

Forward Guidance for Skeptical Markets[†]

Vicente Jimenez-Gimpel

MIT

October 5, 2026

Abstract

This paper develops a theory of forward guidance that captures two key features of central bank communication: announcements are not fully state-contingent, and the private sector is wary of policy mistakes. The central bank tailors announcements along two central dimensions: (i) vagueness—how tightly the announcement constrains future policy; and (ii) data dependence—how aggressively future policy responds to a public signal of the economy. Three main lessons emerge. First, uncertainty affects communication differently depending on its source: sufficiently high uncertainty about either demand or cost-push shocks calls for vaguer guidance, but only demand-side uncertainty increases the optimal degree of data dependence. Second, strict rules are costly: fully precise guidance entails sizable stabilization losses, especially when the Phillips curve is steep. Third, because of the stabilizing role of announcements, optimal forward guidance differs sharply from simply communicating the best forecast of the future optimal policy rate (Delphic guidance).

Keywords: Monetary policy, interest rates, forward guidance, central bank communication, data dependence

JEL Codes: D84, E40, E52, E58, E61

[†]E-mail: jimenezv@mit.edu. First posted version: March 2, 2026. I am deeply grateful to Ricardo J. Caballero and Christian K. Wolf for their continued support and guidance. I thank Tomás Caravello for many insightful suggestions. I also thank Martin Beraja, Luigi Bocola, Luis Felipe Céspedes, Andrés Drenik, Joao Guerreiro, Stephen Morris, Aleksei Oskolkov, Karthik Sastry, Luminita Stevens, Jean Tirole, Javier Turen, Rafael Veiel, Iván Werning, and participants at the MIT Macro Lunch and the VI Alumni Workshop at UC Chile for helpful comments.

1 Introduction

Central banks devote enormous effort to communicating properly, and for good reason: Gurkaynak et al. (2005) find that news about the future policy path—closely associated with FOMC statements—accounts for up to 90% of the explained variation in long term Treasury yields around FOMC announcements. Managing expectations is widely regarded as a crucial element of successful monetary policy (see, e.g., Woodford (2005), Yellen (2013), and Bernanke (2020)). Kevin Warsh’s appointment as Chair of the Federal Reserve has brought the role and design of forward guidance to the forefront of the policy debate. In his August 2026 Jackson Hole address, the Warsh argued that its routine use has “overstayed its welcome”, warning that commitments about future interest rates can restrict policymakers’ flexibility (Warsh, 2026). He also questions whether a simple reaction function can capture the changing considerations relevant to policy decisions. Governor Christopher Waller has emphasized the value of guidance in shaping expectations and financial conditions, while acknowledging the costs of excessive rigidity (Waller, 2026).

Yet standard macroeconomic frameworks offer little guidance on this trade-off. Commitment to a fully state-contingent rule, perfectly understood by the private sector (Woodford, 2003), can anchor expectations while allowing policy to respond to all relevant information. Communication design then reduces to the design of optimal policy rules. In practice, announcements condition on a narrow subset of observables, creating the flexibility concern at the core of the policy debate surrounding forward guidance. Leaving discretion preserves flexibility but gives markets less assurance about future policy. How tightly should such guidance constrain future policy? And how strongly should those constraints respond to imperfect incoming data?

I study these questions in a model where central bank’s announcements can condition only on a subset of the information available when policy is implemented. The central bank commits to a signal-contingent set of interest rates, retaining discretion to choose within that set as additional information becomes available. The private sector is *skeptical* regarding future policy: its expectations reflect the least favorable rate consistent with the announced set. Thus, the central bank weighs the benefits of anchoring expectations against the costs of constraining its response to information that cannot be encoded in the announcement. I characterize optimal announcements along three design margins: vagueness—how much discretion the central bank retains; data dependence—how strongly the announced range of rates responds to public signals; and the center of that range.

I then consider inflation-contingent rate announcements, in line with standard monetary policy rules. Three main lessons emerge. First, uncertainty affects optimal communication

differently depending on its source: sufficiently high uncertainty about either demand or cost-push shocks calls for vaguer guidance, but only demand-side uncertainty calls for guidance that is more strongly anchored to inflation. Second, guidance that leaves no room for discretion entails sizable stabilization costs. Third, optimal guidance differs sharply from the best forecast of future policy because announcements themselves shape expectations and current economic outcomes.

Framework. The paper builds on a standard New Keynesian model without capital. The central bank commits to a signal-contingent *set* of interest rates: a range of policy rates that the central bank can choose from for each realization of a signal of the state of the economy.¹

Crucially, announcements cannot be made fully state-contingent. At the time the central bank sets its interest rate, it has access to additional information encoded in a second signal; thus, the implemented rate can depend on information beyond what is feasible to encode in the announcement. This captures a basic fact: policy-rate decisions depend on the assessment of a very wide range of variables (e.g., different measures of inflation, quit rates, wage growth, job openings), but forward guidance typically conditions on a much narrower subset (e.g., the unemployment rate or inflation) and in a way that does not uniquely pin down future policy.² Fully state-contingent guidance would require conditioning on a very high-dimensional set of variables, including elements that are difficult to verify or communicate—policymakers’ judgment, “soft” information, outputs of large-scale internal forecasting models—causing a mismatch between the coarse, publicly communicable contingencies embedded in guidance and the much richer set of considerations that actual rate decisions depend on.

The private sector shares the central bank’s prior beliefs over fundamentals, but is *ambiguity averse* regarding future policy rates: its policy expectations correspond to the least favorable rate consistent with the announced set. This is motivated by the well-documented disagreements between the private sector and the central bank about appropriate future policy rates (see, e.g., Caballero and Simsek (2022); Amodeo (2025); Sastry (2026)), and by the

¹Treating announcements as binding provides a stylized representation of the constraints that prior guidance places on subsequent policy. This echoes the concern about flexibility emphasized by Warsh (2026) and Waller (2026), discussed above. See also Barkin (2021) and Romer and Romer (2024) on the role of guidance in the Fed’s delayed response to the 2021–2022 inflation surge.

²Yellen (2013) describes the evolution of FOMC forward guidance from unconditional, calendar-based language (e.g., “exceptionally low levels (...) for an extended period” or “at least through mid-2013”) to thresholds tied to unemployment and inflation, emphasizing that even in the latter case, forward guidance was not precise: “...the FOMC’s new forward guidance offers considerable insight into the Committee’s likely reaction function, but I should note that the guidance it provides is not complete. For example, the Committee has not specified exactly how it intends to vary the federal funds rate after liftoff from the effective lower bound, although it has stated that *when the Committee decides to begin to remove policy accommodation, it will take a balanced approach.*”

absence of an agreed-upon probability distribution over future policy, even accounting for the central bank’s announcements. Ambiguity aversion captures the private sector’s aversion to policy uncertainty—a widely acknowledged phenomenon (e.g., [Baker et al. \(2016\)](#); [Bauer et al. \(2022\)](#))—while remaining agnostic about whether the underlying worry stems from differences in data interpretation, in assessments of policy effectiveness, or in preferences. Crucially, it implies that private-sector expectations depend on the endpoints of the announced set, not just its center, a property that is necessary for the commitment-flexibility trade-off at the core of the paper to be operative (see [Section 2.4](#) for a detailed discussion).

The rational expectations benchmark. In the standard rational-expectations benchmark with common information, the design of forward guidance with commitment reduces to the design of optimal policy rules and hence cannot speak to several considerations that are at the heart of central bank communication. The central bank “announces” a mapping from each period’s relevant state vector to that period’s interest rate and then simply implements it (which the private sector perfectly anticipates). Consequently, the trade-off between vagueness and precision is absent, and the trade-off governing optimal data dependence is far less rich: because the rule can condition on the full state, the variables guidance is anchored to are the same as those available at the time of rate-setting—eliminating the central tension studied here, namely that announcements must be anchored to a strict subset of the information the central bank will ultimately act on. Furthermore, the informational role associated with “Delphic” guidance ([Campbell et al., 2012](#)) is absent. With asymmetric information, communication can instead reveal information about the inputs of a known policy rule, complementing the commitment to that rule. Here, the announcement is itself the commitment: it shapes expectations by defining the central bank’s scope for future discretion.

Endogenous signals. In practice, forward guidance is nearly always anchored to endogenous macroeconomic variables, most prominently inflation and unemployment. A variant of the model captures this by letting announcements condition on an endogenous variable—contemporaneous inflation—instead of a signal with an exogenous structure. Because inflation depends on expectations of future rates, and those expectations are anchored by an inflation-contingent announcement, a feedback loop arises: the announcement shapes the very signal on which it conditions. This feedback can be stabilizing or destabilizing depending on the extent of data dependence. When data dependence is positive, extreme worst-case beliefs move inflation in a direction that moderates the perceived mis-centering of the band, making equilibrium unique and self-correcting. A further consequence of the stabilizing feedback is that fully vague announcements—placing no effective constraint on

future policy—may be optimal: unlike in the exogenous-signal model, where widening the band indefinitely causes expectations to diverge, the feedback loop keeps worst-case expectations bounded even as the band grows without limit. The central bank’s choice of data dependence thus determines the equilibrium regime, adding a qualitatively new dimension to the design of optimal communication. In addition, changing the width of the band can also shift its midpoint through its effect on equilibrium inflation, even with the current rate, slope, and intercept held fixed. Because data dependence shapes how strongly the feedback loop dampens or amplifies this re-centering, the interaction between vagueness and data dependence is richer than in the exogenous-signal model.

Main results. The endogenous signals framework yields three main lessons.

Lesson 1: Uncertainty affects optimal communication asymmetrically. When demand uncertainty is high, the central bank optimally issues announcements that are both highly vague and strongly data dependent. By contrast, in environments with high cost-push uncertainty, optimal announcements remain highly vague but exhibit a limited degree of data dependence. Promising higher future rates in response to demand-driven inflation helps restrain both excessive demand and price pressures. Stabilizing cost-push inflation, however, requires pushing output below its natural level, so the central bank optimally accommodates part of the shock. Since guidance conditions on inflation instead of on its sources, a steep response that is useful against demand disturbances also promises substantial tightening when partial accommodation would be preferable. Greater cost-push uncertainty makes this conflict more important, leading the central bank to reduce the responsiveness of its guidance to inflation.

Lesson 2: The cost of (strict) rules is sizable. Constraining the central bank to fully precise announcements—a single rate for each signal realization—entails sizable additional expected losses relative to optimally vague guidance (7.13% in the baseline calibration). These losses increase with the steepness of the Phillips curve over most of the range considered and exhibit a U-shaped pattern as the common weight on inflation stabilization varies. The value of discretion comes from the impossibility of fully state-contingent communication, echoing policymakers’ arguments pushing back against the profession’s strong post-Kydland and Prescott (1977) push toward policy rules that leave no room for discretion (see, e.g., Yellen (2015)).

Lesson 3: Optimal guidance is not a forecast. A forecast aims to predict which policy rate will be appropriate given future information. Forward guidance, however, takes into account how the announcement itself shapes expectations and thus contemporaneous macroeconomic outcomes. To capture this distinction, I compare optimal data dependence with the best

linear forecast of the future unconstrained optimal rate based on inflation (i.e., the *Delphic* guidance benchmark). Quantitatively, optimal guidance responds much more strongly to inflation than this forecasting benchmark owing to its stabilization role. The gap widens with demand uncertainty and narrows with cost-push uncertainty. This formalizes a familiar view in policy circles: proper forward guidance is not simply about disclosing information but a policy instrument in its own right.

The baseline calibration combines the structural New Keynesian benchmark in [Woodford \(2003\)](#), based on [Rotemberg and Woodford \(1997\)](#), with natural-rate and cost-push processes drawn from the same estimated specification in [Cole and Martínez-García \(2023\)](#). The mean natural rate is anchored to the New York Fed’s December 2024 DSGE forecast.

Discussion and related literature. The main macroeconomic literature on forward guidance studies its stabilization power in rational-expectations environments when contemporaneous policy rates are constrained, most often by the zero lower bound ([Eggertsson and Woodford, 2003](#); [Werning, 2012](#); [McKay et al., 2016](#); [Del Negro et al., 2023](#)). A separate line of work applies cheap-talk and information-design frameworks to central bank communication (see, e.g., [Stein \(1989\)](#); [Moscarini \(2007\)](#); [Bassetto \(2019\)](#); [Cieslak et al. \(2020\)](#); [Gáti \(2023\)](#)). However, the strategic-communication frameworks in this literature are typically too stylized to speak to the detailed design of forward guidance—for instance, how strongly guidance should condition on specific observables. Moreover, since these frameworks are not embedded in full-fledged macroeconomic models, relevant variables are typically introduced in a reduced-form way, which limits the scope for quantitative assessments of the mechanisms at play.

On the vagueness margin, a closely related paper is [Lakdawala et al. \(2026\)](#). It studies how precisely central banks should communicate about future policy rates when more precise communication allows agents to make better-informed decisions but also amplifies financial-market responses to policy news. Here, the trade-off arises from the interaction between commitment and market skepticism: more precision anchors market expectations but limits the central bank’s flexibility to respond to information that cannot be encoded in the announcement. Moreover, the central bank jointly chooses how strongly guidance responds to observable signals—its degree of data dependence.

This commitment-based formulation also connects the paper to the delegation literature, which highlights the trade-off between commitment and flexibility when the agent has private information and either has preferences that differ from the principal’s or is time inconsistent under a dual-self interpretation (see, e.g., [Holmstrom \(1984\)](#) for the seminal reference and [Amador et al. \(2006\)](#) for a more modern and general treatment). Within that strand of lit-

erature, the closest references are [Athey et al. \(2005\)](#) and [Kocherlakota \(2016\)](#). [Athey et al. \(2005\)](#) use a dynamic mechanism design approach to find the optimal degree of discretion society should assign the central bank given its inflation bias and extent of private information, finding it can be implemented via an inflation cap. [Kocherlakota \(2016\)](#) revisits the same question using [Holmstrom \(1984\)](#)’s static framework and a reduced-form model of inflation determination, showing that when the central bank has information about the economy that cannot be encoded in a rule—what he terms “nonrulable” information—granting discretion is socially desirable provided the inflation bias is small enough relative to the variance of nonrulable shocks.

These two papers share the core trade-off that drives the vagueness margin in the present paper: the flexibility to respond to information that cannot be encoded in an ex-ante rule is weighed against the cost of discretion. Furthermore, the finding that optimal announcements take the form of intervals ([Proposition 2](#)) is the interest-rate analog of [Athey et al. \(2005\)](#)’s bounded-discretion result; in both cases, the optimal constraint is a region whose width reflects this trade-off. The qualitative message of Lesson 2—that strict rules are costly—echoes [Kocherlakota \(2016\)](#)’s central conclusion. The source of this trade-off differs, however: in those papers, the cost of discretion is caused by an inflation bias—arising from an explicit preference wedge in [Kocherlakota \(2016\)](#) and from dynamic inconsistency in [Athey et al. \(2005\)](#)—so the optimal degree of discretion hinges on its magnitude ([Kocherlakota \(2016\)](#)), for instance, argues that the U.S. inflation bias is negligibly small, which in his framework implies full discretion is optimal). Here, the value of partially tying the central bank’s hands is instead grounded in the private sector’s documented concern about policy mistakes, and hence does not require taking a stand on the size of any inflation bias. Beyond the source of the trade-off, neither paper studies how rules should respond to signals on which policy can be conditioned ex ante—data dependence—and the endogenous-signals extension introduces considerations entirely absent from both frameworks.

This project is also part of a growing literature studying monetary policy with agents that lack model-consistent expectations or, more generally, are boundedly rational ([García-Schmidt and Woodford, 2019](#); [Farhi and Werning, 2019](#); [Angeletos and Huo, 2021](#); [Angeletos and Sastry, 2021](#); [Caballero and Simsek, 2022](#)). This paper studies a novel deviation from rational expectations, namely, ambiguity aversion over future policy rates coupled with only partially state-contingent communication. The former ingredient makes the paper conceptually related to the literature that studies the implications of ambiguity aversion in macroeconomics ([Ilut and Schneider, 2014](#); [Backus et al., 2015](#); [Bianchi et al., 2018](#); [Masolo and Monti, 2021](#)). In most of this literature, ambiguity aversion is about economic fundamentals, not policy variables. Thus, the implications, even for the optimal design of monetary policy

(and hence, forward guidance), are quite different. To the best of my knowledge, the only exceptions are [Michelacci and Paciello \(2020\)](#) and [Masolo and Monti \(2021\)](#), in which agents hold ambiguity-averse preferences over the commitment ability of the monetary authority and over an additive “disturbance” term in an interest rate rule, respectively. In [Michelacci and Paciello \(2020\)](#), the analysis is purely positive: the ambiguity set (i.e., the set of beliefs entertained by the private sector) is taken as given by the central bank and there is no fundamental uncertainty, so none of the trade-offs that shape the optimal design of announcements studied in this paper are present. Similarly, while [Masolo and Monti \(2021\)](#) do consider the optimal policy problem, in their model the private sector’s ambiguity set is also exogenous, so optimal policy in their framework is best read as the choice of a rule that is robust to a given set of belief distortions about policy rates, as opposed to the joint design of policy and communication.

Roadmap. The remainder of the paper is organized as follows. [Section 2](#) presents the baseline model with exogenous signals. [Section 3](#) characterizes and discusses the optimal central bank rate and communication policy. [Section 4](#) develops the variant of the model with endogenous signals. [Section 5](#) presents the numerical results, including the baseline calibration and comparative statics that support the three main lessons. Finally, [Section 6](#) concludes. [Appendix A](#) contains the formal definitions of the model. The proofs of the results in the main text are in [Appendix B](#). [Appendix C](#) provides second-order conditions for optimal announcements in the exogenous-signal model. Proofs of additional technical lemmas can be found in [Appendix D](#).

2 Model

This section lays out the model. First, I describe the physical environment, which corresponds to the standard New Keynesian model without capital. Second, I define what central bank announcements consist of and how they shape private-sector beliefs. Then, I detail the information structure, discussing both agents’ beliefs about exogenous shocks and how the private sector forms expectations about future policy rates following an announcement. Finally, I briefly discuss two key assumptions: how the private sector forms beliefs about future policy, and the constraints on the form of the central bank’s announcements.

2.1 Physical Environment

Consider a standard New Keynesian model without capital in which the private sector (i.e., households and firms) does not hold model-consistent expectations about future policy rates. Time is discrete and indexed by $t \in \{0, 1, \dots\}$. For simplicity, I assume there is no uncertainty from period $t = 2$ onward, so $(r_t^n, u_t) = (\bar{r}^n, 0)$ for $t \geq 2$, where $\bar{r}^n$ is the steady-state value of the natural interest rate. Thus, the linearized versions of the equations describing the private sector's behavior are the usual Dynamic IS and New Keynesian Phillips Curve³

$$\begin{aligned} y_t &= \tilde{\mathbb{E}}_t[y_{t+1}|\mathcal{A}_t] - \sigma(i_t - \tilde{\mathbb{E}}_t[\pi_{t+1}|\mathcal{A}_t] - r_t^n), \\ \pi_t &= \kappa y_t + \beta \tilde{\mathbb{E}}_t[\pi_{t+1}|\mathcal{A}_t] + u_t, \end{aligned}$$

where y_t is the (log) output gap, π_t is (log) gross inflation, σ is the EIS, i_t is the log of the gross yield of a one-period risk-free bond, r_t^n is the natural rate (i.e., the real rate of interest required each period in order to keep output equal to its natural counterpart at all times)⁴, u_t represents cost-push shocks and $\tilde{\mathbb{E}}_t[\cdot|\mathcal{A}_t]$ stands for the private sector's expectation operator after observing a central bank announcement $\mathcal{A}_t$ (more on this later). The above assumes beliefs are the same for households and producers. This can be justified by noting that ultimately the owners of firms are households.⁵

The central bank minimizes a standard quadratic stabilization loss in the output gap and inflation. This criterion is motivated by welfare approximations in standard New Keynesian models (see, e.g., [Woodford \(2003\)](#)). Besides its usual contemporaneous policy-rate decision, the central bank is allowed to make announcements regarding its future policy rates that shape private sector beliefs as discussed below. Hence, the central bank at $t = 0$ solves

$$\begin{aligned} \tilde{V}_0 &= \min_{y_0, \pi_0, i_0, \mathcal{A}_1} \mathbb{E}_0[y_0^2 + \lambda_{cb}\pi_0^2 + \beta V(\mathcal{A}_1)] \\ \text{s.t. } y_0 &= \tilde{\mathbb{E}}_0[y_1|\mathcal{A}_1] - \sigma(i_0 - \tilde{\mathbb{E}}_0[\pi_1|\mathcal{A}_1] - r_0^n) \end{aligned}$$

³Because there is no uncertainty from period $t = 2$ onward and this is common knowledge, no additional terms involving current expectations of future variables appear.

⁴In the standard New Keynesian model, $r_t^n := \rho_t + \sigma^{-1}\mathbb{E}_t\Delta y_{t+1}^n$, where $\rho_t = -\log \beta_t$ and y_t^n is the (log) natural output. As the current model is not concerned with the particular sources of variation in r_t^n , it will be assumed it fluctuates according to an exogenous law of motion. This should be interpreted as a reduced-form representation of the flexible-price real rate in a richer environment—potentially reflecting time variation in preferences, risk premia, and other wedges/frictions and not as the object mechanically implied by β and trend growth in the stripped-down standard model.

⁵The same assumption has been used before in similar contexts (see, e.g., [Ilut and Saijo \(2021\)](#)). If this is not the case, there is an extra layer of complexity as then firms and consumers would hold different expectations which would generate additional terms on the NKPC involving higher-order beliefs. While interesting on its own, this goes beyond the scope of this project.

$$\pi_0 = \kappa y_0 + \beta \tilde{\mathbb{E}}_0[\pi_1 | \mathcal{A}_1] + u_0,$$

where $\lambda_{cb} > 0$ captures the importance of inflation stability relative to output stability for the central bank and

$$\begin{aligned} V(\mathcal{A}_t) = & \min_{y_t, \pi_t, i_t, \mathcal{A}_{t+1}} \mathbb{E}_t \left[y_t^2 + \lambda_{cb} \pi_t^2 + \beta V(\mathcal{A}_{t+1}) \right] \\ \text{s.t. } & y_t = \tilde{\mathbb{E}}_t[y_{t+1} | \mathcal{A}_{t+1}] - \sigma(i_t - \tilde{\mathbb{E}}_t[\pi_{t+1} | \mathcal{A}_{t+1}] - r_t^n) \\ & \pi_t = \kappa y_t + \beta \tilde{\mathbb{E}}_t[\pi_{t+1} | \mathcal{A}_{t+1}] + u_t \\ & i_t \in \mathcal{A}_t \end{aligned}$$

2.2 Central Bank Announcements

A central bank announcement is a signal-contingent rule, $\mathcal{A}_t = \{A_t(\tilde{s}_t)\}_{\tilde{s}_t}$, which assigns to each possible signal realization $\tilde{s}_t$ a non-empty closed set $A_t \subseteq \mathbb{R}$ of policy rates.⁶ Announcements are binding, i.e., at period t and given the realization $\tilde{s}_t$, the central bank's policy rate has to satisfy $i_t \in A_t(\tilde{s}_t)$.

