three current fundamentals (full information) or the three contemporaneous central bank assessments (public and private assessments). [Table 13](#) reports the losses under both policy regimes and the associated gains from commitment in each economy.

Table 13: Losses and Gains from Commitment across Information Environments

Information	Discretion loss	Simple-rule loss	Gain (%)
Full information	1.09414	0.71322	34.81
Public assessments	1.10836	0.73088	34.06
Private assessments	2.29760	0.72762	68.33

The losses under a simple commitment rule are fairly similar across information environments, while our benchmark with private central-bank assessments entails substantially worse loss under discretionary policy.

Importantly, this ranking is not universal. In particular, with a rate-level penalty, the discretionary losses in the full-information and public-assessment environments become arbitrarily large near a critical value of demand persistence. This is because the rate-level penalty weakens discretionary stabilization, while persistent demand disturbances generate reinforcing effects through expected future output and inflation. As demand persistence approaches this critical value, the policy response becomes insufficient to offset this forward-looking reinforcing feedback, and thus the output and inflation responses grow without bound. It turns out that since our model gives the discretionary central bank leverage over the market expectations, its equilibria avoid these pathological features.

G.5.6 A Counterfactual Policy Rule

This exercise starts from the baseline discretionary rule and replaces its coefficients on natural-output assessments with the simple rule’s response. Concretely, write the original discretionary rule as

$$\tilde{i}_t^D = \sum_{\ell=0}^H k_{n,\ell}^D \hat{y}_{t-\ell|t}^n + R_t^D,$$

where $\hat{y}_{t-\ell|t}^n$ is the central bank’s period- t assessment of period- $t - \ell$ natural output, and R_t^D collects the demand-assessment, cost-push-assessment, and past rates terms. The counter-

factual rule is then

$$\tilde{i}_t^{\text{cf}} = k_{n,0}^{\text{simple}} \hat{y}_{t|t}^n + R_t^D,$$

that is, the current natural-output coefficient is replaced and all coefficients on past natural-output assessments are set to zero. The other policy coefficients are unchanged. Crucially, the market acknowledges the new policy rule and thus adjusts its forecasting structure to be consistent with this counterfactual rule. Naturally, in general this does not correspond to an equilibrium, since the rule is not optimal for the central bank.

We consider the effects of a one-standard-deviation demand-signal error in both the discretionary economy and this counterfactual one. Since fundamentals and signal errors are independent across blocks (i.e., natural output, demand, and cost-push), this leaves all natural-output assessments unchanged. Consequently, the implied interest rate path is the same under both rules. Nonetheless, the market’s interpretation of the policy surprise triggered by the policy response to the demand assessment differs across economies, and the resulting output-gap and inflation responses differ as well. [Table 14](#) reports the results.

Table 14: The consequences of a demand-signal error under a counterfactual policy rule

	Private assessments	
	Discretion	Counterfactual
<i>A. Responses on impact</i>		
Policy rate	0.16217	0.16217
Output gap	−0.29860	−0.17532
Inflation	−0.07694	−0.03530
Expected next-period output	−0.26138	−0.14435
Expected next-period inflation	−0.07048	−0.03140
<i>B. Market’s belief-revision of contemporaneous CB assessments</i>		
Natural output (actual: 0)	−0.16227	−0.00964
Demand (actual: 0.07014)	0.04416	0.05665
<i>C. Stationary contribution of demand-signal noise</i>		
Loss	0.34845	0.12852
Policy-rate variance	0.06000	0.06000

Notably, the output-gap contraction falls by 41.3%. The market’s estimate of the central bank’s contemporaneous natural-output assessment falls by only 0.00964 instead of 0.16227. Beyond the effects on impact, the stationary loss contribution of the demand-signal noise falls by 63.1%, while its contribution to the policy rate variance is, naturally, unchanged.

This exercise highlights the relevance of the belief-revision component of monetary transmission: changing the response to one assessment affects the effects of an error in another, even though the policy path is remains unchanged.

H Robustness Exercises

H.1 History-length sensitivity

To assess whether our finite-history approximation is reasonable, we expand the state by increasing the retained-history horizon from 8 to 12 lags. We then rebuild each equilibrium, and re-solve discretion, the simple commitment rule, and the full linear commitment. Consequently, the full linear rule expands from 35 to 51 coefficients because it can condition on four additional lags of each assessment and four additional policy-rate lags. Table 15 reports the resulting values of the key policy and belief-sensitivity statistics.