The constraint that announcements can only be contingent on $\tilde{s}_t$ is motivated by the fact that the central bank's decisions usually are informed by multiple data sources, some of which are salient and verifiable (e.g., jobs reports) while others are not (e.g., DSGE-based forecasts, FOMC members' "hunches", private data). In the model, $\tilde{s}_1$ is taken to stand for the aggregation of the former and $\hat{s}_t$ for the aggregation of the latter.

The central bank can tailor its announcements along two central dimensions: the width of the set $A_t(\tilde{s}_t)$ for a given realization of $\tilde{s}_t$ and the extent of its dependence on the signal $\tilde{s}_t$. The former captures how *vague* announcements are (i.e., to what extent they allow for ex-post discretion), while the latter captures how *data dependent* they are (i.e., how responsive to data releases they are).

Feasible announcements are restricted to satisfy the following condition,

Definition 1. An announcement $\mathcal{A} = \{A(\tilde{s}_t)\}_{\tilde{s}_t}$ is **equivariant** if there exists a scalar $\alpha \in \mathbb{R}$ such that

$$A(x + \delta) = A(x) + \alpha \delta \quad \text{for all } x, \delta \in \mathbb{R},$$

This translation-invariance property (shifting the signal by δ translates the entire announced set by $\alpha\delta$) essentially requires the announcement to depend linearly on the signal $\tilde{s}_t$. While not without loss of generality (more on this in [Subsection 2.4](#) below), this is a reasonable assumption given the model's linear-Gaussian structure. It also echoes the focus

⁶See [Definition 2](#) in the Appendix for the formal details.

on simple rules (e.g., *Taylor rules*) in the monetary economics literature and allows for a natural measure of the “extent of data dependence”: the size of α .⁷

2.3 Beliefs

This subsection describes the beliefs of the central bank and the private sector. The key non-standard ingredient is that the private sector is skeptical about future policy rate choices, and central bank announcements act as a partial commitment device to limit the scope of this skepticism. Formally, this takes the form of ambiguity-averse beliefs over the set of feasible policy rates implied by the announcement.

Central bank. The central bank does not observe contemporaneous shock realizations. Its prior beliefs regarding $t = 0$ shocks are given by

$$\begin{pmatrix} r_0^n \\ u_0 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \bar{r}^n \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_u^2 \end{pmatrix} \right).$$

This is the only information the central bank has regarding exogenous shocks when choosing its announcements and policy decision at $t = 0$. Further, shocks are persistent in the relevant periods. In particular, they follow a mean-reverting AR(1),

$$\begin{aligned} r_1^n &= (1 - \rho_r)\bar{r}^n + \rho_r r_0^n + \epsilon_r \\ u_1 &= \rho_u u_0 + \epsilon_u, \end{aligned}$$

where (ϵ_r, ϵ_u) is a Gaussian vector, independent of (r_0^n, u_0) , with independent components of mean zero and variances $\sigma_{\epsilon_r}^2$ and $\sigma_{\epsilon_u}^2$, respectively. When setting policy rates at $t = 1$, the central bank does not know the realization of the shocks r_1^n and u_1 either, but it does observe the lagged shock values (r_0^n, u_0) and two signals of the contemporaneous shocks,

$$\begin{aligned} \tilde{s}_1 &= \tilde{\omega}_r r_1^n + \tilde{\omega}_u u_1 + \tilde{\epsilon}, \\ \hat{s}_1 &= \hat{\omega}_r r_1^n + \hat{\omega}_u u_1 + \hat{\epsilon}, \end{aligned}$$

where $(\tilde{\epsilon}, \hat{\epsilon})$ is a Gaussian vector, independent of $(r_0^n, u_0, \epsilon_r, \epsilon_u)$, with independent components of mean zero and variances $\tilde{\sigma}^2$ and $\hat{\sigma}^2$, respectively. The positive coefficients $\tilde{\omega}_r, \tilde{\omega}_u, \hat{\omega}_r, \hat{\omega}_u$ are exogenous weights on demand and supply shocks. As mentioned above, the public signal

⁷It's straightforward to verify that if a bounded announcement satisfies equivariance for some α , there is no $\alpha' \neq \alpha$ such that the announcement is also equivariant with respect to it. Thus, the extent of data dependence α of a given announcement is well defined.

$\tilde{s}_1$ is observed by the private sector too while the private signal $\hat{s}_1$ is not. The structure of the signals is common knowledge. Regarding endogenous variables, the central bank is assumed to have rational expectations (i.e., model-consistent beliefs). In particular, it perfectly understands the behavior and belief formation of the private sector.

Private sector. The private sector holds the same prior beliefs as the central bank regarding exogenous shocks, but it does observe their contemporaneous realizations. This is a natural assumption given shocks are assumed to be about variables relevant to their decision problems. In fact, one can always define the actual shocks as the component of them that is known to the private sector. This information asymmetry does matter, as will be clear in what follows (see the discussion below [Proposition 1](#)).

Regarding the central bank's future behavior (i.e., the policy rate i_1), the private sector uses a multiple-priors criterion in the spirit of [Gilboa and Schmeidler \(1989\)](#), with a set of priors shaped by central bank announcements. This formulation follows the distinction in [Gajdos et al. \(2008\)](#) between (objectively) imprecise information and the decision maker's attitude toward that imprecision. Here, the announcement, determines a set of admissible probability distributions. The private sector evaluates future policy using worst-case expected losses over this set. This is the analogue of the identity selection $\varphi(P) = P$ in their framework.

The private sector evaluates future policy using a quadratic stabilization loss, mirroring the functional form of the central bank's objective.⁸

Given an announcement $\mathcal{A}_1$, the private sector considers all joint distributions over $(i_1, r_1^n, u_1, \tilde{s}_1)$ such that (i) the marginal law of $(r_1^n, u_1, \tilde{s}_1)$ equals the fundamental Gaussian conditional distribution given the realization of (r_0^n, u_0) , (ii) the conditional distribution of i_1 given $(r_0^n, u_0, \tilde{s}_1, r_1^n, u_1)$ depends only on $(r_0^n, u_0, \tilde{s}_1)$, and (iii) for each $\tilde{s}_1$, the support of i_1 lies in $A_1(\tilde{s}_1)$.

The first condition is a consistency one: the market's joint beliefs are aligned with its knowledge of the distribution of exogenous variables. The second condition is a simplifying assumption about how the private sector assesses future policy. Concretely, it considers the choices allowed by the announcement using the information shared with the central bank, without anticipating how the central bank will respond to its additional private signal. Finally, the third condition follows from the market understanding that the central bank is

⁸This specification is motivated by curvature properties of household welfare in the underlying (nonlinear) New Keynesian model. Concretely, [Benigno and Woodford \(2005\)](#) establishes conditions under which feasible policy perturbations have negative welfare curvature (which hold around an efficient steady state as is the case in this model). For an announced rate interval contained in a neighborhood where welfare is concave in the policy rate, its minimum is attained at an endpoint. A quadratic stabilization loss captures this feature in a tractable form.

bound by its announcements. Denoting this admissible class by $\mathcal{P}(s_0)$, the private sector’s worst-case beliefs solve⁹

$$\begin{aligned} \sup_{p \in \mathcal{P}(s_0)} \mathbb{E}_p \left[y_1^2 + \lambda_{ps} \pi_1^2 \right] & \tag{1} \\ \text{s.t. } y_1 &= -\sigma(i_1 - r_1^n) \\ \pi_1 &= \kappa y_1 + u_1 \\ \tilde{s}_1 &= \tilde{\omega}_r r_1^n + \tilde{\omega}_u u_1 + \tilde{\epsilon}_1, \end{aligned}$$

where λ_{ps} is the weight the private sector puts on inflation. Note that the private sector anticipates $y_t = \pi_t = 0$ for all $t \geq 2$: it has perfect foresight about the future state of the economy, and the central bank can achieve full stabilization because cost-push shocks are set to their mean, zero.

2.4 Discussion of Key Assumptions

This subsection discusses two assumptions that merit further examination: the formation of private-sector beliefs about future policy and the equivariance restriction on announcements.

On private-sector beliefs. The assumption that the private sector believes future policy maximizes the stabilization loss within the space of feasible rates may seem extreme at first. Indeed, it is the opposite of what is usually assumed: with rational expectations and a standard “benevolent” central bank, the private sector *knows* the central bank will pick rates that minimize the stabilization loss. In practice, disagreements between the private sector and the central bank about appropriate future policy rates are well documented (see, e.g., [Caballero and Simsek \(2022\)](#), [Amodeo \(2025\)](#), [Sastry \(2026\)](#)).¹⁰ This motivates modeling the private sector as entertaining the possibility of adverse policy choices, whether due to different data interpretations, different assessments of the effectiveness of policy, or plain preference differences. The ambiguity aversion formulation is a natural way of representing this for three reasons. First, it reflects that there is no agreed-upon (“objective”) probability distribution of future policy, even accounting for the central bank’s announcements. Second, it captures the private sector’s aversion to policy uncertainty, a widely acknowledged phenomenon (see,

⁹See [Appendix A](#) for a precise formal definition of $\mathcal{P}$.

¹⁰The idea that the Fed can get policy wrong also appears recurrently in the press (e.g., *The Fed Is Behind the Curve on Cutting Interest Rates*, *The Wall Street Journal*, June 13, 2025).

e.g., [Baker et al. \(2016\)](#); [Bauer et al. \(2022\)](#)).¹¹ Third, it abstracts from the precise source of concerns about future policy choices, while remaining consistent with familiar interpretations (for instance, as an approximation to expected-utility behavior under uncertainty about the central bank’s preferences λ_{cb} and high risk aversion).

For the commitment-flexibility trade-off to operate, private-sector expectations following an announcement must depend on the announcement’s width. To see why, suppose agents use only the midpoint of the announced interval to forecast future policy. Then the central bank can make the interval arbitrarily wide—preserving full ex-post flexibility—while keeping expectations pinned down by the midpoint, so the commitment–flexibility trade-off collapses.

Ambiguity aversion is a tractable way to generate this endpoint-sensitive behavior. A more standard alternative would be for the private sector to have expected-utility preferences with announcements consisting of families of *distributions* of policy rates (indexed by realizations of the conditionable signal). In that case, however, sensitivity to endpoint changes that preserve conditional means requires moving beyond a linearized setting: in models with linearized private-sector behavior only conditional means matter, so endpoint adjustments that leave the mean unchanged do not affect the behavioral consequences of a given announcement. Furthermore, this would make the assumption that announcements are binding commitments hard to operationalize, since ex post compliance with a conditional distribution is not meaningfully verifiable from a single realization. Moreover, requiring central bank announcements to specify full conditional probability distributions is far from observed communication practices (see [footnote 2](#)).

On equivariance. As stated above, the space of feasible central bank announcements is restricted to those that satisfy *equivariance* ([Definition 1](#)). This is a natural assumption in the context of policy rules: equivariance formalizes the idea that the announcement reacts to the conditionable signal through a simple linear index, so that shifting the signal by δ shifts the entire announced set by a constant amount. Such a translation-invariance property is the most direct counterpart of the familiar focus on linear reaction functions (e.g., Taylor rules). Furthermore, it allows for a transparent measure of “data dependence” (the slope) and a clean notion of “vagueness” (the width of the set). At the same time, the restriction is not innocuous.

To understand what features of communication are ruled out by studying equivariant announcements only, it is helpful to separate two dimensions of an announcement. In the next section, it is shown that optimal announcements are intervals (and this holds regardless

¹¹Indeed, a formal version of this is a key axiom in the seminal representation theorem for max-min preferences, [Gilboa and Schmeidler \(1989\)](#).

of whether feasible announcements are restricted to be equivariant or not). Thus, for the purposes of understanding the role of equivariance, only interval announcements will be considered. Any interval announcement can be written as

$$A(\tilde{s}_1) = [c(\tilde{s}_1) - \gamma(\tilde{s}_1), c(\tilde{s}_1) + \gamma(\tilde{s}_1)],$$

where $c(\cdot)$ is the center of the interval and $\gamma(\cdot) \geq 0$ is its half-width. Equivariance imposes two conditions: (i) the center is affine, $c(\tilde{s}_1) = \alpha\tilde{s}_1 + b$, and (ii) the width is constant, $\gamma(\tilde{s}_1) \equiv \gamma$. Absent this restriction, the central bank could make the responsiveness of the feasible set vary non-linearly with the signal realization and allocate ex-post flexibility unevenly across signal realizations by letting γ vary with $\tilde{s}_1$.

Why might either of these additional degrees of freedom be valuable for the central bank? The key is that, although the central bank chooses its announcement before observing the first period shocks $s_0 \equiv (r_0^n, u_0)$, the private sector *does* observe s_0 when forming its policy expectations. As a result, even though the functions $c(\cdot)$ and $\gamma(\cdot)$ do not explicitly depend on s_0 , they shape private-sector expectations *as a function of* s_0 via the conditional distribution of $\tilde{s}_1$ and s_1 (next period's shocks) given s_0 . In particular, s_0 shifts the conditional mean (and, more generally, the conditional likelihood) of the public signal $\tilde{s}_1$. Consequently, different values of s_0 entail different probability mass on different regions of the signal space, and a non-affine $c(\cdot)$ can exploit this by “loading” more heavily on the signal regions that are relatively more likely in states where stabilizing expectations is particularly valuable for the central bank. Imposing an affine center, on the other hand, prevents this by “forcing” the announcement to respond with the same marginal sensitivity at all signal realizations.

Similarly, allowing a signal-dependent half-width $\gamma(\tilde{s}_1)$ would let the central bank allocate discretion where it is least costly in terms of $t = 0$ stabilization. Intuitively, widening the announced set increases ex-post flexibility at $t = 1$ (which is valuable as it allows more space to react to the non-conditionable signal $\hat{s}_1$ and the previous period's shocks) yet it also weakens commitment from the perspective of the private sector. The strength of this commitment channel is not in general uniform across realizations of the public signal. In regions of the signal space where the private-sector expectations are especially sensitive to the bounds of the announcement, increasing $\gamma(\tilde{s}_1)$ can substantially worsen perceived policy uncertainty. In other regions (particularly extreme signal realizations) additional width may have a much smaller effect on perceived policy risk while still providing flexibility. A constant γ rules out this natural form of “state-contingent discretion”, forcing the central bank to trade off commitment and flexibility in the same way across all signal realizations.

3 Optimal Policy

This section presents the main theoretical results characterizing the central bank's optimal policy. I proceed in three steps. First, I solve for the optimal rate policy for an arbitrary announcement: at $t = 1$ the central bank sets the closest feasible rate to its best prediction of the full-information target, and at $t = 0$ it chooses i_0 to offset the predictable effects of shocks and announcement-driven fluctuations in expectations on output and inflation. Second, I show that optimal announcements take the form of intervals whose endpoints are affine functions of the public signal realizations. Consequently, each announcement is pinned down by three scalars: data dependence, vagueness and a center shifter. Then, I provide a characterization of the optimal announcement highlighting trade-offs involved in each component's choice.

3.1 Optimal Rate Policy

Recall there is no uncertainty from period $t = 2$ onward, and we thus assume policy is such that it fully stabilizes the economy in those periods (this is achieved by $i_t = \bar{r}^n$, $\mathcal{A}_t = \{\bar{r}^n\}$ for $t \geq 2$), a fact that is perfectly anticipated by the agents (so there is no room to manipulate expectations with announcements at $t \geq 1$). Thus, the first non-trivial problem is the $t = 1$ one. It is easy to see that the following holds. From now on, let $s_0 = (r_0^n, u_0)$.

Lemma 1. *The optimal policy rate at $t = 1$ is given by*

$$i_1^* = \arg \min_{i_1 \in \mathcal{A}_1} ||i_1 - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1, \hat{s}_1]||,$$

where $i_1^{\text{FI}}(\lambda) := r_1^n + \Gamma(\lambda)u_1$ with $\Gamma(\lambda) = \frac{\lambda\kappa}{\sigma(1+\lambda\kappa^2)}$.

Intuitively, as the central bank's loss function is quadratic, if unconstrained it would optimally set the rate equal to the expected value of the full-information optimal rate given its information (in this case, previous shocks' realizations and the draws of the two signals)¹². However, because of the announcement, it is constrained to pick a rate in the set $\mathcal{A}_1(\tilde{s}_1)$, so it optimally implements the projection of that expectation onto that set, i.e., the point in the set that lies closest to $\mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1, \hat{s}_1]$.

¹²That the full information optimal rate is given by that expression is a standard result in the New Keynesian literature (see, e.g., [Woodford \(2003\)](#)). Intuitively, demand shocks are fully stabilized as they push inflation and output in the same direction, while cost-push shocks are only partially stabilized because they affect inflation only and thus counteracting them distorts output, so a trade-off arises that the central bank resolves by partially accommodating cost-push shocks to an extent determined by, among other variables, the relative weight of inflation in its losses.

Regarding the $t = 0$ problem, the following proposition characterizes the optimal policy rate for any private sector's expectation formation rule.

Proposition 1. *Let $c := (1 - \rho_r)\bar{r}^n$ and define the period-0 intercepts*

$$C_y := \sigma(1 + \sigma\kappa)c, \quad C_\pi := \sigma\kappa(1 + \sigma\kappa + \beta)c.$$

Let $s_0 := (r_0^n, u_0)$. For any square-integrable function mapping s_0 to the private sector's expected policy rate i_1 , $(r_0^n, u_0) \mapsto i_1^e$, the optimal policy rate at $t = 0$ is given by

$$i_0^* = \bar{r}^n + \left(1 + \sigma\kappa + \frac{\lambda_{cb}\kappa^2}{1 + \lambda_{cb}\kappa^2}\beta\right) (\bar{r}^n - \mathbb{E}[i_1^e(s_0)]). \quad (2)$$

Furthermore, the corresponding minimized loss at $t = 0$ is

$$\begin{aligned} \Psi(i_1^e) := & \mathbb{E} \left[(C_y + A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e(r_0^n, u_0))^2 + \lambda_{cb} (C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e(s_0))^2 \right] \\ & - \sigma^2 (1 + \lambda_{cb}\kappa^2) (i_0^*)^2, \end{aligned}$$

where $(A_y^r, A_y^u, A_y^{i^e})$ are the loadings of y_0 on natural rate shocks, cost-push shocks and the private sector's policy rate expectations respectively (analogously defined for π_0).

Intuitively, the central bank sets its rate to offset the predictable part of the destabilizing forces that affect inflation and output at $t = 0$. Because cost-push shocks have mean zero, the predictable part of their direct destabilizing effect is zero (that is, disregarding their consequences on beliefs about next-period's policy rate). Despite this, central bank preferences encoded in λ_{cb} matter because of the effects of $t = 0$ shocks on expectations of $t = 1$ variables (both through the exogenous persistence of the natural rate process and the dependence of policy rate beliefs on $t = 0$ shocks). In particular, inasmuch as $t = 0$ shocks affect $\mathbb{E}_0^M[\pi_1]$, the divine coincidence does not hold and the extent of $t = 0$ stabilization of (the predictable component of) demand shocks depends on central bank preferences.¹³

The predictable component is the unconditional mean given the central bank's information when choosing i_0 , it cannot offset each realization of shocks that it has not yet observed. A shifted announcement may make average future-rate expectations nonzero even when the underlying shocks have mean zero. The current rate partly offsets this predictable effect, leaving a residual tradeoff between average output and average inflation stabilization.

Consequently, if $t = 0$ shocks were perfectly observed by the central bank, it could set

¹³This is because unlike in the standard New Keynesian model with rational expectations, the private sector does not believe the central bank will fully stabilize demand shocks at $t = 1$. In a version of that model with demand shocks alone, $\mathbb{E}_0^M[\pi_1] = 0$ and hence the divine coincidence holds.

i_0 to offset a larger share of the effect of announcement-driven expectations on period-0 aggregate demand. Even in that case, however, vagueness would in general not be costless: with one instrument (i_0) and two targets (y_0, π_0), the minimized $t = 0$ loss Ψ remains a quadratic form in terms that include i_1^e , with a strictly positive coefficient on $(i_1^e)^2$.

Widening the band can change both average future-rate expectations and their dispersion across current shocks. The resulting effects on output and inflation may reinforce or offset those of the shocks, so greater vagueness does not necessarily increase current stabilization losses locally. Nevertheless, the current rate cannot generally offset the effects of these expectations on both output and inflation, so the communication trade-off persists even when if current shocks were observed. Hence, information asymmetry is not necessary for the communication trade-off in the exogenous-signal model. It is kept here to preserve the structure of the endogenous-signals extension, where it is essential. In that version of the model, without asymmetry the central bank would infer the conditionable signal (contemporaneous inflation) exactly from its policy choice and observables. As a result, at $t = 0$ it would already know its $t = 1$ policy choice, collapsing the commitment-flexibility trade-off associated with communication.

3.2 The Form of Optimal Announcements

Having characterized the optimal rate policy, I now describe the form of optimal announcements in the following proposition.

Proposition 2. *The optimal communication policy entails an interval announcement, i.e., a map $\tilde{s}_1 \mapsto [L(\tilde{s}_1), U(\tilde{s}_1)]$ with $L, U \in L^1$ and $L(\tilde{s}_1) \leq U(\tilde{s}_1)$ for all $\tilde{s}_1 \in \mathbb{R}$. Furthermore, any equivariant interval announcement takes the affine-band form*

$$A(\tilde{s}_1) = [\alpha\tilde{s}_1 + b - \gamma, \alpha\tilde{s}_1 + b + \gamma]$$

for some $\alpha, b, \gamma \in \mathbb{R}$ with $\gamma \geq 0$.

The above characterization has two distinct components: (i) the optimal announcement is interval-valued; and (ii) the same-slope affine form of its endpoints. To see why intervals are optimal, fix any realization of the public signal and suppress it from notation. Consider an arbitrary non-interval announcement set $\mathcal{A} \subset \mathbb{R}$ for next period's policy rate, and let $\text{co}(\mathcal{A})$ denote its convex hull.

Market side. Because of the market's ambiguity aversion, the private sector evaluates an announcement by the worst-case rate in $\mathcal{A}$. Since the private sector's loss $\ell(i_1)$ is convex in

the policy rate, the worst case is attained at an extreme point of $\mathcal{A}$. Figure 1a illustrates this logic for $\mathcal{A} = [L, \underline{a}] \cup [\bar{a}, U]$: the relevant worst-case outcomes are the extremes L and U . Importantly, $\mathcal{A}$ and its convex hull $[L, U]$ share the same extreme points, so replacing $\mathcal{A}$ by $\text{co}(\mathcal{A})$ leaves the private sector’s worst-case evaluation unchanged.

Central bank side. A non-interval announcement has at least one “gap”, which can prevent the central bank from implementing its preferred rate after the state is realized. In the example above, if $i_1^*(\omega)$ falls in $(\underline{a}, \bar{a})$ —the shaded region in Figure 1b—the central bank must choose one of the boundary rates $\underline{a}$ or $\bar{a}$, creating an avoidable loss. Replacing $\mathcal{A}$ with its convex hull $[L, U]$ fills all such gaps and leaves the endpoints L and U unchanged. This extra flexibility can only (weakly) reduce the central bank’s expected stabilization loss: for each realized state, the central bank can always replicate its choice under $\mathcal{A}$, and in states in which $i_1^* \in (\underline{a}, \bar{a})$ it can choose a rate closer to i_1^* than the boundary points $\underline{a}$ or $\bar{a}$. If i_1^* falls in the hole with positive probability—for instance under Gaussian shocks, which have full support—the improvement is strict, so any non-interval announcement is dominated by its convex hull and an optimal announcement is interval-valued.¹⁴




Figure 1: The optimality of interval announcements.

On the other hand, the fact that the endpoints of the interval are affine in the signal follows from the equivariance assumption. Any non-affine function of the signal would not deliver the translation-invariance that equivariance requires. Further, equivariance requires a common slope too. Otherwise, changing the signal realization would also change the width of the band, breaking translation-invariance.

¹⁴Figure 1 is illustrative. Neither symmetry of $\ell(\cdot)$ nor centering of the distribution of i_1^* at the midpoint of $[L, U]$ is required for the optimality of interval announcements as stated in Proposition 2.

Given the above, an (optimal) announcement is fully pinned down by three scalars: a slope α , an intercept b and a half-width γ . In what follows, I refer to α as the data dependence of a given announcement and to γ as its vagueness (the latter owing to the fact that a higher γ announcement entails a less precise commitment).

Figure 2 illustrates how equivariant interval announcements vary with their main parameters: the slope α , the intercept b and the half-width γ . In **panel (2a)**, α and b are held fixed while γ varies: larger γ widens the admissible band at each signal realization (more vagueness) keeping the slope fixed. In **panel (2b)**, γ and b are held fixed while α varies: larger $|\alpha|$ rotates/tilts the band more strongly with the signal (greater data dependence) keeping the width fixed.




Figure 2: Equivariant interval announcements.

As an intermediate step toward characterizing the optimal announcement, the following Lemma describes the form that the private sector's rate expectations take.

Lemma 2. *For any affine interval announcement (α, b, γ) , the private sector's expected policy rate is given by*

$$i_1^e(\alpha, b, \gamma; s_0) = \alpha \mu_{\tilde{s}}(s_0) + b + \gamma \theta_{\alpha, b}(s_0),$$

where $\mu_{\tilde{s}}(s_0) := \mathbb{E}[\tilde{s}_1 \mid s_0]$,

$$\theta_{\alpha, b}(s_0) = \begin{cases} 2\Phi\left(\frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + b - a_{ps}^0(s_0)}{|\alpha - a_{ps}|\sigma_{\tilde{s}}}\right) - 1, & \alpha \neq a_{ps}, \\ 1, & \alpha = a_{ps}, \quad b \geq a_{ps}^0(s_0), \\ -1, & \alpha = a_{ps}, \quad b < a_{ps}^0(s_0). \end{cases}$$

where $\sigma_{\tilde{s}}^2 = \text{Var}(\tilde{s}_1 \mid s_0)$ and $a_{ps}, a_{ps}^0(s_0)$ are s.t. $\mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0, \tilde{s}_1] = a_{ps}\tilde{s}_1 + a_{ps}^0(s_0)$.

Intuitively, in line with the general expression in (1), for every possible signal realization, the private sector expects the policy rate to be the endpoint of the announced interval that is farthest from its expected preferred rate conditional on the realizations of $(s_0, \tilde{s}_1)$. Hence, in general the expected rate is not the interval’s center but a weighted average of the two endpoints, with weights given by how likely each endpoint is predicted to be worse for the private sector. Put differently, for a given realization of current shocks, expected future policy equals the expected announcement midpoint plus a worst-case adjustment. This adjustment is positive when the upper endpoint is more likely to be worse and negative when the lower endpoint is more likely to be worse. If the conditional mean gap between the announcement midpoint and the private sector’s (conditional) preferred rate is zero, and there is positive residual uncertainty about that gap, symmetry makes both endpoints equally likely to be worse and the adjustment vanishes. The relevant conditional standard deviation is $|\alpha - a_{ps}|\sigma_{\tilde{s}}$: uncertainty about the signal matters insofar as the midpoint and the preferred rate respond differently to it. Fixing this conditional mean gap, greater residual uncertainty moves this “expected selector” $\theta_{\alpha,b}(s_0)$ toward zero. Two limiting cases illustrate this. If this residual uncertainty goes to zero (and the conditional mean gap is non-zero), the private sector knows which endpoint is worse: $\theta_{\alpha,b}(s_0)$ approaches 1 for a positive gap and -1 for a negative one. By contrast, as this uncertainty grows without bound, which endpoint is worse becomes fully unpredictable and due to symmetry, $\theta_{\alpha,b}(s_0) \rightarrow 0$ so the expected rate approaches the expected announcement midpoint.

Finally, note that since announcements cannot condition directly on current shocks, in general they cannot eliminate conditional misalignment in every state. Nevertheless, the central bank can align the *unconditional mean midpoint* with the common mean preferred rate (in the $\bar{r}^n = 0$ case, this is done choosing $b = 0$, as explained below). Whether it wants to do so depends on how misalignment changes current expectations and future policy constraints (see the discussion following the characterization of optimal announcements in [Proposition 3](#) below).

The central bank’s problem involves choosing three announcement parameters (α, b, γ) that jointly determine both contemporaneous expectations and future policy flexibility. Each parameter entails a different dimension of the tension between $t = 0$ and $t = 1$ objectives: *data dependence* α governs how next period’s feasible set (and current private sector’s expectations) respond to signals; *the center* b shifts the average location of the announcement band; and *vagueness* γ controls the width of the band, trading off commitment value against flexibility. The optimal choice balances these competing forces.

3.3 Optimal Announcements

The mean natural rate $\bar{r}^n$ can be normalized to zero without changing optimal data dependence, vagueness, or the central bank's minimized losses. I adopt this normalization below to simplify the exposition. All a nonzero mean would entail is a shift in the level of rates and a corresponding adjustment in the announcement intercept b , as described at the end of this subsection.

Proposition 3 (Optimal announcements). *Normalize $\bar{r}^n = 0$ and let $W_1 := \beta\sigma^2(1 + \lambda_{cb}\kappa^2)$. After optimizing the current rate as in [Proposition 1](#), the central bank minimizes*

$$\mathcal{J}(\alpha, b, \gamma) = \Psi(\alpha, b, \gamma) + W_1 \mathbb{E} \left[(|z_{\alpha,b}| - \gamma)_+^2 \right]$$

over $(\alpha, b, \gamma) \in \mathbb{R}^2 \times [0, \infty)$, where $\Psi(\alpha, b, \gamma)$ is the minimized $t = 0$ loss induced by the announcement, and

$$z_{\alpha,b} := (\alpha - a_{cb})\tilde{s}_1 + b - a_{cb}^0(s_0, \hat{s}_1).$$

with $a_{cb}, a_{cb}^0(s_0, \hat{s}_1)$ such that

$$a_{cb}\tilde{s}_1 + a_{cb}^0(s_0, \hat{s}_1) = \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1, \hat{s}_1].$$

Further, define the $t = 1$ alignment target as

$$\tilde{\alpha}_1 := a_{cb} + \frac{\text{Cov}(\tilde{s}_1, a_{cb}^0(s_0, \hat{s}_1))}{\text{Var}(\tilde{s}_1)}. \quad (3)$$

The objective function is continuously differentiable for $\gamma > 0$, $\alpha \neq a_{ps}$. Any finite local minimizer with $\alpha \neq a_{ps}$ satisfies¹⁵

$$0 = \frac{\partial \Psi}{\partial \alpha} + 2W_1 \text{Var}(\tilde{s}_1)(\alpha - \tilde{\alpha}_1)\mathbb{P}(|z_{\alpha,b}| > \gamma), \quad (4)$$

$$0 = \frac{\partial \Psi}{\partial b} + 2W_1 \mathbb{E} [(|z_{\alpha,b}| - \gamma)_+ \text{sgn}(z_{\alpha,b})], \quad (5)$$

$$0 \leq \frac{\partial \Psi}{\partial \gamma} - 2W_1 \mathbb{E} [(|z_{\alpha,b}| - \gamma)_+]. \quad (6)$$

The width inequality holds with equality when $\gamma > 0$ and uses the right derivative at $\gamma = 0$.

Finally, if both shock processes have nonzero persistence and positive variances, and $\tilde{\omega}_u \neq \Gamma(\lambda_{ps})\tilde{\omega}_r$, there exists a global minimum at finite (α, b, γ) .

¹⁵Sufficient conditions for a strict local minimum are given in [Appendix C](#).

The announcement’s location: predictability versus mean alignment. Under the $\bar{r}^n = 0$ normalization, setting $b = 0$ aligns the expected midpoint of the announcement with the unconditional mean of both the central bank’s and the private sector’s preferred rates.¹⁶ For any slope α and half-width γ , this minimizes future losses from the announced constraint. Furthermore, the normalization implies the Gaussian distribution of shocks is centered at zero. Its symmetry, then, implies that every realization of shocks has an equally likely counterpart with its signs reversed. Since stabilization losses are quadratic, shifting the announcement upward or downward by the same amount yields the same expected total loss (after optimally adjusting the current rate i_0). Thus, under the normalization, the central bank has no preferred direction of misalignment. This does not, however, imply that it always prefers a centered announcement: mean misalignment in either direction may improve stabilization by making the worst endpoint more predictable.

The reason is that the announcement’s intercept b also shapes the private sector’s worst-case selection. Shifting the midpoint toward one side makes the same endpoint increasingly likely to be the worse one under the central bank’s beliefs. This reduces the ex-ante volatility of the selection across $t = 0$ shocks and thus can improve current stabilization (note, however, that it also changes the selector’s covariance with shocks, which may be stabilizing or destabilizing). The central bank weighs these effects against a systematic distortion in expected policy and a less “well-positioned” feasible set for future decisions.

The current policy rate partly offsets the predictable effect of misaligned expectations, as characterized in [Proposition 1](#). Naturally, in general it can not eliminate both mean output and mean inflation deviations, leaving a residual stabilization cost. Thus, this misalignment of the average announcement is attractive only when its benefits for current stabilization outweigh this residual mean cost and the additional future constraint losses. When it is optimal to choose a centered announcement and the selection is nondegenerate, symmetry gives $\mathbb{E}[i_1^e] = 0$ and hence $i_0^* = 0$.

Note that these incentives do not necessarily lead the central bank to introduce a small amount of mean misalignment: exact mean alignment, $b = 0$, is optimal over a range of parameters (i.e., is not knife-edge). Moreover, an aligned announcement can be locally optimal even when a different policy combining greater vagueness with substantial mean misalignment delivers lower losses. In such cases, the central bank effectively compares two communication strategies: an aligned announcement with a vaguer, mean-misaligned one. The latter makes worst-case selection more predictable, while its greater vagueness

¹⁶This is because ex-post, the only source of disagreement regards how much to accommodate cost-push shocks due to the potentially different weight both agents put in inflation stabilization, and those shocks have mean zero as well.

preserves flexibility for future stabilization despite the misalignment. However, this is not always the case: under some parameters, even small levels of mean misalignment can improve stabilization.

Data dependence: tracking the preferred rate versus shaping shock transmission.

Period $t = 0$ incentives shaping the optimal extent of data dependence can be split into two components: reducing volatility from the signal-driven band location and reducing volatility from the worst-case selection. To isolate the former, define

$$\tilde{\alpha}_0 := \arg \min_{\alpha \in \mathbb{R}} \Psi(\alpha, 0, 0). \quad (7)$$

This target minimizes the period-0 stabilization loss generated by direct shocks and the signal-dependent band center, before taking into account the worst-case selection. Its explicit expression is given in [Appendix B.5](#).

A second set of effects on $t = 0$ losses operates through worst-case selection: changing the slope changes its covariance with shocks and with the signal-driven midpoint, as well as the selector’s own response to current states. These effects can reinforce or offset the pull toward $\tilde{\alpha}_0$, depending on the signal structure and preferences.

When the announcement is misaligned, the slope also changes average worst-case expectations by changing how consistently one endpoint is selected. This adds a mean-stabilization consideration, even though the public signal has mean zero. It is absent for centered announcements, where symmetry keeps average expectations at zero.

Regarding $t = 1$ incentives, future stabilization continues to pull data dependence toward $\tilde{\alpha}_1$, the best linear prediction coefficient of the central bank’s conditional preferred rate on the public signal. The strength of this motive is proportional to the probability that the announcement binds ex-post. Importantly, the target accounts both for the central bank’s direct use of the public signal (represented by a_{cb} , the ex-post preferred weight on the conditionable signal) and for the information in current shocks and the non-conditionable signal that can be projected onto it (represented by the remaining “correction”) due to the unfeasibility of fully state contingent announcements. Misalignment changes the frequency of binding and hence the weight on this target, but does not change the target itself.

Optimal data dependence balances these future-alignment incentives against current stabilization and worst-case selection effects.

Vagueness: calming markets vs. retaining flexibility. Fixing the slope and intercept of a given announcement, the marginal current stabilization cost rises with vagueness

(provided the selector θ is not identically zero). Further, greater vagueness changes current stabilization through the covariance of worst-case selection with shocks and through its variation across states. When the band is misaligned, it also increases the systematic distortion in expected policy by moving the more frequently selected endpoint farther from the midpoint. The current rate i_0 partly offsets that average effect, but, as discussed above, leaves a residual stabilization cost.

At $t = 1$, widening the band reduces losses whenever the preferred rate lies outside it. The marginal benefit is proportional to the expected distance of that rate from the nearest feasible endpoint.

This illustrates the commitment–flexibility tradeoff: the central bank balances the current consequences of weaker constraints on beliefs against the ability to implement a more appropriate future rate. With a misaligned announcement, vaguer announcements help accommodate both the information the announcement could not ex ante condition on and the deliberate shift in its midpoint (see discussion above). See [Figure 3](#) for a graphical illustration of the role vagueness plays.




Figure 3: The value of vagueness. The left panel shows the central bank losses as a function of the policy misalignment $z_{\alpha,b}$ (the midpoint of the announcement minus the preferred rate) for two values of vagueness $\gamma_1 > \gamma_2$. The right panel shows the probability of the announcement binding whenever the misalignment has mean zero.

Beyond the zero mean natural rate. The zero-mean normalization leaves optimal data dependence, vagueness, and minimized loss unchanged in the following sense. If $(i_{0,0}^*, \alpha^*, b_0^*, \gamma^*)$ is an optimal policy in the zero-mean case, the corresponding policy with mean natural rate $\bar{r}^n$ is given by

$$i_{0,\bar{r}^n}^* = i_{0,0}^* + \bar{r}^n, \quad b_{\bar{r}^n}^* = b_0^* + (1 - \alpha^* \tilde{\omega}_r) \bar{r}^n.$$

Thus, the announcement midpoint can be written

$$\alpha^* \tilde{s}_1 + b_{\tilde{r}^n}^* = \tilde{r}^n + \alpha^* (\tilde{s}_1 - \tilde{\omega}_r \tilde{r}^n) + b_0^*.$$

The normalized intercept b_0^* represents systematic misalignment from the common unconditional mean preferred rate. A nonzero mean natural rate simply shifts levels without changing the tradeoff determining this misalignment.

In the next section, I modify the model to let the central bank condition on an endogenous variable: inflation. While the trade-off between stabilizing expectations today and preserving flexibility tomorrow remains, communication now also changes the variable on which its own commitment depends, which introduces several new and important elements to the analysis.

4 Endogenous Signals

In the baseline model above, the signal on which announcements are conditioned has a fixed exogenous structure. However, both in practice and in standard macroeconomic models, central banks often anchor their communication on observables that are determined in equilibrium and thus depend on policy and announcements themselves (e.g., inflation, unemployment rates, layoff rates). This section addresses this by studying optimal announcements with inflation as the conditioning variable.

I first outline the modified environment and timing ([Section 4.1](#)), then characterize the inflation-expectations feedback that this choice creates: because announcements are tied to π_0 , and π_0 itself depends on worst-case policy expectations, inflation and beliefs are jointly determined. [Proposition 4](#) characterizes this joint determination, showing how different announcement choices can either work as a stabilizing force or a destabilizing one, and what implications this has for equilibrium uniqueness. I then show the existence of optimal announcements and give necessary conditions for finite stabilizing optima ([Proposition 5](#)), and contrast it with the exogenous-signals baseline. Three differences stand out. First, arbitrarily vague announcements may be optimal as the feedback loop limits $t = 0$ volatility. Second, making announcements vaguer doesn't just widen the range of possible rates—it also can shift where that range is centered. Third, the value of vagueness depends on how responsive the announcement is to inflation (i.e., the extent of data dependence). These results set the stage for the key lessons that follow and clarify how the endogeneity of the conditioning variable reshapes the trade-offs governing central bank communication.

4.1 Inflation-Anchored Announcements

The changes relative to the baseline model in [Section 2](#) are: (i) the central bank's announcements are contingent on contemporaneous inflation, an equilibrium object; and (ii) exogenous signals are removed, so at $t = 1$ the only new information available to the central bank is the realization of (π_0, y_0) .¹⁷ Inflation is the natural choice for a policy anchor, as it is publicly observed and routinely used in central bank communication. Exogenous signals are shut down purely for expositional simplicity: keeping them as further information available at $t = 1$ leaves the qualitative results unchanged. The sequence of events in periods $t \in \{0, 1\}$ is as follows.

- The central bank sets the contemporaneous policy rate i_0 and makes an announcement regarding next period's rate conditional on current period's inflation, $\mathcal{A}_1(\pi_0)$.
- Shocks (r_0^n, u_0) realize.
- Given (r_0^n, u_0) , policy i_0 and expectations of their future counterparts, the private sector determines (π_0, y_0) (observed by the central bank).
- The central bank sets its policy rate $i_1 \in \mathcal{A}_1(\pi_0)$.
- Shocks (r_1^n, u_1) realize.
- Given (r_1^n, u_1) and policy i_1 , the private sector determines (π_1, y_1) .

It is worth noting that, while information asymmetry and shocks' persistence are present in both models, their significance for the analysis is different. In the exogenous-signals model, removing either shock persistence or asymmetric information has the same qualitative implication for optimal communication: the response of market beliefs to $t = 0$ policy and announcements becomes fully predictable to the central bank. Under no persistence, this happens because beliefs no longer depend on realized $t = 0$ shocks; under symmetric information, it happens because the central bank directly observes those shocks. In both cases, the central bank can use i_0 and announcements more effectively for stabilization.¹⁸ Even then, full stabilization at $t = 0$ is generically infeasible, so optimal announcements still trade off expectations-management against $t = 1$ stabilization.

¹⁷Formally, the latter is equivalent to letting $\sigma_{\tilde{\epsilon}}, \sigma_{\epsilon} \rightarrow \infty$ in the baseline model, so exogenous signals are effectively uninformative.

¹⁸The two cases are not identical: without asymmetric information, the central bank can also tailor rates and announcements to offset the direct effect of $t = 0$ shocks (in addition to their effects via expectations). This distinction is secondary here because it is not specific to the expectations-management channel.

On the other hand, in the endogenous-signals model, removing shock persistence makes π_0 uninformative about the relevant $t = 1$ state. From the perspective of $t = 1$ losses alone, the central bank should therefore set $\alpha = 0$. Yet, because of the stabilization effects at $t = 0$, it might still be optimal to set $\alpha \neq 0$. Removing asymmetric information also has two effects. First, the response of market expectations to $t = 0$ policy and announcements becomes fully predictable (again as in the exogenous-signals model). Second, the communication trade-off changes fundamentally: the central bank knows (π_0, y_0) as a function of its policy decisions, and since this is the relevant information set for choosing i_1 , it already knows at $t = 0$ which i_1 it will set ex post.

Finally, note that the arguments proving the optimality of affine interval announcements (within the equivariance class, see [Proposition 2](#)) in the baseline model apply mutatis mutandis to this variant. Thus, the analysis is restricted to announcements of that form.

4.2 The Inflation-Beliefs Feedback Loop

Recall that the private sector is ambiguity-averse regarding future policy rates: it evaluates announcements by the worst-case rate in the announced set, expecting whichever endpoint of the interval is farther from its preferred full-information rate (as formalized in [Section 2.3](#)). Because of this ambiguity aversion, the private sector's expectations of i_1 for given shocks (r_0^n, u_0) depend on what it expects inflation to be. Two differences from the baseline setup arise. First, conditional on shocks, π_0 is observed without noise, so the private sector knows exactly which end of the band is worse. More importantly, π_0 itself is shaped by private-sector behavior and depends on the private sector's worst-case rate expectation, while that expectation is in turn anchored by the announcement through $\alpha\pi_0$. Expectations and inflation are thus jointly determined, creating a feedback loop that does not arise with exogenous signals. As is standard in macroeconomics, I take the representative agent to stand for an aggregation of atomistic agents. Thus each agent takes the "aggregate policy rate expectations" (and thus, inflation) as given when forming its policy rate expectations, even though the former is in equilibrium determined by the latter. The following proposition characterizes the joint equilibrium determination of π_0 and the private sector's expectations following an announcement.

Proposition 4 (Equilibrium with inflation-anchored announcements). *For any policy $p = (i_0, \alpha, b, \gamma)$ with $\alpha \neq \alpha^* := \frac{1}{A_\pi^{i^e}}$, let $\Theta_p(\pi, s_0) \subseteq [-1, 1]$ denote the aggregate worst-case rate selection given inflation π and shocks s_0 . A $t = 0$ equilibrium is a pair (π_0, ϑ_0) such that*

$$\vartheta_0 \in \Theta_p(\pi_0, s_0) \quad \text{and} \quad \pi_0 = \frac{C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} (b + \gamma \vartheta_0)}{1 - \alpha A_\pi^{i^e}},$$

where the coefficients are as defined in *Proposition 1*. Equivalently, define the equilibrium selection set

$$\mathcal{E}_p(s_0) := \left\{ \vartheta \in [-1, 1] : \exists \pi \text{ s.t. } \vartheta \in \Theta_p(\pi, s_0) \text{ and } \pi = \frac{C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} (b + \gamma \vartheta)}{1 - \alpha A_\pi^{i^e}} \right\}.$$

Then (π_0, ϑ_0) is an equilibrium if and only if $\vartheta_0 \in \mathcal{E}_p(s_0)$ and π_0 is given by the above display.