Table 15: History-Length Sensitivity

Statistic	Policy	$H = 8$	$H = 12$
<i>A. Commitment gains</i>			
Loss	Discretion	2.29760	2.25528
	Commitment: simple rule	0.72762	0.72773
	Full commitment	0.70776	0.70660
Full-rule gain / discretionary loss	Full commitment	69.20%	68.67%
Commitment: simple rule (share of F.C. gain)	Commitment: simple rule	98.75%	98.64%
<i>B. Expectations management</i>			
$\beta_{1:20}^P$	Discretion	-2.01174	-1.77076
	Commitment: simple rule	-1.26820	-1.26885
	Full commitment	-0.29370	-0.18435
Policy-surprise SD	Discretion	0.11439	0.11638
	Commitment: simple rule	0.03055	0.03054
	Full commitment	0.04796	0.04842
Fitted revision SD	Discretion	0.23013	0.20608
	Commitment: simple rule	0.03875	0.03875
	Full commitment	0.01409	0.00893

Note: The policy surprise and fitted revision are defined as in Table 5.

The full-commitment gain changes by 0.53 percentage points of discretionary loss, and the simple commitment rule share of that gain changes by 0.11 percentage points. The ranking of size of belief sensitivities is unchanged and the reduction under full commitment becomes stronger. Furthermore, policy-surprise standard deviations move by less than 2% in every policy regime. The simple commitment rule allocation is essentially unchanged. Some discretionary magnitudes are more sensitive: its output-gap variance falls by 10.6 percent and its policy-rate variance rises by 10.1%. Thus, the main results are robust to the history-length choice.

H.2 Rate-Level Stabilization Weight

In our baseline specification, we fix $\vartheta_i = 1$ instead of leaving it free for the estimation to determine its value. We do this because our model lacks conventional sources of endogenous persistence and thus even with the high estimated persistence of fundamentals, the observed persistence of the policy rate pushes ϑ_i to implausibly high levels. Concretely, when this parameter was left free, the estimation selected a value of ϑ_i above two. This is naturally problematic, since it entails the central bank puts more than twice as much weight on “rate stability” as on either inflation or output-gap stabilization. We interpret that result as ϑ_i absorbing missing endogenous persistence channels (i.e., misspecification) in the estimation, not as a proper identification of the preferences of the central bank.

To assess the sensitivity of the main results to this choice, we fix $\vartheta_i \in \{0.25, 0.5, 1.5, 2\}$ and re-estimate the remaining parameters using the same targets and weighting scheme as in the baseline. Then, we solve for the two commitment policy regimes in each estimated economy. Table 16 reports the results.

Table 16: Varying the Rate-Level Penalty ϑ_i

ϑ_i	Commitment gain	Simple rule share	R_β	R_i	Fitted revision SD ratio
0.25	70.85%	81.20%	0.002	1.492	0.004
0.50	70.55%	90.56%	0.055	1.076	0.059
1.00	69.20%	98.75%	0.146	0.419	0.061
1.50	67.91%	90.73%	0.875	0.150	0.131
2.00	75.25%	93.28%	4.830	0.192	0.926

Notes: At each value of ϑ_i , the remaining structural and information parameters are re-estimated using the baseline targets, weights, and bounds, and the discretionary equilibrium and both commitment problems are solved for that economy. The row for $\vartheta_i = 1$ is the baseline estimate. As defined in the main text, R_β is the ratio of absolute $\beta_{1,20}^P$ under commitment to discretion, and R_i is the analogous ratio of policy-surprise standard deviations. The fitted revision SD ratio is their product.

The commitment gain remains large across these alternative estimates, ranging from 67.91% to 75.25% of discretionary loss. The simple commitment rule captures between 81.20% and 98.75% of that gain. Thus, both the substantial cost of discretion and the fact that it is largely not a consequence of history dependence are robust to changes in the rate-level weight.

The associated changes in expectations sensitivity and policy-surprise volatility are a bit more subtle. At lower rate-level penalties, commitment sharply reduces the expected cumulative-price response to policy surprises while allowing greater surprise volatility than discretion. At $\vartheta_i = 2$, it instead combines an expectations response 4.83 times as large with a policy-surprise standard deviation only 19.2% as large. Fitted revision volatility remains

lower under full commitment in every case, highlighting the effectiveness of commitment in reducing the volatility of expectations. These are consistent with more responsive expectations being more useful when rate adjustments are costly (note, in particular, that R_β is increasing in ϑ_i) as discussed in the main text. At the same time, the main results remain robust. It should be stressed, though, that we are not simply changing ϑ_i but fully re-estimating the model to fit the data for each level of the rate-level penalty. Thus, the economies are not fully comparable.