If $\gamma = 0$, the two endpoints coincide. The equilibrium expected rate and inflation are uniquely determined for every $\alpha \neq \alpha^*$ by

$$i_1^e = \frac{\alpha(C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0) + b}{1 - \alpha A_\pi^{i^e}}, \quad \pi_0 = C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} i_1^e.$$

The endpoint worst-case selector is immaterial in this case, $\gamma \vartheta_p = 0$. For $\alpha \neq 0$, define

$$T(s_0; p) := -\frac{1}{\alpha A_\pi^{i^e}} \left((\alpha A_\pi^r - (1 - \alpha A_\pi^{i^e}) \rho_r) r_0^n + (\alpha A_\pi^u - (1 - \alpha A_\pi^{i^e}) \Gamma(\lambda_{ps}) \rho_u) u_0 + b - (\alpha \sigma \kappa) i_0 - (1 - \rho_r) \bar{r}^n \right).$$

If $\gamma > 0$, the equilibrium selection set is characterized as follows.

- If $\alpha > 0$ or $\alpha < \alpha^*$, the equilibrium selection set is a singleton $\mathcal{E}_p(s_0) = \{\vartheta_p(s_0)\}$ with

$$\vartheta_p(s_0) = \begin{cases} -1 & \text{if } T(s_0; p) < -\gamma \\ \frac{T(s_0; p)}{\gamma} & \text{if } -\gamma \leq T(s_0; p) \leq \gamma \\ +1 & \text{if } T(s_0; p) > \gamma. \end{cases}$$

- If $\alpha \in (\alpha^*, 0)$, then

$$\mathcal{E}_p(s_0) = \begin{cases} \{-1\} & \text{if } T(s_0; p) > \gamma \\ \left\{-1, \frac{T(s_0; p)}{\gamma}, +1\right\} & \text{if } -\gamma \leq T(s_0; p) \leq \gamma \\ \{+1\} & \text{if } T(s_0; p) < -\gamma. \end{cases}$$

- If $\alpha = 0$, then

$$\mathcal{E}_p(s_0) = \begin{cases} \{-1\} & \text{if } b < \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) | s_0] \\ [-1, +1] & \text{if } b = \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) | s_0] \\ \{+1\} & \text{if } b > \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) | s_0]. \end{cases}$$

The object $T(s_0; p)$ can be interpreted as a scaled mis-centering index. It measures how

far the midpoint of the announcement band, $\alpha\pi_0 + b$, is from the private sector's preferred full-information rate, excluding the worst-case selection component. The scaling factor reflects both how strongly inflation reacts to worst-case expectations (through $\alpha A_\pi^{i^e}$) and how this response is amplified or dampened by the feedback loop (through $1 - \alpha A_\pi^{i^e}$). In other words, it measures how much a shift in the worst-case endpoint ultimately moves inflation and hence the perceived mis-centering. Naturally, instances where $T \ll 0$ in the policy regions where the equilibrium is unique amount to situations where the private sector believes the mean of the band is too low, so it naturally makes the worst-case selection the lower bound. Conversely, large values of T lead to the upper bound being selected.

Notably, and unlike in the exogenous-signals model, the market's expected worst-case selection ϑ_p is not a smooth function of its perceived misalignment. In the unique equilibrium regions this selection is continuous and has kinks at the thresholds. Under the convention used below, the selected equilibrium in the multiplicity region instead jumps at those thresholds. This is because now the private sector knows exactly which endpoint of the band is worse, as the private sector faces no uncertainty regarding π_0 (unlike $\tilde{s}_1$ in the exogenous signals model). Consequently, inflation-anchored announcements induce a piecewise worst-case selection. For high enough misalignment, the selection is insensitive to marginal changes in misalignment itself, and hence locally independent of data dependence or band centering. By contrast, for intermediate levels of misalignment, the contribution of the worst-case selection to expected rates, $\gamma\vartheta_p = T(s_0; p)$, is locally insensitive to vagueness. This follows from the fact that following a marginal change in the announcement's vagueness, the fractions of agents expecting each endpoint adjust, offsetting the widening of the band. Nevertheless, all three dimensions of the announcement affect the probability that the equilibrium worst-case selection falls in each region.

Stabilizing vs. destabilizing announcements. When $\alpha > 0$ or $\alpha < \alpha^*$, the feedback through expectations works as a stabilizer. In the former case, if expectations lean toward the upper endpoint of the band, inflation decreases and thus the mis-centering decreases too, preventing the private sector from wanting to move its expectations further in that direction. An analogous mechanism operates if expectations lean toward the lower endpoint of the band. In the latter case, where $\alpha < \alpha^*$, the logic is similar: if expectations lean toward the upper endpoint of the band, inflation increases, and thus, because $\alpha^* < 0$, the mis-centering decreases. In this scenario, though, the stabilization relies on the feedback through policy rate expectations being so strong as to reverse the “direct” effect of shocks on inflation. This stabilizing effect of policy ensures the uniqueness of the equilibrium worst-case

selection.¹⁹

By contrast, when $\alpha \in (\alpha^*, 0)$, feedback through policy expectations is destabilizing: leaning toward one endpoint increases the mis-centering in the same direction, so both corners can be self-confirming. Naturally, this leads to multiple equilibria in some regions of the shocks.

Furthermore, when $\alpha = 0$, expectations of future policy rates are independent of inflation, so the equilibrium reduces to a simple comparison of the band's center b with the private sector's preferred rate.

For nonzero data dependence and nonzero vagueness, an intermediate equilibrium exists when $|T| \leq \gamma$. Its aggregate worst-case selector is T/γ , the fraction of private sector agents choosing the upper endpoint minus the fraction choosing the lower endpoint. This intermediate split is unique due to the feedback into the evaluation of the worst-case (although in the multiplicity region there are also equilibria where the entire private sector agrees on a worst-case endpoint). At zero data dependence, intermediate selections require the tie $b = \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) | s_0]$, which is a null event under the nondegeneracy assumption.

In summary, in this variant of the baseline model the announcement not only constrains the future policy rate but also shapes the distribution of the signal itself. This gives rise to regimes where worst-case beliefs are pinned down uniquely (negative feedback) and regimes where self-confirming multiplicity arises (positive feedback). Importantly, both the direction and the strength of these feedbacks are shaped by announcements, and thus become a key component in the determination of the optimal announcement policy.

Since the $t = 0$ equilibrium is not necessarily unique (Proposition 4), the policy problem studied below requires an equilibrium-selection rule. I use the convention stated below throughout the analytical and numerical analysis. At zero data dependence, the tie $b = \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) | s_0]$ has probability zero given the Gaussian specification, so its tie-break does not affect expected losses. The main numerical figures select positive data dependence, where the equilibrium is unique.

Selection convention. Proposition 4 implies that for each policy $p = (i_0, \alpha, b, \gamma, s_0)$ the set of $t = 0$ equilibrium worst-case selections is $\mathcal{E}_p(s_0) \subseteq [-1, 1]$. For the remainder of the

¹⁹Note that this mechanism is quite close to the standard one in New Keynesian models with rational expectations, albeit in those models it does not only operate intertemporally but also contemporaneously, as it is assumed that rates react to contemporaneous values of endogenous variables (chiefly among them, inflation).

paper, fix a measurable selection for $\gamma > 0$ ²⁰

$$\vartheta_p(s_0) \in \mathcal{E}_p(s_0)$$

as follows: when $\mathcal{E}_p(s_0)$ is a singleton, $\vartheta_p(s_0)$ is its unique element; when $\alpha \in (\alpha^*, 0)$ and $|T(s_0; p)| \leq \gamma$, select the intermediate equilibrium $\vartheta_p(s_0) = T(s_0; p)/\gamma$; and when $\alpha = 0$ and $b = \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0]$, set $\vartheta_p(s_0) = 0$ (otherwise $\vartheta_p(s_0) = \pm 1$ as in [Proposition 4](#)).

With this selection convention in place, equilibrium outcomes define a single-valued ex-ante loss for the central bank. The following proposition establishes the existence of a solution allowing limiting outcomes and gives necessary first-order conditions for finite regular optima *in the stabilizing regions*. Policies in the multiplicity region remain feasible, but their first-order conditions involve additional terms because the selected equilibrium jumps at the switching thresholds.

4.3 Optimal Inflation-Anchored Announcements

In what follows, I present and discuss the necessary first-order conditions for the optimal communication and rate policy at $t = 0$, highlighting the differences with the baseline model. Note that with inflation-anchored announcements, the optimal policy rate at $t = 0$ cannot be solved in closed form as in the baseline model (see [Proposition 1](#)), because it affects $t = 1$ losses via shaping π_0 . Hence, the first-order condition for the policy rate is characterized jointly with those corresponding to the announcement parameters.

Proposition 5 (Optimal inflation-anchored announcements). *Assume the equilibrium-selection convention following [Proposition 4](#) and define the policy domain $\mathcal{D} = \mathbb{R} \times \mathbb{R} \setminus \{\alpha^*\} \times \mathbb{R} \times \mathbb{R}_+$. For $p = (i_0, \alpha, b, \gamma) \in \mathcal{D}$, let (y_0, π_0) denote the selected equilibrium outcomes and write*

$$\tilde{\Psi} := \mathbb{E}[y_0^2 + \lambda_{cb}\pi_0^2], \quad z := \alpha\pi_0 + b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0].$$

Up to a policy-independent constant, expected loss is finite, continuous and given by

$$\tilde{\mathcal{J}}(p) = \tilde{\Psi} + W_1 \mathbb{E}[(|z| - \gamma)_+^2], \quad W_1 := \beta\sigma^2(1 + \lambda_{cb}\kappa^2).$$

There exist a solution to the central bank's problem: the minimum expected loss is attained either by a policy with finite parameters or by a limit of a sequence of finite policies with $\gamma \rightarrow \infty$. The proof characterizes these limits and gives a sufficient condition for a finite

²⁰As mentioned above, if $\gamma = 0$, the selection is immaterial since $\gamma\vartheta = 0$ and thus the value of ϑ is irrelevant.

minimizer.

Suppose

$$\rho_r(1 + (\sigma\kappa + \beta)\rho_u) \neq \lambda_{ps}\kappa^2(\rho_u - \rho_r(1 - \rho_u)). \quad (8)$$

Then, any finite local minimizer with $\alpha \in (-\infty, \alpha^*) \cup (0, \infty)$ satisfies

$$\begin{aligned} \frac{\partial \tilde{\mathcal{J}}}{\partial i_0} &= \frac{\partial \tilde{\mathcal{J}}}{\partial \alpha} = \frac{\partial \tilde{\mathcal{J}}}{\partial b} = 0, \\ \frac{\partial \tilde{\mathcal{J}}}{\partial \gamma} &\geq 0, \quad \gamma \frac{\partial \tilde{\mathcal{J}}}{\partial \gamma} = 0. \end{aligned} \quad (9)$$

All derivatives are evaluated at the candidate, with the vagueness derivative taken from the right when $\gamma = 0$. Their expanded expressions can be found in (21) in the appendix.

Key differences with the baseline model. Beyond the direct consequence of using π_0 as the conditionable variable (i.e., tying the information structure to the model’s physical environment), the use of inflation-anchored announcements changes the central bank’s problem in three substantive ways.

1. *Fully vague announcements.* The limit $\gamma \rightarrow \infty$ corresponds to fully vague guidance, which places no effective constraint on future policy. With exogenous signals this is always dominated: widening the band indefinitely generates increasingly extreme policy expectations, making current stabilization losses grow without bound.²¹ With stabilizing inflation-anchored announcements, by contrast, as announcements get vaguer, an equilibrium in which all agents expect the same endpoint requires increasingly large shock realizations. For $\alpha > 0$, the empirically relevant case, this is because overly high or overly low aggregate rate expectations move inflation in the opposite direction, hence feeding back through the mean-band term to reduce the magnitude of those expectations. Hence, for each fixed shock realization, the equilibrium is in the intermediate once the band is sufficiently wide. Further, within this intermediate region, inflation and aggregate rate expectations are insensitive to γ , since widening the band moves both endpoints out symmetrically, and the difference between the fractions of agents at the two endpoints varies inversely with γ so that the average worst-case expectation doesn’t move. Consequently, as $\gamma \rightarrow +\infty$, the central bank’s overall loss converges to a finite constant, making fully vague announcements potentially optimal. However, this is not necessarily the case, since by picking fully vague announcements the worst-case

²¹See the proof of Proposition 3, where it is shown that $t = 0$ stabilization losses diverge as $\gamma \rightarrow \infty$. As stated in the proposition, this requires some assumptions, albeit they are mild and hold generically.

selection $\gamma\vartheta_p$ is not truncated and thus induces potentially high output and inflation volatility at $t = 0$.²²

2. *An additional effect of vagueness.* With exogenous signals, making the announcement vaguer while keeping its other dimensions unchanged leaves the midpoint of the band unchanged for every signal realization. Thus, vagueness affects $t = 1$ losses only via changing the flexibility the central bank has to set rates. With inflation-anchored announcements, vagueness can also change equilibrium inflation and hence the midpoint $\alpha\pi_0 + b$, even if (i_0, α, b) are kept fixed. When the announcement is stabilizing (so the policy-expectations loop pushes misalignment back toward the target), widening the band tends to move the midpoint toward the central bank’s target in the very states where the announced band would bind. Intuitively, when $t = 0$ shocks are, say, large and positive, the private sector’s worst-case rates are excessively low, pushing π_0 upwards and thus moving the band to the right and allowing the central bank to pick a rate closer to its unconstrained optimum (which will be large and positive because shocks were too high). While this increases the value of vagueness at $t = 1$, the “endogenous re-centering” lowers the probability of announcements being binding ex-post (and the size of the loss if they are) and thus generates less need for flexibility. Hence, whether this pushes optimal vagueness upward or downward is unclear (one can think of it as the “price” of flexibility in terms of vagueness decreases, so there are both “income” and “substitution” effects at play that go in opposite directions). Importantly, this relies on the perceptions of misalignment by the central bank and the private sector being reasonably aligned (i.e., $\lambda_{ps} \approx \lambda_{cb}$).
3. *The interaction between vagueness and data dependence.* Because data dependence mutes the sensitivity of inflation to shocks via the beliefs feedback loop, it shapes how strong the endogenous band re-centering via the worst-case mechanism above is. For positive data dependence ($\alpha > 0$), two effects work in opposite directions. Higher α dampens the impact of beliefs (the worst-case term $\gamma\vartheta_p$) on inflation π_0 , so inflation responds less to a given change in policy expectations. However, because the band midpoint is $C = \alpha\pi_0 + b$, the induced responsiveness of the band center to the worst-case component moves in the opposite direction and increases with α for $\alpha > 0$. As a result, the marginal value of γ for flexibility depends on α in a richer way than with exogenous signals (see discussion below [Proposition 3](#)).

²²For policies that induce multiple equilibria, this logic and the consequent possibility of fully vague announcements apply given the equilibrium selection criterion above.

5 Numerical Analysis

This section uses the endogenous-signals model to draw policy lessons from illustrative numerical exercises. While quantitative in nature, the analysis is not meant to deliver definitive numerical estimates, but rather to highlight how the model’s mechanisms translate into policy-relevant implications. I proceed in four steps. First, I present the baseline calibration. Second, I present and discuss the optimal announcement implied by this calibration. Third, I assess the stabilization cost of strict rules (the additional losses from constraining the central bank to make fully precise announcements). Finally, I compare optimal data dependence to the Delphic (best-forecast) benchmark and discuss the implications for optimal communication.

5.1 Calibration

The New Keynesian parameters follow [Woodford \(2003\)](#), itself based on [Rotemberg and Woodford \(1997\)](#): $\kappa = 0.024$, $\beta = 0.99$, and $\sigma = 6.25$. The welfare-based weight in [Woodford \(2003, p. 400, equation \(2.22\)\)](#) entails a $\lambda = \theta/\kappa$, where $\theta = 7.88$ is the elasticity of substitution across goods. Thus, since I abstract from preference disagreements between the central bank and the private sector, I set $\lambda_{cb} = \lambda_{ps} = 7.88/0.024 \approx 328.33$.

Both shock processes are calibrated using the no-price-indexation specification of [Cole and Martínez-García \(2023\)](#). For the natural rate, I derive the dynamics implied by their model, using its estimated productivity process and reported preference parameters. The resulting persistence is $\rho_r \approx 0.3813$ and the unconditional standard deviation is $\sigma_r \approx 0.3062\%$ quarterly.²³ For the cost-push process, the same specification gives $\rho_u = 0.75$ and a quarterly innovation standard deviation of 0.08% . The corresponding unconditional standard deviation is $\sigma_u = 0.08\%/\sqrt{1 - 0.75^2} \approx 0.1209\%$. As in this model, the two innovations are independent.

The natural-rate mean is set to a quarterly value of $\bar{r}^n = 0.525\%$, obtained by dividing the New York Fed’s December 2024 forecast of the end-2025 annualized natural rate, 2.1% by four ([Cho et al., 2024](#)). Note, however, that as discussed above this choice is immaterial for the key dimensions of optimal announcements: changing $\bar{r}^n$ translates rate levels without affecting optimal data dependence or vagueness. Cost-push shocks’ unconditional mean is set to zero, following the conventional symmetric specification (see, e.g. [Woodford \(2003\)](#); [Galí \(2015\)](#)).

²³The calculation uses equations (11)–(13) on p. 538 of [Cole and Martínez-García \(2023\)](#), evaluated at the posterior means in Table 10 of the online appendix for the no-price-indexation specification, with the fixed discount factor and habit parameter given in note 26 (p. 564). At the reported rounded parameter values, the implied natural-rate process is AR(1).

Table 1 summarizes the baseline calibration.

Parameter	Interpretation	Value	Source / rationale
κ	Phillips curve slope	0.024	Woodford (2003).
β	Discount factor	0.99	Woodford (2003).
σ	Intertemporal elasticity of substitution	6.25	Woodford (2003).
$\bar{r}^n$	Mean natural rate	0.525%	NY Fed December 2024 forecast.
ρ_r	Natural rate persistence	0.3813	Derived from the model in Cole and Martínez-García (2023).
σ_r	Natural rate standard deviation	0.3062%	Derived from the model in Cole and Martínez-García (2023).
$\bar{u}$	Mean cost-push shock	0%	Symmetric benchmark.
ρ_u	Cost-push shock persistence	0.75	Cole and Martínez-García (2023).
σ_u	Cost-push shock standard deviation	0.1209%	Cole and Martínez-García (2023).
λ_{cb}	Central bank's relative weight on inflation	328.33	Welfare-implied weight.
λ_{ps}	Private sector's relative weight on inflation	328.33	Equal to λ_{cb} (no preference disagreement).

Table 1: Baseline calibration. Rates and shock standard deviations are in quarterly units, σ_r and σ_u are unconditional standard deviations.

In the exercises that follow, I systematically vary key parameters to explore how optimal communication depends on structural features of the economy. Specifically, I examine how results change with the volatility of natural rate shocks (σ_r), the volatility of cost-push shocks (σ_u), the relative weight on inflation stabilization (λ), and the slope of the Phillips curve (κ). This sensitivity analysis helps identify which features of the economic environment are most important for communication design.

5.2 Optimal Announcements

The optimal communication policy for next quarter's rate implied by the endogenous-signals model is characterized by

$$\alpha^* \approx 2.85$$

and

$$\gamma^* \approx 0.00156,$$

which corresponds roughly to an annualized interval width of 125 basis points.

A natural empirical counterpart to the model’s degree of data dependence is the estimated inflation coefficient in Taylor rules. The canonical paper in this literature, [Clarida et al. \(2000\)](#), reports a point estimate of 2.15, which is close to our implied optimal coefficient. While the comparison is not exact (their estimate refers to the coefficient on expected inflation one period ahead, and their sample ends in 1996Q4), the similarity is reassuring.

On the other hand, a plausible empirical counterpart to the degree of vagueness is the range of policy-rate projections for the end of 2025 reported by FOMC participants in the December 2024 Summary of Economic Projections, which spans 125 basis points in annualized terms.²⁴ A market-based alternative can be built using the Kansas City Fed’s *Measure of Policy Rate Uncertainty* (KC PRU), an option-implied measure of uncertainty about short-term interest rates one year ahead ([Bundick et al., 2024](#)). Its December 18, 2024 value implies an approximate 95% range of 448 basis points under a Gaussian approximation (annualized).

However, these measures incorporate uncertainty or disagreement about future economic conditions. The model’s notion of vagueness is narrower: it captures the range of discretion the central bank retains for a given realization of its conditioning variable (inflation, in this case). Moreover, the SEP measure reflects an aggregation of individual projections and not a collective commitment²⁵, while the market-based measure reflects risk-adjusted beliefs, not the underlying probability assessments of future policy rates (as all measures built from asset prices). Nonetheless, the fact that the calibration-implied level of optimal vagueness is of the same order of magnitude as these two empirical measures is reassuring.

I next study how optimal announcements vary with demand and supply (cost-push) uncertainty. In each case, we consider standard deviation values from one-half to twice its baseline value, holding the remaining parameters fixed.

Demand uncertainty and the design of guidance. Greater demand uncertainty makes large spending and inflation fluctuations more likely, increasing the value of a stronger *anticipated* policy response. Higher data dependence shifts both endpoints of the announced band upward when inflation is positive and downward when it is negative. Thus, even a skeptical private sector anticipates more tightening in an expansion and more easing in a contraction. Through expectations of future activity and inflation, these responses restrain current demand pressures. Since demand shocks move output and inflation in the same direction, this feedback helps stabilize both. In this calibration, the announced midpoint already tracks

²⁴Available at <https://www.federalreserve.gov/monetarypolicy/fomcprojtabl20241218.htm>.

²⁵In addition, these projections concern a one-year horizon approximately, not the following quarter as in the model.

demand-induced movements in the preferred future rate relatively closely. Greater demand uncertainty consequently raises the current stabilization benefit of stronger data dependence much more than its marginal cost from restricting future policy at the original announcement. The central bank therefore increases α , accepting additional future misalignment in exchange for better current stabilization.

Optimal vagueness is approximately flat at low levels of demand uncertainty and then increases substantially (Figure 4b). This pattern is explained by the interaction between uncertainty and data dependence. When data dependence is low, positive demand shocks raise the preferred future rate more than the announced midpoint, shifting skeptical expectations toward the lower endpoint. Vaguer announcements thus make current demand pressures worse, so greater demand uncertainty favors narrower announcements for fixed data dependence. The central bank also increases α , however. Stronger data dependence moderates this effect and dampens the transmission of wider bands to current expectations. It also strengthens the flexibility provided by additional width through the endogenous adjustment of the band's midpoint. Near the lower end of the plotted range, these effects approximately offset one another. As demand uncertainty rises further, the stronger adjustment in data dependence shifts the balance toward vaguer announcements.

In short, higher demand uncertainty calls for announcements that are more responsive to data and—over most of the plotted range—vaguer.




Figure 4: Higher demand uncertainty (σ_r) raises both the optimal responsiveness of the announced range to inflation and optimal vagueness over most of the plotted range. The latter ($2\gamma^*$) is measured in annualized basis points. The plotted range is 0.5–2 times the baseline standard deviation, the dotted line marks the baseline. All other parameters are held at their baseline values.

Cost-push uncertainty and the design of guidance. When uncertainty about cost-push shocks increases, the logic of optimal data dependence policy partially reverses. Greater cost-push uncertainty increases the importance of inflation movements that the central bank would prefer to partially accommodate. A steep inflation-contingent announcement then faces two related costs. Through expectations, the promised tightening reduces current inflation but can impose substantial output losses. Further, at $t = 1$, the announcement can also constrain the central bank to more tightening than it would choose when balancing inflation against output stabilization. Since the announced response cannot distinguish demand-driven inflation from cost-push inflation, greater cost-push uncertainty shifts the optimal balance toward a less aggressive common response. Optimal data dependence consequently falls (Figure 5a).

Interestingly, the shape of the response of vagueness is similar to that under heightened demand uncertainty, though driven by different mechanisms. Higher cost-push volatility directly increases the value of flexibility more than its cost, which on its own would make announcements more vague. However, the accompanying decline in data dependence works in the opposite direction, lowering the marginal value of that flexibility. For small levels of σ_u , this decline in α is sharp and weakly dominates the direct effect, so optimal vagueness initially falls. As σ_u rises further, the decline in data dependence flattens out, weakening the counteracting force. The direct effect then prevails, and vagueness increases again. Thus, while both sources of uncertainty produce a similar non-monotonic pattern in γ , the mechanisms behind it are different. In the demand case, vagueness eventually rises because α increases; in the supply case, it does so because α stops falling so sharply (Figure 5b).

Taken together, high supply uncertainty leads the central bank to adopt a **less commitment based communication strategy**, dialing back data dependence and offering less explicit guidance about future rates. The optimal response is to step back, allowing more discretion and reducing the extent to which communication links future policy to specific inflation-related data releases.

Together, these results highlight that not all uncertainty calls for the same communication strategy. While very high uncertainty—whether from demand or supply—increase the benefit of vague announcements that preserve flexibility to react, the optimal degree of data dependence (inflation-anchoring) depends crucially on the source of uncertainty. When uncertainty is mainly demand-driven, the central bank should underscore the sensitivity of its future decisions to incoming inflation data. By contrast, when uncertainty is primarily supply-driven, optimal communication calls for downplaying the relevance of inflation in future policy.

A second lesson also emerges from these results. Data dependence and vagueness cannot

Cost-push uncertainty and the design of guidance




Figure 5: Higher cost-push uncertainty (σ_u) lowers the optimal responsiveness of the announced range to inflation. Optimal vagueness, on the other hand, initially falls and then rises. The latter ($2\gamma^*$) is measured in annualized basis points. The standard deviation is measured in quarterly percentage points. The plotted range is 0.5–2 times the baseline standard deviation, the dotted line marks the baseline. All other parameters are held at their baseline values.

be treated in isolation. If data dependence were held fixed at a non-optimal level, or moved in the opposite direction, the optimal degree of vagueness in a given economic environment would differ substantially. The reverse effect is much weaker locally: near the baseline optimum, optimal data dependence changes little in response to small deviations in vagueness. This asymmetry is important for interpretation. Once signals are endogenous, communication design must be understood as a joint problem: data dependence determines how valuable flexibility becomes and, consequently, how vague optimal announcements should be.

5.3 The cost of strict rules

This section revisits the classic debate on rules versus discretion in monetary policy through the lens of the model. Since Milton Friedman’s advocacy for exogenous policy rules—such as his proposal of a constant growth rate for the money supply (Friedman, 1960)—the question of whether policy should follow predetermined rules or retain discretionary flexibility has been central to macroeconomics. The seminal contribution of Kydland and Prescott (1977) formalized the case against discretion: tying the hands of the policymaker can improve outcomes by preventing it from deviation from what was previously optimal after private expectations have been formed. Subsequent work, notably Taylor (1993) and the literature it inspired, gave this debate a concrete application through simple monetary

policy rules that respond to inflation deviations from target and the output gap. Despite this, central bankers have often defended discretion, emphasizing the need for flexibility amid evolving data and unforeseen disturbances (see, e.g., [Blinder \(1998\)](#), [Greenspan \(2004\)](#)).

Standard macroeconomic frameworks render this tension largely moot, as their structure directly implies a strict preference for rules: commitment —interpreted as the ability to announce and credibly implement a policy rule that private agents internalize in their expectations— generically yields superior outcomes relative to discretionary policy. The present model allows revisiting this question in a richer way. Because feasible announcements are not fully state-contingent, the policymaker’s desire for flexibility emerges endogenously, while the wary forward-looking private sector benefits from precision. In this sense, the proposed framework bridges the academic case for commitment with the practitioner’s concerns and preference for discretionary judgment. The exercise below measures the stabilization cost of strict rules: announcements that prescribe a unique future rate for each realization of the conditionable signal ($\gamma = 0$), leaving no discretion to respond to additional information. The benchmark is optimally vague guidance.

While stylized, the model with inflation-anchored announcements is particularly well suited for this purpose: it embeds the key elements that make both rules and discretion valuable in the now standard framework for studying monetary policy. The optimal announcement balances the stabilizing benefits of commitment against the value of retaining flexibility to respond to information beyond its conditioning variable, yielding a middle ground between strict commitment and full discretion.

In the baseline parameterization, a regime of zero discretion raises expected losses by 7.13% relative to the optimally vague benchmark, suggesting that the “cost” of rules may be substantial. While this figure should not be taken literally given the stylized nature of the model, it underscores the importance of flexibility and motivates further quantitative explorations.

Beyond the baseline magnitude, it is informative to examine how the stabilization cost of rules depends on structural characteristics of the economy. The next exercise explores this dependence, identifying the environments in which preserving some discretion is most valuable. In what follows, I focus on how the cost of rules varies as the slope of the Phillips curve κ and the weight on inflation λ change.

Weight on inflation stabilization. [Figure 6](#) shows a U-shaped relationship. This follows from the fact that at very low inflation weights, the central bank uses the misalignment

strategy described in the exogenous-signal model.²⁶ Concretely, it communicates vaguely—thus leaving substantial future flexibility—while shifting its midpoint enough to get the private sector to almost always select the same worst rate. The current rate offsets most of the resulting predictable effect on output, leaving an average inflation distortion that is not too costly given the weight on inflation is low. As inflation receives more weight, this strategy becomes less attractive and the selected policies become more precise and more closely centered, thus lowering the cost of rules.

Beyond this initial decline, a greater concern for inflation (higher λ) raises optimal data dependence, as the feedback loop it induces via policy expectations dampens inflation volatility substantially. This, in turn, increases the marginal value of vagueness in terms of $t = 1$ flexibility (see the discussion of demand uncertainty above). Over the high- λ range (i.e., beyond the “misalignment” zone), the response of data dependence greatly amplifies the increase in optimal vagueness.

Since there is a wide range of both empirically and normatively relevant values for λ (from the pure “dual mandate” weights to the baseline welfare-grounded calibration), understanding how communication design depends on this parameter is important for interpreting the model’s policy lessons.

²⁶Average misalignment is measured by $\alpha \mathbb{E}[\pi_0] + b - \bar{r}^n$. With inflation-anchored announcements, the former term is a function of the chosen policy, and thus the extent to which the announcement is aligned on average depends on all its dimensions, not only on b as in the (normalized) exogenous signal baseline.

The importance of inflation and the cost of strict rules




Figure 6: The stabilization cost of fully precise guidance initially falls and then rises as the common weight on inflation stabilization (λ) increases. The figure reports the percentage increase in expected policy loss under fully precise guidance ($\gamma = 0$), relative to optimal guidance, with the remaining policy choices reoptimized. The inflation weight varies jointly for the central bank and the private sector. All other parameters are held at their baseline values.

The slope of the Phillips curve. See Figure 7 for the effect of the Phillips curve slope on the cost of strict rules. I consider values of κ in $[0.005, 0.048]$, extending the half-to-twice-baseline range downward to include flatter estimates in the literature.²⁷ When the slope is very small, the consequently flat Phillips curve makes the cost of the inflation distortion associated with vague and shifted announcements described above relatively low. As the curve steepens, the current stabilization costs of this strategy rise, and the central bank moves to a communication strategy with more precise announcements, thus contributing to the initial decline in the cost of strict rules. Beyond this initial decline, two forces raise the value of ex post flexibility as inflation becomes more sensitive to real activity in both periods.

²⁷A half to twice the baseline value gives $[0.012, 0.048]$. The lower endpoint additionally accommodates the flatter estimates in Hazell et al. (2022): 0.0062 for the full sample and 0.0050–0.0055 after 1990. While these estimates relate (nontradeable) inflation to unemployment, I use them as illustrative output-gap benchmarks under the authors’ model mapping between employment and output.

First, at $t = 0$, the higher responsiveness of inflation to shocks makes expectations—and hence current inflation—more volatile. To stabilize these fluctuations, the central bank optimally increases data dependence, as the feedback it creates through beliefs helps dampen the sensitivity of inflation to shocks. Near the baseline, this rise in data dependence feeds back into the design of announcements by strengthening the marginal $t = 1$ flexibility value of vagueness (see the discussion of the comparative static of the optimal policy with respect to changes in demand uncertainty above), pushing optimal vagueness upward.

Second, at $t = 1$, the direct payoff to ex-post adjustment increases: a larger κ raises the loss from a given deviation of the future rate from its preferred value.²⁸ Together, these channels help explain why the stabilization cost of strict rules increases with the slope of the Phillips curve over most of the range considered.




Figure 7: The stabilization cost of fully precise guidance initially falls and then rises as the Phillips curve steepens (κ increases), holding the weight on inflation stabilization fixed. The figure reports the percentage increase in expected policy loss under fully precise guidance ($\gamma = 0$), relative to optimal guidance, with the remaining policy choices reoptimized. The range is $\kappa \in [0.005, 0.048]$, the dotted line marks the baseline. All other parameters are held at their baseline values.

²⁸Admittedly, this sort of “mechanical” effect of a change in κ would suggest adjusting the weight λ as κ varies (indeed, the mapping in the standard New Keynesian welfare benchmark features $\lambda \propto 1/\kappa$). However, here λ is held fixed to isolate the direct effect of κ on optimal communication design.

5.4 Optimal guidance vs. best forecast

The model also provides a natural environment to study the distinction between two well-known forms of central bank communication: *Odyssean* guidance, which commits to a future course of action, and *Delphic* guidance, which conveys the central bank’s best forecast of its future policy given current information (Campbell et al., 2012).

In the present framework, Odyssean guidance takes the form of a commitment to a range of future policy rates. A natural question is how similar the responsiveness of this range to inflation should be to a best-rate-forecast coefficient on inflation. While current inflation helps predict which rate will be appropriate tomorrow, the extent to which the announcement itself is anchored to inflation shapes expectations and consequently current economic outcomes. Optimal data dependence balances these stabilization effects against the cost of restricting future policy choices.

To quantify this distinction, I compare optimal data dependence with the forecasting benchmark. Concretely, I define the Delphic benchmark as the slope of the best linear projection of the central bank’s unconstrained preferred future rate on current inflation,

$$\alpha^d := \frac{\text{Cov}(\mathbb{E}[i_1^{\text{FI}}|\pi_0, y_0], \pi_0)}{\text{Var}(\pi_0)},$$

that is, the regression coefficient of the central bank’s unconstrained preferred $t = 1$ policy rate on today’s inflation.²⁹

This structure allows for a systematic analysis of when and why optimal data dependence diverges from the forecasting benchmark (purely Delphic guidance). In the baseline parameterization, optimal data dependence is approximately $\alpha^* = 2.85$, compared with a forecast coefficient of approximately $\alpha^d = 1.26$. The optimal announced range thus responds substantially more strongly to inflation than what a best forecast would suggest. This pattern illustrates the importance of the stabilization role of communication: making the announced range more responsive to inflation can counteract current inflationary pressures through expectations of future policy.

The following numerical exercises show how the Odyssean–Delphic gap responds to changes in demand and cost-push uncertainty. These exercises use the same 0.5–2-times-baseline uncertainty ranges as [Section 5.2](#).

²⁹Note that π_0 depends on the announcement (α, b, γ) through worst-case expectations. Hence, when computing α^d for different parameters below, the process for inflation is taken to be the one corresponding to the optimal announcement for each parameter configuration.

Demand uncertainty and the Odyssean–Delphic gap. Figure 8 plots optimal data dependence α^* and the Delphic projection coefficient α^d as functions of demand uncertainty σ_r .

Two features stand out. First, α^* is above α^d throughout the entire plotted range. Second, the gap between them widens as σ_r increases. The former reflects the two roles of data dependence. Inflation helps predict the central bank’s preferred future rate, and aligning the announcement with that rate helps decrease the expected cost from having to deviate from the ex post preferred rate. The forecasting (Delphic) benchmark α^d captures this role. However, optimal guidance also takes into account the contemporaneous stabilization benefits of greater responsiveness to inflation. In these exercises, this additional motive favors stronger data dependence, helping explain why $\alpha^* > \alpha^d$. The fact that the gap widens as σ_r increases, on the other hand, is explained by the fact that the stabilization consequences of high data dependence are particularly salient for demand shocks: a positive demand shock raises both output and inflation, while the prospect of higher future policy rates restrains current aggregate demand and thus allows the central bank to dampen the rise in both.




Figure 8: Higher demand uncertainty (σ_r) widens the gap between the optimal responsiveness of the announced range to inflation (α^*) and the forecast coefficient (α^d). The forecast benchmark is the best linear projection of the central bank’s unconstrained preferred future rate on current inflation, evaluated in the equilibrium induced by optimal guidance. The standard deviation is measured in quarterly percentage points. The plotted range is 0.5–2 times the baseline standard deviation, the dotted line marks the baseline. All other parameters are held at their baseline values.

Cost-push uncertainty and the Odyssean–Delphic gap. Figure 9 shows that while α^* remains above α^d throughout (the stabilization motive highlighted above does not disappear), the gap between them narrows as supply uncertainty σ_u increases. In this calibration, cost-push shocks cause more current inflation relative to the change in the preferred future rate than demand shocks do. This follows from a standard feature of the New Keynesian model: the preferred policy partially accommodates cost-push inflation to limit output losses, while counteracting demand pressures helps stabilize both objectives (“divine coincidence”). As cost-push shocks account for more of inflation’s variation, a given inflation increase is thus associated with a smaller preferred rate adjustment, lowering the forecasting coefficient. The same trade-off discourages committing to an announcement aggressively anchored to inflation. Furthermore, the narrowing gap reflects that stabilization considerations lead optimal guidance to incorporate a smaller additional response to inflation beyond the forecasting

benchmark. In this sense, the gap provides an (imperfect) indicator of the stabilization motive's influence on the chosen rule, reflecting the balance between current stabilization and future policy flexibility.

Cost-push uncertainty and the Odyssean–Delphic gap.




Figure 9: Higher cost-push uncertainty (σ_u) lowers both coefficients and narrows the gap between optimal data dependence (α^*) and the forecast coefficient (α^d). Data dependence measures the responsiveness of the announced range to inflation. The forecast benchmark is the best linear projection of the central bank's unconstrained preferred future rate on current inflation, evaluated in the equilibrium induced by optimal guidance. The standard deviation is measured in quarterly percentage points. The plotted range is 0.5–2 times the baseline standard deviation, the dotted line marks the baseline. All other parameters are held at their baseline values.

Taken together, these results support interpreting the gap between optimal data dependence and the forecasting benchmark—between Odyssean and Delphic guidance—as an imperfect indicator of the stabilization motive in optimal communication, whose importance varies with the source of uncertainty. When uncertainty is primarily demand-driven, forward guidance should emphasize a stronger, data-dependent response that harnesses the stabilizing power of expectations. When uncertainty is mainly supply-driven, by contrast, a strong inflation-contingent commitment risks forcing the central bank to suppress inflation pressures at an excessive output cost. Optimal guidance thus allows more accommodation

through a smaller response to inflation, closer to the forecasting benchmark.

6 Conclusion

This paper develops a framework to study optimal central bank communication when feasible announcements are not fully state-contingent and private agents are wary of policy mistakes. Within this setting, optimal communication is shown to take the form of signal-contingent intervals determined by three dimensions: how strongly to anchor future policy on signals that only partially capture the relevant state, how vague to make its commitments and where to center its announcements. When announcements are anchored to an endogenous variable, the choice of data dependence takes on additional significance: it governs whether the feedback between expectations and inflation is stabilizing or destabilizing, and thus whether there are multiple equilibria.

With inflation as the conditionable variable (arguably the most salient and policy-relevant anchor in practice), the analysis delivers three main lessons. First, optimal communication responds asymmetrically to the source of uncertainty: while high uncertainty calls for vague announcements regardless of whether it is about demand or cost-push shocks, it calls for strong data dependence only when it is about demand shocks. Second, strict rule-like communication is costly: constraining announcements to be fully precise entails sizable additional stabilization losses under a standard calibration. Third, optimal forward guidance differs sharply from the best forecast of next-period’s optimal rate (Delphic guidance), as the central bank internalizes the stabilizing role of expectations. More broadly, the analysis underscores the importance of jointly designing how precise and how data-responsive announcements should be: their interaction is rich enough that comparative statics obtained by varying one dimension while holding the other fixed can reverse when all the dimensions of the announcement and the current rate are jointly optimized.

The analysis also suggests natural directions for future work. This paper’s framework takes commitments to be fully credible and binding, but in practice credibility evolves over time depending on the correspondence between past announcements and subsequent actions, and earlier announcements for distant horizons interact with the design of new announcements as those horizons approach. Understanding optimal communication in a dynamic setting featuring these considerations is a clear next step. Similarly, the model takes inflation as the conditionable variable, but in practice the choice of what to condition on is itself a policy decision. Questions such as whether it is better to anchor guidance on inflation, unemployment, or some combination thereof—and what form that anchor should take—can be addressed with relatively minor extensions of the present framework.

References

- Amador, M., Werning, I. and Angeletos, G.-M. (2006), ‘Commitment vs. flexibility’, *Econometrica* **74**(2), 365–396.
- Amodeo, F. (2025), Mind the gap: Disagreement and credible monetary policy. Working Paper.
- Angeletos, G.-M. and Huo, Z. (2021), ‘Myopia and Anchoring’, *American Economic Review* **111**(4), 1166–1200.
- Angeletos, G.-M. and Sastry, K. A. (2021), ‘Managing expectations: Instruments versus targets’, *The Quarterly Journal of Economics* **136**(4), 2467–2532.
- Athey, S., Atkeson, A. and Kehoe, P. J. (2005), ‘The optimal degree of discretion in monetary policy’, *Econometrica* **73**(5), 1431–1475.
- Backus, D., Ferriere, A. and Zin, S. (2015), ‘Risk and ambiguity in models of business cycles’, *Journal of Monetary Economics* **69**, 42–63.
- Baker, S. R., Bloom, N. and Davis, S. J. (2016), ‘Measuring Economic Policy Uncertainty’, *The Quarterly Journal of Economics* **131**(4), 1593–1636.
- Barkin, T. I. (2021), ‘Talking about outcomes’, Speech. Forecasters Club of New York, New York, NY.
- Bassetto, M. (2019), ‘Forward guidance: Communication, commitment, or both?’, *Journal of Monetary Economics* **108**.
- Bauer, M. D., Lakdawala, A. and Mueller, P. (2022), ‘Market-Based Monetary Policy Uncertainty’, *Economic Journal* **132**(644), 1290–1308.
- Benigno, P. and Woodford, M. (2005), ‘Inflation stabilization and welfare: The case of a distorted steady state’, *Journal of the European Economic Association* **3**(6), 1185–1236.
- Bernanke, B. S. (2020), ‘The new tools of monetary policy’, *American Economic Review* **110**(4), 943–83.
- Bianchi, F., Ilut, C. L. and Schneider, M. (2018), ‘Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle’, *The Review of Economic Studies* **85**(2), 810–854.