H.3 Pure Inverse-Variance Weighting

To assess the sensitivity of our results to the choice of weighting, we re-estimate the model using only inverse-variance weights. Table 17 reports the resulting moment fit and policy statistics, comparing them with the baseline weighting.

Using only inverse-variance weights sharply improves the fit of the residual variance of inflation and output gap while materially weakening the fit on the market-expectations sensitivity (see underlined). At the main estimate, the macro block accounts for 98.55% of the estimation loss in the pure inverse-weighting specification, and the output-gap residual variance alone for 73.84% of it. This lowers the macro contribution from 27.256 to 2.325 but raises the contribution of the two Tealbook moments and the inflation-swaps reaction from .400 to 3.096. In particular, the fitted swaps reaction moves from -2.012 to -1.033 , while the data target is -2.449 .

The large commitment gains and the fact that a simple rule captures a majority of them survive: $G = 64.19\%$ and $S = 76.78\%$. The expectations-management orderings do not: full commitment raises absolute belief sensitivity, policy-surprise volatility, and fitted belief-revision volatility relative to discretion. The unconditional variances of (x, π, i) nevertheless fall from $(6.844, 3.334, 1.693)$ to $(1.368, 2.598, .285)$. Output-gap, inflation, and rate-level stabilization account for approximately $(72, 10, 18)$ percent of the total loss reduction of 3.810. Thus, under this calibration, commitment actually *raises* the strength of the expectations channel of monetary policy since its baseline discretion one is too weak. This is consistent with the discussion in the main text, since in this richer environment the sensitivity of expectations has its use and it might overwhelm its cost. However, note that this is associated with a much worse fit of the sensitivity of the expectations channel to the data.⁵³

⁵³Repeating the counterfactual in Appendix G.5.6 at this estimate also changes allocations while preserving the nominal-rate path following demand disturbances. The loss contribution of demand innovations falls by 9.8%, while that of demand-signal errors *rises* by 29.8%, with a net reduction in their combined loss contribution. This reflects that, while the specific ways in which commitment manages the interpretation of policy decisions may vary, the importance of this mechanism remains significant.

Table 17: Inverse-Variance Weights Only: A Comparison

Panel A: Moment fit

Moment	Data	Baseline weighting		Inverse-variance weights only	
		Model	Model–data	Model	Model–data
AR(1): inflation	.1321	.2440	.1119	.2379	.1058
AR(1): output gap	.8943	.9219	.0276	.9162	.0220
AR(1): interest rate	.9643	.9121	−.0522	.9462	−.0180
Residual variance: inflation	2.7672	1.3397	−1.4275	3.1455	.3783
Residual variance: output gap	1.1785	.2946	−.8839	1.0988	−.0798
Residual variance: interest rate	.1650	.2031	.0381	.1771	.0121
Forecast-revision variance share: inflation	.2414	.2680	.0266	.2892	.0478
Forecast-revision variance share: real GDP growth	.5876	.6296	.0420	.6911	.1035
Market-expectations sensitivity $\beta_{1:20}^P$	−2.4486	−2.0117	.4369	−1.0329	<u>1.4157</u>

Panel B: Sensitivity of expectations across policy regimes

Specification	Policy	Loss	$\beta_{1:20}^P$	Policy-surprise SD	Fitted revision SD
Baseline weighting	Discretion	2.29760	−2.01174	.11439	.23013
	Simple rule	.72762	−1.26820	.03055	.03875
	Full commitment	.70776	−.29370	.04796	.01409
Inverse-variance weights only	Discretion	5.93546	−1.03294	.11350	.11724
	Simple rule	3.01006	−1.89547	.16979	.32183
	Full commitment	2.12533	−1.34918	.17133	.23115

Panel C: Headline results

Specification	G	S	R_β	R_i	Fitted revision SD ratio
Baseline weighting	69.20%	98.75%	.146	.419	.061
Inverse-variance weights only	64.19%	76.78%	1.306	1.510	1.972

Notes: G is the full-commitment gain as a percentage of discretionary loss and S is the share of that gain captured by the simple commitment rule. Policy-surprise and fitted-revision SDs are defined as in [Table 5](#). The three ratios compare full commitment with discretion.