- Blinder, A. S. (1998), *Central Banking in Theory and Practice*, The Lionel Robbins Lectures, MIT Press, Cambridge, MA.
- Bundick, B., Smith, A. L. and Van der Meer, L. (2024), ‘Introducing the Kansas City Fed’s measure of policy rate uncertainty (KC PRU)’, *Economic Review* **109**(7).
- Caballero, R. J. and Simsek, A. (2022), ‘Monetary policy with opinionated markets’, *American Economic Review* **112**(7), 2353–2392.
- Campbell, J. R., Evans, C. L., Fisher, J. D. M. and Justiniano, A. (2012), ‘Macroeconomic Effects of Federal Reserve Forward Guidance’, *Brookings Papers on Economic Activity* **43**, 1–80.
- Cho, S., Del Negro, M., Diagne, I., Gundam, P., Lee, D. and Pacula, B. (2024), ‘The New York Fed DSGE model forecast—december 2024’, Liberty Street Economics, Federal Reserve Bank of New York. December 20.
- Cieslak, A., Malamud, S. and Schrimpf, A. (2020), ‘Policy Announcement Design’.
- Clarida, R., Galí, J. and Gertler, M. (2000), ‘Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory’, *Quarterly Journal of Economics* **115**(1).
- Cole, S. J. and Martínez-García, E. (2023), ‘The effect of central bank credibility on forward guidance in an estimated New Keynesian model’, *Macroeconomic Dynamics* **27**(2), 532–570.
- Del Negro, M., Giannoni, M. P. and Patterson, C. (2023), ‘The Forward Guidance Puzzle’, *Journal of Political Economy Macroeconomics* **1**(1), 43–79.
- Eggertsson, G. and Woodford, M. (2003), ‘The zero bound on interest rates and optimal monetary policy’, *Brookings Papers on Economic Activity*.
- Farhi, E. and Werning, I. (2019), ‘Monetary policy, bounded rationality, and incomplete markets’, *American Economic Review* **109**(11), 3887–3928.
- Friedman, M. (1960), *A Program for Monetary Stability*, Fordham University Press.
- Gajdos, T., Hayashi, T., Tallon, J.-M. and Vergnaud, J.-C. (2008), ‘Attitude toward imprecise information’, *Journal of Economic Theory* **140**(1), 27–65.
- 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.
- Gáti, L. (2023), ‘Talking over time—dynamic central bank communication’, *Journal of Money, Credit and Banking* **55**(5), 1147–1176.
- Gilboa, I. and Schmeidler, D. (1989), ‘Maxmin expected utility with non-unique prior’, *Journal of Mathematical Economics* **18**(2), 141–153.
- Greenspan, A. (2004), ‘Risk and Uncertainty in Monetary Policy’, *American Economic Review* **94**(2), 33–40.
- Gurkaynak, R. S., Sack, B. and Swanson, E. T. (2005), ‘Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements’, *International Journal of Central Banking* **1**(1), 55–93.
- Hazell, J., Herreño, J., Nakamura, E. and Steinsson, J. (2022), ‘The Slope of the Phillips curve: Evidence from U.S. States’, *The Quarterly Journal of Economics* **137**(3), 1299–1344.
- Holmstrom, B. (1984), On the Theory of Delegation, in M. Boyer and R. E. Kihlstrom, eds, ‘Bayesian Models in Economic Theory’.
- Ilut, C. L. and Schneider, M. (2014), ‘Ambiguous Business Cycles’, *American Economic Review* **104**(8), 2368–2399.
- Ilut, C. and Saijo, H. (2021), ‘Learning, confidence, and business cycles’, *Journal of Monetary Economics* **117**, 354–376.
- Kocherlakota, N. (2016), ‘Rules versus Discretion: A Reconsideration’, *Brookings Papers on Economic Activity* **47**(2), 1–55.
- Kydland, F. E. and Prescott, E. C. (1977), ‘Rules Rather than Discretion: The Inconsistency of Optimal Plans’, *Journal of Political Economy* .
- Lakdawala, A., Seo, J. and Shim, M. (2026), ‘On the optimal precision of central bank communication’.
- Masolo, R. M. and Monti, F. (2021), ‘Ambiguity, Monetary Policy and Trend Inflation’, *Journal of the European Economic Association* **19**(2), 839–871.
- McKay, A., Nakamura, E. and Steinsson, J. (2016), ‘The power of forward guidance revisited’, *American Economic Review* **106**(10).

- Michelacci, C. and Paciello, L. (2020), ‘Ambiguous Policy Announcements’, *The Review of Economic Studies* **87**(5), 2356–2398.
- Moscarini, G. (2007), ‘Competence implies credibility’, *American Economic Review* **97**(1).
- Romer, C. D. and Romer, D. H. (2024), ‘Did the Federal Reserve’s 2020 Policy Framework Limit Its Response to Inflation? Evidence and Implications for the Framework Review’, *Brookings Papers on Economic Activity* **55**(2), 59–87.
- Rotemberg, J. J. and Woodford, M. (1997), ‘An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy’, *NBER Macroeconomics Annual* **12**, 297–346.
- Sastry, K. A. (2026), ‘Disagreement about monetary policy’, *American Economic Journal: Macroeconomics* **18**(2), 486–522.
- Stein, J. C. (1989), ‘Cheap talk and the Fed: A theory of imprecise policy announcements’, *American Economic Review* **79**(1).
- Taylor, J. B. (1993), ‘Discretion versus policy rules in practice’, *Carnegie-Rochester Conference Series on Public Policy* **39**, 195–214.
- Waller, C. J. (2026), ‘Two thoughts on the transmission of monetary policy’, Speech. Challenges for Monetary Policy Transmission in a Changing World, Bank of Italy, Rome, Italy, July 6.
- Warsh, K. (2026), ‘In our time’, Speech. Jackson Hole Economic Policy Symposium, Jackson Hole, Wyoming, August 28.
- Werning, I. (2012), Managing a liquidity trap: Monetary and fiscal policy, Technical Report 17344, NBER.
- Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.
- Woodford, M. (2005), ‘Central bank communication and policy effectiveness’, *Proceedings - Economic Policy Symposium - Jackson Hole* (August), 399–474. Federal Reserve Bank of Kansas City.
- Yellen, J. L. (2013), ‘Challenges confronting monetary policy’, Speech. National Association for Business Economics Policy Conference, Washington, D.C.
- Yellen, J. L. (2015), ‘Letter to House Speaker Paul Ryan and House Minority Leader Nancy Pelosi’. November 16, 2015.

Appendix

A	Technical Definitions	54
B	Proofs of Main Results	54
B.1	Proof of Lemma 1	54
B.2	Proof of Proposition 1	55
B.3	Proof of Proposition 2	56
B.4	Proof of Lemma 2	58
B.5	Proof of Proposition 3	59
B.6	Proof of Proposition 4	70
B.7	Proof of Proposition 5	72
C	Second-Order Conditions: Exogenous Signals	85
D	Proofs of Additional Technical Lemmas	90
D.1	Proof of Lemma 3	90
D.2	Proof of Lemma 4	91
D.3	Proof of Lemma 5	92
D.4	Proof of Lemma 6	97
D.5	Proof of Lemma 7	98

A Technical Definitions

Definition 2. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting the jointly Gaussian random variables defined above and let $\mathcal{K}(\mathbb{R})$ be the set of non-empty closed subsets of $\mathbb{R}$ equipped with the Effros σ -algebra $\mathcal{E}$. An **announcement** is a measurable map $A : (\mathbb{R}, \mathcal{B}(\mathbb{R})) \rightarrow (\mathcal{K}(\mathbb{R}), \mathcal{E})$. Given a realization of the signal $\tilde{s}$, the corresponding feasible set is denoted by $A(\tilde{s})$.

Remark 1. If $A(\tilde{s}) = [L(\tilde{s}), U(\tilde{s})]$ for Borel-measurable $L, U : \mathbb{R} \rightarrow \mathbb{R}$ and $L \leq U$, then A is Effros-measurable.

Definition 3. Fix $s_0 = (r_0^n, u_0)$. Let $F(\cdot \mid s_0)$ denote the conditional law of $(r_1^n, u_1, \tilde{\epsilon}_1)$, and define the measurable map

$$g(r_1^n, u_1, \tilde{\epsilon}_1) := \tilde{\omega}_r r_1^n + \tilde{\omega}_u u_1 + \tilde{\epsilon}_1.$$

Write $F_{(r_1^n, u_1, \tilde{s}_1)}(\cdot \mid s_0)$ for the pushforward of $F(\cdot \mid s_0)$ under $(r_1^n, u_1, \tilde{\epsilon}_1) \mapsto (r_1^n, u_1, g(\cdot))$, and let $F_{\tilde{s}_1}(\cdot \mid s_0)$ denote its marginal on $\tilde{s}_1$. Given an announcement $A : \mathbb{R} \rightarrow \mathcal{K}(\mathbb{R})$, **the set of admissible private-sector beliefs** $\mathcal{P}(s_0)$ is the set of probability measures on $\mathbb{R}^4$ over $(i_1, r_1^n, u_1, \tilde{s}_1)$ such that there exists a measurable stochastic kernel K s.t.

- $p(B \times C) = \int_C K(B \mid \tilde{s}_1) F_{(r_1^n, u_1, \tilde{s}_1)}(dr_1^n, du_1, d\tilde{s}_1 \mid s_0) \quad \forall B \in \mathcal{B}(\mathbb{R}), C \in \mathcal{B}(\mathbb{R}^3),$
- $\text{supp } K(\cdot \mid \tilde{s}_1) \subseteq \mathcal{A}_1(\tilde{s}_1) \quad \text{for } F_{\tilde{s}_1}(\cdot \mid s_0)\text{-a.e. } \tilde{s}_1.$

B Proofs of Main Results

B.1 Proof of **Lemma 1**

Note that i_1 affects only $t = 1$ outcomes and, further, for $t \geq 2$, $y_t = \pi_t = 0$. Thus, at $t = 1$ the CB faces a standard New Keynesian one-period optimal policy problem save for the interest rate constraint. As usual, $t = 1$ losses can be written as

$$W_1(i_1 - i_1^{\text{FI}}(\lambda_{cb}))^2 + \text{t.i.p.},$$

where $W_1 := \beta\sigma^2(1 + \lambda_{cb}\kappa^2)$ and, as usual, t.i.p. stands for terms independent of policy. Denote by $\mathcal{I}_1 = \sigma(\tilde{s}_1, \hat{s}_1, r_0^n, u_0)$ the information available at time $t = 1$ to the central bank. For any $\mathcal{I}_1$ -measurable i ,

$$\mathbb{E}[(i_1 - i_1^{\text{FI}}(\lambda_{cb}))^2 \mid \mathcal{I}_1] = (i_1 - i_1^{\text{unc}})^2 + \text{Var}(i_1^{\text{FI}}(\lambda_{cb}) \mid \mathcal{I}_1),$$

where $i_1^{\text{unc}} = \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid \mathcal{I}_1]$ is the unconstrained optimal rate. Because the second term is independent of i_1 , minimizing over $i_1 \in A(\tilde{s}_1)$ entails minimizing the squared distance to i_1^{unc} , hence the implemented rate is the projection of i_1^{unc} into $A(\tilde{s}_1)$,

$$i_1^*(s_0, \tilde{s}_1, \hat{s}_1) = \arg \min_{i_1 \in A_1(\tilde{s}_1)} \|i_1 - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1, \hat{s}_1]\|.$$

B.2 Proof of Proposition 1

Let i_1^e denote the private sector's ex-ante mean policy rate (naturally this is a function of the realizations of r_0^n, u_0 and the central bank announcement, but we do not need to be explicit about those at this point). Substituting the IS and NKPC at $t = 1$ into the IS at $t = 0$ and using that $E_t \pi_{t+k} = \mathbb{E}_t y_{t+k} = 0$ for all $t \geq 1$ and $k \geq 1$ we have that for an arbitrary period $t = 0$ announcement $\mathcal{A}_1$

$$\begin{aligned} y_0 &= \tilde{\mathbb{E}}_0[y_1 | \mathcal{A}_1] - \sigma(i_0 - \tilde{\mathbb{E}}_0[\pi_1 | \mathcal{A}_1] - r_0^n) \\ &= \tilde{\mathbb{E}}_0[y_1 | \mathcal{A}_1] - \sigma(i_0 - \tilde{\mathbb{E}}_0[\kappa y_1 + \beta \tilde{\mathbb{E}}_1[\pi_2 | \mathcal{A}_2] + u_1 | \mathcal{A}_1] - r_0^n) \\ &= (1 + \sigma\kappa)\tilde{\mathbb{E}}_0[y_1 | \mathcal{A}_1] - \sigma(i_0 - \beta \tilde{\mathbb{E}}_0[\tilde{\mathbb{E}}_1[\pi_2 | \mathcal{A}_2] | \mathcal{A}_1] - \rho_u u_0 - r_0^n) \\ &= (1 + \sigma\kappa)\tilde{\mathbb{E}}_0[y_1 | \mathcal{A}_1] - \sigma i_0 + \sigma r_0^n + \sigma \rho_u u_0 \\ &= (1 + \sigma\kappa)\tilde{\mathbb{E}}_0[\tilde{\mathbb{E}}_1[y_2 | \mathcal{A}_2] - \sigma(i_1 - \tilde{\mathbb{E}}_1[\pi_2] - r_1^n) | \mathcal{A}_1] - \sigma i_0 + \sigma r_0^n + \sigma \rho_u u_0 \\ &= -\sigma(1 + \sigma\kappa)\tilde{\mathbb{E}}_0[i_1 - r_1^n | \mathcal{A}_1] - \sigma i_0 + \sigma r_0^n + \sigma \rho_u u_0 \\ &= \sigma(1 + (1 + \sigma\kappa)\rho_r)r_0^n + \sigma \rho_u u_0 - \sigma i_0 - \sigma(1 + \sigma\kappa)i_1^e + C_y \end{aligned}$$

Similarly for the NKPC,

$$\begin{aligned} \pi_0 &= \kappa y_0 + \beta \tilde{\mathbb{E}}_0[\pi_1 | \mathcal{A}_1] + u_0 \\ &= \kappa y_0 + \beta \tilde{\mathbb{E}}_0[\kappa y_1 + \beta \tilde{\mathbb{E}}_1[\pi_2 | \mathcal{A}_2] + u_1 | \mathcal{A}_1] + u_0 \\ &= \kappa(y_0 + \beta \tilde{\mathbb{E}}_0[y_1 | \mathcal{A}_1]) + (1 + \beta \rho_u)u_0. \\ &= \kappa(y_0 + \beta \tilde{\mathbb{E}}_0[\tilde{\mathbb{E}}_1[y_2] - \sigma(i_1 - \tilde{\mathbb{E}}_1[\pi_2] - r_1^n) | \mathcal{A}_1]) + (1 + \beta \rho_u)u_0 \\ &= \kappa(y_0 - \beta \sigma(\tilde{\mathbb{E}}_0[i_1 | \mathcal{A}_1] - c - \rho_r r_0^n)) + (1 + \beta \rho_u)u_0 \\ &= \sigma\kappa(1 + (1 + \sigma\kappa + \beta)\rho_r)r_0^n + (1 + (\sigma\kappa + \beta)\rho_u)u_0 - \sigma\kappa i_0 - \sigma\kappa(1 + \sigma\kappa + \beta)i_1^e + C_\pi. \end{aligned}$$

Thus, defining

$$\begin{aligned} A_y^r &:= \sigma(1 + (1 + \sigma\kappa)\rho_r), & A_y^u &:= \sigma\rho_u, & A_y^{i^e} &:= -\sigma(1 + \sigma\kappa), \\ A_\pi^r &:= \sigma\kappa(1 + (1 + \sigma\kappa + \beta)\rho_r), & A_\pi^u &:= 1 + (\sigma\kappa + \beta)\rho_u, & A_\pi^{i^e} &:= -\sigma\kappa(1 + \sigma\kappa + \beta). \end{aligned}$$

we have

$$\begin{aligned} y_0 &= C_y + A_y^r r_0^n + A_y^u u_0 - \sigma i_0 + A_y^{i^e} i_1^e, \\ \pi_0 &= C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} i_1^e. \end{aligned}$$

Define

$$\delta_r := A_y^r + \lambda_{cb} \kappa A_\pi^r, \quad \delta_u := A_y^u + \lambda_{cb} \kappa A_\pi^u, \quad \delta_{i^e} := A_y^{i^e} + \lambda_{cb} \kappa A_\pi^{i^e}, \quad \delta_c := C_y + \lambda_{cb} \kappa C_\pi.$$

Let $R := \delta_r r_0^n + \delta_u u_0 + \delta_{i^e} i_1^e + \delta_c$ and $C := (C_y + A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e)^2 + \lambda_{cb} (C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e)^2$. Then, the expected period-0 loss under the central bank's beliefs is

$$\mathbb{E}[y_0^2 + \lambda_{cb} \pi_0^2] = \sigma^2 (1 + \lambda_{cb} \kappa^2) i_0^2 - 2\sigma i_0 \mathbb{E}[R] + \mathbb{E}[C].$$

As this is clearly a strictly convex function of i_0 , minimizing over it yields

$$i_0^* = \frac{\mathbb{E}[R]}{\sigma(1 + \lambda_{cb} \kappa^2)} \quad \text{and} \quad \min_{i_0} \mathbb{E}[y_0^2 + \lambda_{cb} \pi_0^2] = \mathbb{E}[C] - \frac{\mathbb{E}[R]^2}{1 + \lambda_{cb} \kappa^2}.$$

Thus, as $\mathbb{E}[u_0] = 0$, the minimized ex-ante period-0 loss is the functional

$$\begin{aligned} \Psi(i_1^e) &:= \mathbb{E} \left[(C_y + A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e(r_0^n, u_0))^2 + \lambda_{cb} (C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e(r_0^n, u_0))^2 \right] \\ &\quad - \frac{(\delta_r r_0^n + \delta_c + \delta_{i^e} \mathbb{E}[i_1^e(r_0^n, u_0)])^2}{1 + \lambda_{cb} \kappa^2}. \end{aligned}$$

B.3 Proof of Proposition 2

The proof consists of two steps. First, it is shown that intervals weakly dominate any other shape of the announcement set. Second, it is proved that equivarence implies feasible announcements have a constant interval width (i.e., independent of $\tilde{s}_1$) and an interval's center that is an affine function of $\tilde{s}_1$.

- (i) *The optimal announcement is an interval.* Denoting by i_1^* the rate chosen at period $t = 1$ by the central bank and by i_1^{unc} the rate chosen if unconstrained by the previous announcement, let $\Delta_i := i_1^* - i_1^{unc}$. Substituting the expressions for the private sector behavior at $t = 1$ and using that $E_t \pi_{t+k} = \mathbb{E}_t y_{t+k} = 0$ for all $t \geq 1$ and $k \geq 1$, we can write expected $t = 1$ losses as

$$\sigma^2 \mathbb{E} \left((i_1^* - r_1^n)^2 + \lambda_{cb} \kappa^2 \left(i_1^* - \left(r_1^n + \frac{u_1}{\sigma \kappa} \right) \right)^2 \right)$$

$$\begin{aligned}
&= \sigma^2 \left(1 + \lambda_{cb} \kappa^2\right) \mathbb{E}[\Delta_i^2] + 2\sigma^2 E_0 \left[\Delta_i \left(i_1^{unc} - r_1^n + \lambda_{cb} \kappa^2 \left(i_1^{unc} - \left(r_1^n + \frac{u_1}{\sigma \kappa} \right) \right) \right) \right] \\
&\quad + \sigma^2 E_0 \left[\left(i_1^{unc} - r_1^n \right)^2 + \lambda_{cb} \kappa^2 \left(i_1^{unc} - \left(r_1^n + \frac{u_1}{\sigma \kappa} \right) \right)^2 \right].
\end{aligned}$$

The second term equals zero by the Law of Iterated Expectations: Δ_i is measurable with respect to the information available at $t = 1$, and the conditional expectation of the term in parenthesis is zero by definition of i_1^{unc} . The last term is independent of communication policy and thus can be disregarded for the purposes of solving for the optimal announcement. Thus, we can write

$$\mathbb{E}[y_1^2 + \lambda_{cb} \pi_1^2] = \sigma^2(1 + \lambda_{cb} \kappa^2) \mathbb{E}[(i_1^* - i_1^{unc})^2] + t.i.c.p.$$

Naturally, i_1^* corresponds to the point in the feasible set that is closest to i_1^{unc} . Now, let $\mathcal{A}$ be an arbitrary announcement. Consider the modified announcement $\tilde{\mathcal{A}}$ where, for each signal realization $\tilde{s}_1$, $\tilde{A}_t(\tilde{s}_1) = \text{co } A_t(\tilde{s}_1)$ (i.e., we take the convex hull of every feasible set). Given the above, period $t = 1$ losses are given by the expected squared distance to the feasible set and thus weakly decrease under convexification (if the target random variable i_1^{unc} has full support and the convexification is non-trivial, it strictly decrease). On the other hand, $t = 0$ losses depend on the announced interval only through the belief it induces on the private sector. As the extent of data dependence α is unchanged, the modified announcement can only differentially affect beliefs by inducing a different worst-case rate. Because the private sector losses are a convex function of next period's interest rate and the feasible set is a non-empty compact set³⁰, its maximizer (the worst-case rate) is always on the boundary of the feasible set and thus the maximizer over A_t is the same as the maximizer over its convex hull $\tilde{A}_t$. Thus, $t = 0$ losses are unchanged. Hence, the modified announcement $\tilde{\mathcal{A}}$ weakly dominates $\mathcal{A}$. Because $A_t(\tilde{s}_1) \subseteq \mathbb{R}$, the convex hull of a feasible set is always an interval, thus finishing the argument.

- (ii) *The Optimal Interval Announcement has Affine Ends with Common Slope.* Let $A(x) = [L(x), U(x)]$ with $L(x) \leq U(x)$ be an interval announcement. Equivariance implies there exists $\alpha \in \mathbb{R}$ such that

$$[L(x + \delta), U(x + \delta)] = [L(x), U(x)] + \alpha\delta = [L(x) + \alpha\delta, U(x) + \alpha\delta]$$

³⁰As $A_t(\tilde{s}_1) \subseteq \mathbb{R}$, it is compact if and only if it is closed and bounded. It is closed by assumption and it's straightforward to verify it's never optimal for the central bank to pick an announcement that entails unbounded sets. Thus, it is without loss of generality to consider announcements whose feasible sets A_t are bounded for every $\tilde{s}_1$.

for all $x, \delta \in \mathbb{R}$. Hence,

$$L(x + \delta) = L(x) + \alpha\delta, \quad U(x + \delta) = U(x) + \alpha\delta.$$

Fix $x, y \in \mathbb{R}$ and take $\delta = y - x$. Then

$$L(y) - \alpha y = L(x) - \alpha x, \quad U(y) - \alpha y = U(x) - \alpha x.$$

Thus $d_L(x) := L(x) - \alpha x$ and $d_U(x) := U(x) - \alpha x$ are constant on $\mathbb{R}$; call the constants γ_L and γ_U . Thus,

$$L(x) = \alpha x + \gamma_L, \quad U(x) = \alpha x + \gamma_U,$$

and $A(x) = [\alpha x + \gamma_L, \alpha x + \gamma_U]$ for all x .

B.4 Proof of Lemma 2

The proof relies on the following lemmas. Their proofs can be found in [Appendix D](#).

Lemma 3. *There exist $a_{ps}, a_{cb} \in \mathbb{R}$ and affine maps a_{ps}^0, a_{cb}^0 such that*

$$\begin{aligned} \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid \tilde{s}_1, r_0^n, u_0] &= a_{ps}\tilde{s}_1 + a_{ps}^0(r_0^n, u_0), \\ \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid \tilde{s}_1, \hat{s}_1, r_0^n, u_0] &= a_{cb}\tilde{s}_1 + a_{cb}^0(r_0^n, u_0, \hat{s}_1). \end{aligned}$$

Lemma 4. *Consider an interval announcement with $L, U \in L^1$ such that $L(s) \leq U(s)$ for all $s \in \mathbb{R}$. For each realization s of $\tilde{s}_1$, the private sector's worst-case beliefs about i_1 place probability one on a maximizer of the expected quadratic loss conditional on $(s_0, \tilde{s}_1)$ over the compact set $[L(s), U(s)]$, i.e.,*

$$i_1^w(s, s_0) \in \arg \max_{a \in [L(s), U(s)]} \mathbb{E} \left[\left(a - i_1^{\text{FI}}(s_1, \lambda_{ps}) \right)^2 \mid s_0, \tilde{s}_1 = s \right].$$

In particular, letting $c(s) := \frac{L(s)+U(s)}{2}$ and $\gamma(s) := \frac{U(s)-L(s)}{2} \geq 0$,

$$i_1^w(s, s_0) = c(s) + \gamma(s) \text{sign} \left(c(s) - a_{ps}s - a_{ps}^0(s_0) \right)$$

with the convention that if $a_{ps}s + a_{ps}^0(s_0) = c(s)$ the upper endpoint is selected. Consequently, the private sector's expected $t = 1$ policy rate is

$$i_1^e(s_0) := \mathbb{E}[i_1^w \mid s_0] = \mathbb{E}[c(\tilde{s}_1) \mid s_0] + \mathbb{E} \left[\gamma(\tilde{s}_1) \text{sign} \left(c(\tilde{s}_1) - a_{ps}\tilde{s}_1 - a_{ps}^0(s_0) \right) \mid s_0 \right].$$

Proof of Lemma 2. Applying Lemma 4 to an interval of the form $[\alpha\tilde{s}_1 + b - \gamma, \alpha\tilde{s}_1 + b + \gamma]$, it's straightforward to get

$$i_1^e(s_0; \alpha, b, \gamma) = \alpha\mu_{\tilde{s}}(s_0) + b + \gamma\theta_{\alpha,b}(s_0),$$

where

$$\theta_{\alpha,b}(s_0) = \mathbb{E} \left[\text{sign} \left((\alpha - a_{ps})\tilde{s}_1 + b - a_{ps}^0(s_0) \right) \mid s_0 \right].$$

Let $Z := (\alpha - a_{ps})\tilde{s}_1 + b - a_{ps}^0(s_0)$ and first consider the case where $\alpha \neq a_{ps}$. Clearly,

$$Z \mid s_0 \sim N(\mu_Z, V_Z).$$

Thus,

$$\begin{aligned} \mathbb{E} [\text{sign}(Z) \mid s_0] &= \mathbb{P}(Z > 0) - \mathbb{P}(Z < 0) \\ &= 2\mathbb{P}(Z > 0) - 1 \\ &= 2\Phi\left(\frac{\mu_Z}{\sqrt{V_Z}}\right) - 1, \end{aligned}$$

where the probabilities used are conditional on s_0 . Further, note that

$$\mu_Z = (\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + b - a_{ps}^0(s_0)$$

and

$$V_Z = (\alpha - a_{ps})^2 \text{Var}(\tilde{s}_1 \mid s_0)$$

Finally, if $\alpha = a_{ps}$, then from the private sector's perspective (i.e., conditional on s_0) the numerator is deterministic: $Z = b - a_{ps}^0(s_0)$. The upper endpoint is selected when $Z \geq 0$ (where the equality is included by convention as stated), and the lower endpoint is selected otherwise. $\square$

B.5 Proof of Proposition 3

The proof relies on the following additional lemma. Its proof can be found in Appendix D.

Lemma 5. *Let $s_0 := (r_0^n, u_0)$ and*

$$\mathcal{D} := \{(\alpha, b, \gamma) \in \mathbb{R}^3 : \gamma > 0, \alpha \neq a_{ps}\}.$$

Consider the function

$$\mathcal{J}(\alpha, b, \gamma) := \Psi(\alpha, b, \gamma) + W_1 \mathbb{E} \left[(|z_{\alpha,b}| - \gamma)_+^2 \right], \quad z_{\alpha,b} = (\alpha - a_{cb})\tilde{s}_1 + (b - a_{cb}^0),$$

with all the constants as defined in [Proposition 3](#) and Ψ as defined in [Proposition 1](#) evaluated at i_1^e as defined in [Lemma 2](#). Then, $\mathcal{J}$ is continuously differentiable on $\mathcal{D}$ and its derivatives with respect to α , b , and γ are given by

$$\begin{aligned} & \mathbb{E} \left[G_{\alpha,b,\gamma}(s_0) \left(\mu_{\tilde{s}}(s_0) + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha}(s_0) \right) \right] + 2W_1 \mathbb{E} \left[(|z_{\alpha,b}(s_0, \tilde{s}_1, \hat{s}_1)| - \gamma)_+ \text{sign}(z_{\alpha,b}(s_0, \tilde{s}_1, \hat{s}_1)) \tilde{s}_1 \right], \\ & \mathbb{E} \left[G_{\alpha,b,\gamma}(s_0) \left(1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b}(s_0) \right) \right] + 2W_1 \mathbb{E} \left[(|z_{\alpha,b}(s_0, \tilde{s}_1, \hat{s}_1)| - \gamma)_+ \text{sign}(z_{\alpha,b}(s_0, \tilde{s}_1, \hat{s}_1)) \right], \\ & \mathbb{E} [G_{\alpha,b,\gamma}(s_0) \theta_{\alpha,b}(s_0)] - 2W_1 \mathbb{E} [(|z_{\alpha,b}(s_0, \tilde{s}_1, \hat{s}_1)| - \gamma)_+], \end{aligned}$$

where³¹

$$\begin{aligned} G_{\alpha,b,\gamma}(s_0) &:= 2A_y^{i^e} \left(C_y + A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e(s_0; \alpha, b, \gamma) \right) \\ &\quad + 2\lambda_{cb} A_\pi^{i^e} \left(C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e(s_0; \alpha, b, \gamma) \right) \\ &\quad - \frac{2\delta_{i^e}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{i^e} \mathbb{E}[i_1^e(s_0; \alpha, b, \gamma)]), \\ \theta_{\alpha,b}(s_0) &= 2\Phi(\nu(s_0)) - 1, \quad \nu(s_0) := \frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + b - a_{ps}^0(s_0)}{|\alpha - a_{ps}|\sigma_{\tilde{s}}}, \end{aligned}$$

$$\begin{aligned} \frac{\partial \theta_{\alpha,b}}{\partial b} &= \frac{2}{|\alpha - a_{ps}|\sigma_{\tilde{s}}} \phi(\nu), \\ \frac{\partial \theta_{\alpha,b}}{\partial \alpha} &= 2\phi(\nu(s_0)) \left(\frac{(b - a_{ps}^0(s_0)) \text{sign}(a_{ps} - \alpha)}{(\alpha - a_{ps})^2 \sigma_{\tilde{s}}} \right) \end{aligned}$$

and $\mu_{\tilde{s}}(s_0) := \mathbb{E}[\tilde{s}_1 \mid s_0]$. Moreover, for any (α, b) with $\alpha \neq a_{ps}$, all three partials extend continuously to $\gamma = 0$ (interpreting $\partial \mathcal{J} / \partial \gamma$ as the right derivative in that point).

Proof of [Proposition 3](#).

Notation and optimized current loss. Throughout the proof of the proposition, normalize $\bar{r}^n = 0$. To economize notation in the proof, write

$$S := \tilde{s}_1, \quad M := \mathbb{E}[S \mid s_0],$$

³¹The coefficients in $G_{\alpha,b,\gamma}$ below are defined in [Appendix B.2](#).

$$\begin{aligned}
T &:= \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, S, \hat{s}_1], & F &:= \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0], \\
\theta &:= \theta_{\alpha, b}, & I &:= \alpha M + b + \gamma \theta, \\
z &:= \alpha S + b - T, & e &:= \mathbb{E}[\theta], \quad m := \mathbb{E}[I] = b + \gamma e.
\end{aligned}$$

The variables S, M, T, F have mean zero, while $\mathbb{E}[z] = b$. Further, [Proposition 1](#) gives

$$i_0^* = -\chi m, \quad \chi := 1 + \sigma \kappa + \frac{\lambda_{cb} \kappa^2 \beta}{1 + \lambda_{cb} \kappa^2}.$$

Also define

$$\begin{aligned}
Y &:= A_y^r r_0^n + A_y^u u_0, & P &:= A_\pi^r r_0^n + A_\pi^u u_0, \\
H &:= A_y^{i^e} Y + \lambda_{cb} A_\pi^{i^e} P, & Q &:= (A_y^{i^e})^2 + \lambda_{cb} (A_\pi^{i^e})^2, \\
w &:= \frac{\lambda_{cb} \sigma^2 \kappa^2 \beta^2}{1 + \lambda_{cb} \kappa^2}, & K &:= \mathbb{E}[Y^2 + \lambda_{cb} P^2].
\end{aligned}$$

Substituting i_0^* into the current loss and using that $Q - \sigma^2(1 + \lambda_{cb} \kappa^2) \chi^2 = w$ yields

$$\Psi(\alpha, b, \gamma) = K + 2 \text{Cov}(H, I) + Q \text{Var}(I) + w m^2, \quad (10)$$

where $Q, w > 0$. The marginal effect of a (square-integrable) perturbation $I + \varepsilon h$ on this minimized loss is $\mathbb{E}[Gh]$, where

$$G := 2(H + Q(I - m) + w m), \quad (11)$$

which is the zero-mean specialization of $G_{\alpha, b, \gamma}$ in [Lemma 5](#) (note that it already accounts for the optimal adjustment of i_0).

The objective function and $t = 1$ losses. By [Lemma 1](#), the future rate is the projection of T onto $[\alpha S + b - \gamma, \alpha S + b + \gamma]$. Define

$$\ell := (|z| - \gamma)_+, \quad \mathcal{R} := \ell \text{sign}(z), \quad p := \mathbb{P}(|z| > \gamma).$$

Thus $\mathcal{R} = i_1^* - T$ and $\ell = |i_1^* - T|$. A conditional squared-loss decomposition around T gives that future losses can be written as $W_1 \mathbb{E}[\ell^2]$, where $W_1 = \beta \sigma^2(1 + \lambda_{cb} \kappa^2)$ plus a communication-policy independent term. The reduced objective function is then

$$\mathcal{J}(\alpha, b, \gamma) = \Psi(\alpha, b, \gamma) + W_1 \mathbb{E}[\ell^2], \quad (12)$$

as in the proposition statement.

Differentiability and first-order conditions. **Lemma 5** establishes continuous differentiability for $\gamma > 0$ and $\alpha \neq a_{ps}$, together with the continuous extension of the partial derivatives to $\gamma = 0$, using a right derivative in γ . On this domain, let $\theta_\alpha := \frac{\partial \theta_{\alpha,b}}{\partial \alpha}$ and $\theta_b := \frac{\partial \theta_{\alpha,b}}{\partial b}$. Since the worst-case selector $\theta_{\alpha,b}$ does not depend on vagueness,

$$I_\alpha = M + \gamma \theta_\alpha, \quad I_b = 1 + \gamma \theta_b, \quad I_\gamma = \theta.$$

Equation (11) gives $\Psi_j = \mathbb{E}[GI_j]$ for each policy coordinate $j \in \{\alpha, b, \gamma\}$. Furthermore, differentiating the $t = 1$ loss we get

$$\frac{\partial}{\partial \alpha}(W_1 \mathbb{E}[\ell^2]) = 2W_1 \mathbb{E}[\mathcal{R}S], \quad \frac{\partial}{\partial b}(W_1 \mathbb{E}[\ell^2]) = 2W_1 \mathbb{E}[\mathcal{R}], \quad \frac{\partial}{\partial \gamma}(W_1 \mathbb{E}[\ell^2]) = -2W_1 \mathbb{E}[\ell].$$

To compute $\mathbb{E}[\mathcal{R}S]$, first consider the case where $\sigma_z^2 := \text{Var}(z) > 0$ and write

$$g(x) := (|x| - \gamma)_+ \text{sign}(x), \quad \mathcal{R} = g(z).$$

Since (S, z) is jointly Gaussian, $\mathbb{E}[S] = 0$, and $\mathbb{E}[z] = b$, its conditional mean is simply the standard linear regression expression,

$$\mathbb{E}[S | z] = \frac{\text{Cov}(S, z)}{\sigma_z^2}(z - b).$$

Furthermore,

$$g(x) = \begin{cases} x + \gamma, & x < -\gamma, \\ 0, & |x| \leq \gamma, \\ x - \gamma, & x > \gamma \end{cases}$$

and thus g is absolutely continuous with $g'(x) = \mathbf{1}_{\{|x| > \gamma\}}$ well defined almost everywhere.

We have

$$\begin{aligned} \mathbb{E}[(z - b)g(z)] &= \int_{-\infty}^{\infty} (x - b)g(x)f_z(x)dx \\ &= -\sigma_z^2 \int_{-\infty}^{\infty} g(x)f'_z(x)dx \\ &= \sigma_z^2 \int_{-\infty}^{\infty} g'(x)f_z(x)dx \\ &= \sigma_z^2 \mathbb{P}(|z| > \gamma) = \sigma_z^2 p, \end{aligned}$$

where the second equality uses that since f_z is the Gaussian density of z ,

$$f'_z(x) = -\left(\frac{x-b}{\sigma_z^2}\right) f_z(x)$$

and the third equality uses integration by parts and the fact that the boundary terms vanish.³² The same argument applies at $\gamma = 0$, when $g(x) = x$. Taking conditional expectations and using $z = \alpha S + b - T$, it follows that

$$\begin{aligned} \mathbb{E}[\mathcal{R}S] &= \mathbb{E}[g(z) \mathbb{E}[S \mid z]] \\ &= \frac{\text{Cov}(S, z)}{\sigma_z^2} \mathbb{E}[(z-b)g(z)] \\ &= \text{Cov}(S, z)p \\ &= (\alpha \text{Var}(S) - \text{Cov}(S, T))p \\ &= \text{Var}(S)(\alpha - \tilde{\alpha}_1)p, \end{aligned}$$

where

$$\tilde{\alpha}_1 := \frac{\text{Cov}(S, T)}{\text{Var}(S)} = a_{cb} + \frac{\text{Cov}(S, a_{cb}^0(s_0, \hat{s}_1))}{\text{Var}(S)}.$$

Now, if z is degenerate, $\text{Var}(z) = 0$ implies $z = \mathbb{E}[z] = b$ almost surely. Hence, $\mathcal{R} = g(b)$ is constant, and

$$\mathbb{E}[\mathcal{R}S] = g(b) \mathbb{E}[S] = 0, \quad \text{Cov}(S, z) = \mathbb{E}[S(z-b)] = 0.$$

Thus, $\mathbb{E}[\mathcal{R}S] = \text{Cov}(S, z)p$ holds too. Consequently,

$$\frac{\partial \mathcal{J}}{\partial \alpha} = \mathbb{E}[G(M + \gamma\theta_\alpha)] + 2W_1 \text{Var}(S)(\alpha - \tilde{\alpha}_1)p, \quad (13)$$

$$\frac{\partial \mathcal{J}}{\partial b} = \mathbb{E}[G(1 + \gamma\theta_b)] + 2W_1 \mathbb{E}[\mathcal{R}], \quad (14)$$

³²Concretely, fix $L > \gamma$. Integrating by parts,

$$\int_{-L}^L g(x) f'_z(x) dx = [g(x) f_z(x)]_{-L}^L - \int_{-L}^L g'(x) f_z(x) dx.$$

Applying this separately on the intervals delimited by $-\gamma$ and γ , the boundary terms at those points vanish since $g(-\gamma) = g(\gamma) = 0$. At the outer boundaries,

$$|g(x) f_z(x)| \leq \frac{|x|}{\sqrt{2\pi}\sigma_z} \exp\left(-\frac{(x-b)^2}{2\sigma_z^2}\right) \longrightarrow 0 \quad \text{as } |x| \rightarrow \infty.$$

Thus, as $L \rightarrow \infty$, the remaining boundary term also vanishes.

$$\frac{\partial \mathcal{J}}{\partial \gamma} = \mathbb{E}[G\theta] - 2W_1 \mathbb{E}[\ell]. \quad (15)$$

At a (finite) local minimum with $\alpha \neq a_{ps}$, $\frac{\partial \mathcal{J}}{\partial \alpha} = \frac{\partial \mathcal{J}}{\partial b} = 0$. Since vagueness is non negative by definition, the third condition is the complementary slackness one, $\frac{\partial \mathcal{J}}{\partial \gamma} \geq 0$ and $\gamma \frac{\partial \mathcal{J}}{\partial \gamma} = 0$. These are the necessary conditions in the statement of the proposition.

Current stabilization and the slope targets. To relate these conditions to the discussion in the text, define

$$\begin{aligned} \xi_r &:= 2(A_y^{i^e} A_y^r + \lambda_{cb} A_\pi^{i^e} A_\pi^r), \\ \xi_u &:= 2(A_y^{i^e} A_y^u + \lambda_{cb} A_\pi^{i^e} A_\pi^u), \\ \xi_i &:= 2((A_y^{i^e})^2 + \lambda_{cb} (A_\pi^{i^e})^2) = 2Q. \end{aligned}$$

Since the conditional expectation of the conditionable signal equals $M = \tilde{\omega}_r \rho_r r_0^n + \tilde{\omega}_u \rho_u u_0$, it's straightforward to see differentiating (10) that the target minimizing $\Psi(\alpha, 0, 0)$, whenever $\text{Var}(M) > 0$, is simply

$$\tilde{\alpha}_0 = -\frac{\text{Cov}(H, M)}{Q \text{Var}(M)} = -\frac{\xi_r \tilde{\omega}_r \rho_r \sigma_r^2 + \xi_u \tilde{\omega}_u \rho_u \sigma_u^2}{\xi_i \text{Var}(M)}. \quad (16)$$

For a general (α, b, γ) announcement, we can write

$$\begin{aligned} \frac{1}{2} \Psi_\alpha &= Q \text{Var}(M)(\alpha - \tilde{\alpha}_0) + Q\gamma \text{Cov}(M, \theta) \\ &\quad + \gamma \text{Cov}(H + QI, \theta_\alpha) + wm\gamma \mathbb{E}[\theta_\alpha], \end{aligned} \quad (17)$$

$$\frac{1}{2} \Psi_b = \gamma \text{Cov}(H + QI, \theta_b) + wm(1 + \gamma \mathbb{E}[\theta_b]), \quad (18)$$

$$\frac{1}{2} \Psi_\gamma = \text{Cov}(H + Q\alpha M, \theta) + Q\gamma \text{Var}(\theta) + wme, \quad (19)$$

where we used that $\mathbb{E}[H] = \mathbb{E}[M] = 0$ and $I - m = \alpha M + \gamma(\theta - e)$.

The first term in (17) is the pull toward $\tilde{\alpha}_0$, while the other terms account for worst-case selection. The terms involving wm represent the residual cost of average rate expectations (accounting for the partial offset of the i_0 adjustment).

Symmetry and mean alignment. For $\alpha \neq a_{ps}$, the symmetry of the Gaussian distribution gives

$$\theta_{\alpha, -b}(-s_0) = -\theta_{\alpha, b}(s_0).$$

The corresponding expected rate and optimal current rate change sign, and the reflected future gap is $-z$. Thus, since the central bank's loss is quadratic,

$$\mathcal{J}(\alpha, -b, \gamma) = \mathcal{J}(\alpha, b, \gamma).$$

Thus $\mathcal{J}_b(\alpha, 0, \gamma) = 0$, and at $b = 0$ symmetry gives $e = m = 0$ and $i_0^* = 0$. Further, these also hold at $\alpha = a_{ps}$ when the indifference between endpoints has probability zero, as it is the case under the condition below.

The future loss alone is even and convex in b , and thus it is minimized at zero for any fixed (α, γ) . If $\text{Var}(z) > 0$, its curvature is

$$\frac{\partial^2}{\partial b^2}(W_1 \mathbb{E}[(|z| - \gamma)_+^2]) = 2W_1 \mathbb{P}(|z| > \gamma) > 0$$

at every finite vagueness γ . Consequently, any gain from mean misalignment, $b \neq 0$, must come from current stabilization.

Existence of a (finite) solution. To prove the existence of a finite solution to the central bank's problem, assume the additional conditions stated in the proposition, i.e., both shock processes have nonzero persistence and positive variances, and what we call a *nonalignment condition*, $\tilde{\omega}_u \neq \Gamma(\lambda_{ps})\tilde{\omega}_r$

Note that the latter condition implies that $\tilde{\omega}_u$ and $\tilde{\omega}_r$ are not both zero, so together with the persistence and variance conditions, $\text{Var}(M) > 0$. Further, the nonalignment condition implies that F is not a linear function of M and thus $\text{Var}(F | M) > 0$ as well.

Now, we show that this implies that the objective function has compact sublevel sets. Young's inequality applied to (10) implies

$$\Psi \geq K - \frac{2 \text{Var}(H)}{Q} + \frac{Q}{2} \text{Var}(I) + wm^2.$$

Now, fix a finite loss threshold $B > 0$ and consider a nonempty $\mathcal{J}$ sublevel set, that is, the collection of feasible announcements satisfying $\mathcal{J}(\alpha, b, \gamma) \leq B$. Since the future penalty is nonnegative, $\mathcal{J} \geq \Psi$, and the inequality above implies

$$\frac{Q}{2} \text{Var}(I) + wm^2 \leq B - K + \frac{2 \text{Var}(H)}{Q}.$$

Both terms on the left are nonnegative and have strictly positive coefficients. Hence, neither the variance nor the squared mean of I can grow without bound without the total loss $\mathcal{J}$

going above B . Note that $\|I\|_2^2 := \mathbb{E}[I^2] = \text{Var}(I) + m^2$. Consequently,

$$\|I\|_2^2 \leq \frac{B - K + 2 \text{Var}(H)/Q}{\min\{Q/2, w\}} < \infty,$$

and note that since the bound depends only on B and the model parameters, it holds uniformly over all announcements in the sublevel set.

Let $C := \alpha M + b$. Note that $F = a_{ps}M + a_{ps}^0(s_0)$, so the conditional mean in the numerator of the worst-case selector is

$$C - F = (\alpha - a_{ps})M + b - a_{ps}^0(s_0).$$

For $\alpha \neq a_{ps}$, the expression for the selector in [Lemma 2](#) can then be written as

$$\theta = 2\Phi\left(\frac{C - F}{|\alpha - a_{ps}|\sigma_{\tilde{s}}}\right) - 1.$$

Since Φ is strictly increasing and $\Phi(0) = 1/2$, θ has the same sign as $C - F$ whenever the latter is nonzero. At $\alpha = a_{ps}$, the selector equals 1 for a positive gap and -1 for a negative gap, so the same sign property holds. By construction, the worst-case adjustment $\gamma\theta$ moves the expected rate farther from F , so

$$|I - F| = |C - F + \gamma\theta| = |C - F| + \gamma|\theta| \geq |C - F|.$$

Further, this also holds when $C - F = 0$ since then $I - F = \gamma\theta$, using the upper-endpoint tie convention at $\alpha = a_{ps}$ ([Lemma 2](#)).

Squaring the pointwise inequality and taking expectations gives $\|C - F\|_2 \leq \|I - F\|_2$. Applying the triangle inequality twice then yields

$$\begin{aligned} \|C\|_2 &\leq \|C - F\|_2 + \|F\|_2 \\ &\leq \|I - F\|_2 + \|F\|_2 \\ &\leq \|I\|_2 + 2\|F\|_2. \end{aligned}$$

The random variable F has a finite second moment and does not depend on the announcement. Hence, the uniform bound on $\|I\|_2$ established above also bounds $\|C\|_2$ uniformly on the sublevel set. Finally, since $\mathbb{E}[M] = 0$,

$$\|C\|_2^2 = \mathbb{E}[(\alpha M + b)^2] = \alpha^2 \text{Var}(M) + b^2.$$

Both terms on the right are nonnegative and $\text{Var}(M) > 0$, so

$$|\alpha| \leq \frac{\|C\|_2}{\sqrt{\text{Var}(M)}}, \quad |b| \leq \|C\|_2.$$

Consequently, α and b are bounded uniformly on every sublevel set.

To bound γ , on top of the above we must rule out the possibility that $\|\theta_{\alpha,b}\|_2$ approaches zero as (α, b) vary within the bounds just established. Fix a closed and bounded rectangle $\mathcal{K}_B$ containing all (α, b) in the sublevel set. We first show that $(\alpha, b) \mapsto \theta_{\alpha,b}$ is continuous in L^2 on this rectangle.

To do so, consider an arbitrary sequence $(\alpha_n, b_n) \rightarrow (\alpha, b)$. If $\alpha \neq a_{ps}$, the denominator within the worst-case selector converges to a strictly positive number. Continuity of the cdf of a Gaussian random variable thus implies $\theta_{\alpha_n, b_n}(s_0) \rightarrow \theta_{\alpha, b}(s_0)$ pointwise in s_0 . If instead $\alpha = a_{ps}$, the law of total variance implies

$$\text{Var}(a_{ps}M + b - F) \geq \mathbb{E}[\text{Var}(a_{ps}M + b - F \mid M)] = \mathbb{E}[\text{Var}(F \mid M)] > 0.$$

Thus, $a_{ps}M + b - F$ is a nondegenerate Gaussian variable, and consequently

$$\mathbb{P}(a_{ps}M + b - F = 0) = 0.$$

Thus, the numerator $\alpha_n M + b_n - F$ converges to a nonzero number almost surely. Along terms with $\alpha_n \neq a_{ps}$, the denominator $|\alpha_n - a_{ps}| \sigma_{\tilde{s}}$ goes to zero, so the CDF argument tends to $+\infty$ or $-\infty$ depending on the sign of the limiting numerator. Thus, the worst-case selector converges to $\text{sign}(a_{ps}M + b - F)$, which equals $\theta_{a_{ps}, b}$. Terms with $\alpha_n = a_{ps}$ have the same limit directly from the sign formula, since the numerator has its limiting sign eventually. Thus, the worst-case selectors converge almost surely in both cases.

To further show the convergence in L^2 , note that every selector lies in $[-1, 1]$, so

$$0 \leq |\theta_{\alpha_n, b_n} - \theta_{\alpha, b}|^2 \leq 4.$$

Bounded convergence applied to this squared difference yields

$$\|\theta_{\alpha_n, b_n} - \theta_{\alpha, b}\|_2^2 = \mathbb{E} \left[|\theta_{\alpha_n, b_n} - \theta_{\alpha, b}|^2 \right] \rightarrow 0.$$

Thus, by the reverse triangle inequality,

$$||\theta_{\alpha_n, b_n}\|_2 - \|\theta_{\alpha, b}\|_2| \leq \|\theta_{\alpha_n, b_n} - \theta_{\alpha, b}\|_2 \rightarrow 0,$$

so the scalar function $(\alpha, b) \mapsto \|\theta_{\alpha,b}\|_2$ is continuous as well.

Next, this norm is strictly positive at every point of $\mathcal{K}_B$. To see this, note that for $\alpha \neq a_{ps}$, a zero norm requires $\theta_{\alpha,b} = 0$ almost surely. Since $2\Phi(x) - 1 = 0$ only when $x = 0$, this in turn requires $F = \alpha M + b$ almost surely, contradicting $\text{Var}(F \mid M) > 0$. At $\alpha = a_{ps}$, the worst-case selector has absolute value one almost surely, so its L^2 norm is one. A continuous function on the compact rectangle $\mathcal{K}_B$ attains its minimum. Since its value is positive at every point, this minimum is strictly positive. Consequently,

$$\varepsilon_B := \min_{(\alpha,b) \in \mathcal{K}_B} \|\theta_{\alpha,b}\|_2 > 0.$$

Since this is a uniform lower bound over the rectangle, it also holds over all announcements in the sublevel set indexed by B

Finally, $I - C = \gamma\theta_{\alpha,b}$ and $\gamma \geq 0$ imply

$$\gamma\varepsilon_B \leq \gamma\|\theta_{\alpha,b}\|_2 = \|I - C\|_2 \leq \|I\|_2 + \|C\|_2.$$

Both norms on the right have already been bounded uniformly on the sublevel set. Thus, dividing by the positive scalar ε_B gives a uniform upper bound on γ as well.

It remains to show that $\mathcal{J}$ is continuous, so that the bounded sublevel sets established above are also closed. Consider any sequence of feasible announcements $(\alpha_n, b_n, \gamma_n) \rightarrow (\alpha, b, \gamma)$. Write $\theta_n := \theta_{\alpha_n, b_n}$ and $I_n := \alpha_n M + b_n + \gamma_n \theta_n$, with θ and I denoting the corresponding quantities at the limiting announcement. The worst-case selector continuity argument above applies at any finite (α, b) , including $\alpha = a_{ps}$ provided the nonalignment condition holds, and gives $\|\theta_n - \theta\|_2 \rightarrow 0$. Since $\|\theta_n\|_2 \leq 1$,

$$\begin{aligned} \|I_n - I\|_2 &\leq |\alpha_n - \alpha| \|M\|_2 + |b_n - b| + |\gamma_n - \gamma| + \gamma\|\theta_n - \theta\|_2 \\ &\longrightarrow 0. \end{aligned}$$

Thus, the sequence of induced expected rates converges in L^2 . This implies convergence of their means and second moments, since

$$\begin{aligned} |\mathbb{E}[I_n] - \mathbb{E}[I]| &\leq \|I_n - I\|_2 \longrightarrow 0, \\ |\mathbb{E}[I_n^2] - \mathbb{E}[I^2]| &\leq \|I_n - I\|_2 (\|I_n\|_2 + \|I\|_2) \longrightarrow 0, \end{aligned}$$

where the second inequality follows by writing $I_n^2 - I^2 = (I_n - I)(I_n + I)$ and applying the Cauchy-Schwarz inequality, and the norms $\|I_n\|_2$ are bounded because $I_n \rightarrow I$ in L^2 . Thus,

the sequence of variances converges as well. Moreover, $\mathbb{E}[H] = 0$ and Cauchy-Schwarz give

$$|\text{Cov}(H, I_n) - \text{Cov}(H, I)| \leq \|H\|_2 \|I_n - I\|_2 \longrightarrow 0.$$

Consequently, every term in the $t = 0$ loss decomposition (10) converges, proving continuity of Ψ .

Now, for $t = 1$ loss, write $z_n := \alpha_n S + b_n - T$ and $\ell_n := (|z_n| - \gamma_n)_+$. The absolute-value and positive-part functions each satisfy $|f(x) - f(y)| \leq |x - y|$. Thus,

$$|\ell_n - \ell| \leq \left| |z_n| - |z| \right| + |\gamma_n - \gamma| \leq |z_n - z| + |\gamma_n - \gamma|.$$

Taking L^2 norms then gives

$$\begin{aligned} \|\ell_n - \ell\|_2 &\leq \|z_n - z\|_2 + |\gamma_n - \gamma| \\ &\leq |\alpha_n - \alpha| \|S\|_2 + |b_n - b| + |\gamma_n - \gamma| \longrightarrow 0. \end{aligned}$$

As above, convergence in L^2 implies $\mathbb{E}[\ell_n^2] \rightarrow \mathbb{E}[\ell^2]$. Hence, the $t = 1$ loss $W_1 \mathbb{E}[\ell^2]$ is continuous too. Together, these arguments prove the continuity of $\mathcal{J} = \Psi + W_1 \mathbb{E}[\ell^2]$ on the entire feasible set $\mathbb{R}^2 \times [0, \infty)$, including at $\alpha = a_{ps}$ and $\gamma = 0$.

Putting the above together, we get that for any finite threshold $B > 0$, the sublevel set

$$\{(\alpha, b, \gamma) \in \mathbb{R}^2 \times [0, \infty) : \mathcal{J}(\alpha, b, \gamma) \leq B\}$$

is closed in $\mathbb{R}^3$. In addition, the preceding bounds on α , b , and γ make this set bounded. Thus, it is compact by the Heine-Borel theorem.

Finally, choose $B = 1 + \mathcal{J}(0, 0, 0)$, a finite value since at this announcement $I = 0$, $z = -T$, and thus $\mathcal{J}(0, 0, 0) = K + W_1 \mathbb{E}[T^2] < \infty$. Its sublevel set is clearly nonempty, since it contains $(0, 0, 0)$. By the extreme value theorem, $\mathcal{J}$ attains a minimum on this compact set. By construction, every feasible announcement outside the set has loss strictly greater than B , while the attained minimum is at most B . Thus, the minimizer on the sublevel set is also a global minimizer over the entire feasible set, establishing the existence of a finite solution.

A nonzero mean natural rate. To establish the translation result for $\bar{r}^n \neq 0$, define the centered variables

$$r_t^{n,c} := r_t^n - \bar{r}^n, \quad i_t^c := i_t - \bar{r}^n, \quad \tilde{s}_1^c := \tilde{s}_1 - \tilde{\omega}_r \bar{r}^n, \quad \hat{s}_1^c := \hat{s}_1 - \hat{\omega}_r \bar{r}^n.$$

Naturally, the centered shocks and signals have the same distributions as in the zero-mean model. Subtracting $\bar{r}^n$ from both policy rates and natural rates leaves output, inflation, and both agents' losses unchanged. Moreover,

$$\alpha \tilde{s}_1 + b - \bar{r}^n = \alpha \tilde{s}_1^c + b^c, \quad b^c := b + (\alpha \tilde{\omega}_r - 1) \bar{r}^n,$$

while the announcement γ is unchanged. This gives a one-to-one correspondence between feasible policies in the normalized and general cases, preserving worst-case selections and losses. Thus, any optimal policy at $\bar{r}^n = 0$ has a corresponding optimum in the general case with the same data dependence and vagueness and

$$i_{0,\bar{r}^n}^* = i_{0,0}^* + \bar{r}^n, \quad b_{\bar{r}^n}^* = b_0^* + (1 - \alpha^* \tilde{\omega}_r) \bar{r}^n.$$

□

B.6 Proof of Proposition 4

Fix $p = (i_0, \alpha, b, \gamma)$ with $\gamma \geq 0$ and $\alpha \neq \alpha^* := 1/A_\pi^{i^e}$. At $\alpha = \alpha^*$, equilibrium requires

$$C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} (b + \gamma \vartheta_0) = 0.$$

For finite γ , this condition cannot hold for all shock realizations: $C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} b$ is a nondegenerate Gaussian random variable, while $|A_\pi^{i^e} \gamma \theta_0| \leq |A_\pi^{i^e}| \gamma$. Thus there is positive probability of states with no equilibrium. Thus, we exclude this α value from the finite policy domain to ensure equilibrium existence almost surely. This does not rule out finite limiting losses as data dependence approaches α^* jointly with an increasingly vague announcement. Those limits are addressed in Proposition 5.

It has already been established (see the proof of Proposition 1) that

$$\pi_0 = C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} i_1^e.$$

At $t = 0$, an equilibrium is a pair (π_0, ϑ_0) such that $\vartheta_0 \in [-1, 1]$ and

$$i_1^e = \alpha \pi_0 + b + \gamma \vartheta_0$$

as the argument of Lemma 2 applies analogously, substituting the exogenous signal by π_0 .

Substituting and solving yields that, for any $\vartheta \in [-1, 1]$,

$$\pi_0(\vartheta) = \frac{C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} (b + \gamma \vartheta)}{1 - \alpha A_\pi^{i^e}}.$$

It remains to characterize the set $\mathcal{E}_p(s_0) \subseteq [-1, 1]$ of equilibrium values of ϑ_0 .

Fix $p = (i_0, \alpha, b, \gamma, s_0)$ and a candidate $\vartheta \in [-1, 1]$. Define

$$D_{\text{mis}}(\vartheta; s_0) = \alpha \pi_0(\vartheta) + b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0]$$

Given ϑ , the two ends of the interval correspond to policy rates with deviations from the private sector's preferred mean equal to $D_{\text{mis}}(\vartheta; s_0) \pm \gamma$. Hence the upper end $+\gamma$ is (strictly) worse than the lower end $-\gamma$ if and only if

$$|D_{\text{mis}}(\vartheta; s_0) + \gamma| > |D_{\text{mis}}(\vartheta; s_0) - \gamma| \iff D_{\text{mis}}(\vartheta; s_0) > 0,$$

and the lower end is strictly worse if and only if

$$|D_{\text{mis}}(\vartheta; s_0) - \gamma| > |D_{\text{mis}}(\vartheta; s_0) + \gamma| \iff D_{\text{mis}}(\vartheta; s_0) < 0.$$

Agents are indifferent between the two ends if and only if $D_{\text{mis}}(\vartheta; s_0) = 0$.

Since $\pi_0(\vartheta)$ is affine in ϑ , so is $D_{\text{mis}}(\vartheta; s_0)$. In particular, one can rewrite the indifference condition $D_{\text{mis}}(\vartheta; s_0) = 0$ as

$$\vartheta = \frac{T(s_0; p)}{\gamma},$$

where $T(s_0; p)$ is as defined in the statement of the proposition.

Thus, the unique ϑ that supports indifference satisfies $\vartheta \in [-1, 1]$ if and only if $|T(s_0; p)| \leq \gamma$; equivalently, the “intermediate” equilibrium value is $\vartheta = T/\gamma$ whenever $|T| \leq \gamma$.

Next, consider equilibria such that one end is strictly worse than the other. Evaluating $D_{\text{mis}}(\vartheta; s_0)$ at $\vartheta = +1$ and $\vartheta = -1$ and using the sign of

$$\frac{\alpha A_\pi^{i^e}}{1 - \alpha A_\pi^{i^e}}$$

yields the following cases.

Case (i): $\alpha > 0$ or $\alpha < \alpha^*$. In this region $\alpha A_\pi^{i^e} / (1 - \alpha A_\pi^{i^e}) < 0$, and the strict best-response conditions imply

$$\vartheta_0 = +1 \text{ is an equilibrium} \iff T(s_0; p) \geq \gamma,$$

$$\vartheta_0 = -1 \text{ is an equilibrium } \iff T(s_0; p) \leq -\gamma.$$

When $|T(s_0; p)| \leq \gamma$, neither extreme is strictly worse, and the unique equilibrium is the indifference value $\vartheta_0 = T(s_0; p)/\gamma$, with the interpretation that a share of the market picks the upper extreme and the rest pick the lower extreme. Hence $\mathcal{E}_p(s_0)$ is a singleton and equals the expression in the statement.

Case (ii): $\alpha \in (\alpha^*, 0)$. In this region $\alpha A_\pi^{i^e}/(1 - \alpha A_\pi^{i^e}) > 0$, so the strict best-response conditions are reversed, implying

$$\begin{aligned} \vartheta_0 = -1 \text{ is an equilibrium } &\iff T(s_0; p) \geq -\gamma, \\ \vartheta_0 = +1 \text{ is an equilibrium } &\iff T(s_0; p) \leq \gamma. \end{aligned}$$

Moreover, when $|T(s_0; p)| \leq \gamma$, the indifference value $\vartheta_0 = T(s_0; p)/\gamma$ is also feasible and hence constitutes an additional equilibrium. This yields the set-valued characterization in the statement.

Case (iii): $\alpha = 0$. When $\alpha = 0$, inflation does not enter the announcement, so there is no feedback through π_0 and

$$D_{\text{mis}}(\vartheta; s_0) = b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0]$$

is independent of ϑ . Hence $\vartheta_0 = +1$ is the unique equilibrium if $b > \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0]$, $\vartheta_0 = -1$ is the unique equilibrium if $b < \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0]$, and if equality holds then any $\vartheta_0 \in [-1, 1]$ is an equilibrium.

Collecting the three cases yields the characterization of $\mathcal{E}_p(s_0)$ in the statement of the proposition. Given any $\vartheta_0 \in \mathcal{E}_p(s_0)$, equilibrium inflation is then given by $\pi_0(\vartheta_0)$ above.

B.7 Proof of Proposition 5

Proof of Proposition 5.

- (i) *Objective Function.* Fix $p = (i_0, \alpha, b, \gamma) \in \mathcal{D}$. Under the selection convention following Proposition 4, for each s_0 define

$$i_1^e(p; s_0) := \alpha \pi_0(p; s_0) + b + \gamma \vartheta_p(s_0),$$

where $\pi_0(p, \cdot)$ is the associated $t = 0$ equilibrium inflation given the selection and $\vartheta_p(\cdot)$ is the selected equilibrium worst-case choice. Using the expressions for y_0 and π_0 from

Proposition 1 (which remain valid with endogenous signals once i_1^e is defined as above), $t = 0$ losses are

$$\begin{aligned}\tilde{\Psi}(i_0, \alpha, b, \gamma) &= \mathbb{E} \left[(C_y + A_y^r r_0^n + A_y^u u_0 - \sigma i_0 + A_y^{i^e} i_1^e)^2 \right] \\ &\quad + \lambda_{cb} \mathbb{E} \left[(C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} i_1^e)^2 \right].\end{aligned}$$

At $t = 1$ the central bank chooses the policy rate subject to the announcement constraint; equivalently, it implements the projection of its unconstrained choice onto the announced interval. Hence, up to policy-independent terms, the $t = 1$ loss equals

$$W_1 \mathbb{E} \left[\left(\left| \alpha \pi_0(i_0, \alpha, b, \gamma; s_0) + b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0] \right| - \gamma \right)_+^2 \right].$$

Define

$$z_{\alpha, b, \gamma}(s_0) := \alpha \pi_0(i_0, \alpha, b, \gamma; s_0) + b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0].$$

Hence,

$$\tilde{\mathcal{J}}(i_0, \alpha, b, \gamma) = \tilde{\Psi}(i_0, \alpha, b, \gamma) + W_1 \mathbb{E} \left[(|z_{\alpha, b, \gamma}| - \gamma)_+^2 \right].$$

- (ii) *Existence.* The proof proceeds in five steps. First, we identify a family of limiting outcomes associated with increasingly vague announcements ($\gamma \rightarrow +\infty$). Second, we establish continuity of expected loss at finite policies. Third, we examine sequences in which α approaches its excluded value α^* or policy parameters become unbounded while losses remain bounded. We show that their limiting losses cannot fall below the minimum loss within the family identified in the first step. Fourth, we show that every outcome in this family can be approached by finite announcements with $\gamma \rightarrow \infty$ and find the lowest loss within the family. Finally, we compare this limiting minimum with finite-policy losses and conclude: if some finite policy does better, a finite minimizer exists, otherwise, the best limiting outcome attains the infimum.

Let

$$V := \inf_{p \in \mathcal{D}} \tilde{\mathcal{J}}(p)$$

and, for notational convenience, define

$$\begin{aligned}A &:= A_\pi^{i^e} < 0, & B &:= A_y^{i^e}, \\ F_{i_0}(s_0) &:= C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0, \\ G_{i_0}(s_0) &:= C_y + A_y^r r_0^n + A_y^u u_0 - \sigma i_0, \\ P(s_0) &:= \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0], & Q(s_0) &:= \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0].\end{aligned}$$

Also letting $I := i_1^e$, we can write

$$\pi_0 = F_{i_0} + AI, \quad y_0 = G_{i_0} + BI.$$

Set $D := 1 - \alpha A$ and $q := \gamma \vartheta_p$, with $q = 0$ when $\gamma = 0$ (the equilibrium selection does not matter in that case). For $\alpha \neq 0$ ($\alpha \neq \alpha^*$ throughout), from [Proposition 4](#), we know the threshold and the private sector's expected rate satisfy

$$T = -\frac{\alpha F_{i_0} + b - DP}{\alpha A}, \quad I = \frac{\alpha F_{i_0} + b + q}{D}.$$

Furthermore, given our equilibrium selection convention, $q = T$ when $|T| \leq \gamma$ and $|q| = \gamma$ otherwise. Thus, $|q| \leq |T|$. Moreover, letting $C := \alpha\pi_0 + b$, every equilibrium satisfies

$$I \in [C - \gamma, C + \gamma], \quad |C - P| \leq |I - P|, \quad (|C - Q| - \gamma)_+ \leq |I - Q|. \quad (20)$$

Indeed, when $C \neq P$, the worst-case endpoint is $I = C + \gamma \text{sign}(C - P)$; when $C = P$, the middle inequality is immediate. On the other hand, the right inequality follows from $|C - Q| \leq |C - I| + |I - Q| \leq \gamma + |I - Q|$ and the nonnegativity of $|I - Q|$.

Fix (i_0, α, b) with $\alpha \neq 0$. Since T is an affine function of Gaussian random variables, as $\gamma \rightarrow \infty$ it eventually lies in the intermediate region as $\gamma \rightarrow \infty$. In that region, $C = P$, and thus

$$\pi_0 \longrightarrow \frac{P - b}{\alpha}, \quad I \longrightarrow \frac{(P - b)/\alpha - F_{i_0}}{A}.$$

The bound $|q| \leq |T|$ gives a common linear Gaussian envelope for I, y_0, π_0 , uniformly in γ . It also gives an integrable quadratic envelope for the $t = 1$ loss, given the third inequality in [\(20\)](#) above. That loss tends to zero state by state, since eventually $C = P$ as $\gamma \rightarrow \infty$. Consequently, dominated convergence implies that expected $t = 0$ loss converges to the loss associated with these limiting outcomes, while expected $t = 1$ loss tends to zero. Hence, $\lim_{\gamma \rightarrow \infty} \tilde{\mathcal{J}}(i_0, \alpha, b, \gamma)$ exists and is finite.

To account also for joint limits of (α, b, γ) , let k and h denote the slope and intercept of limiting inflation as an affine function of $P(s_0)$. Write the limiting outcomes and loss as

$$\begin{aligned} \pi_L(s_0) &:= kP(s_0) + h, & I_L(s_0) &:= \frac{\pi_L(s_0) - F_{i_0}(s_0)}{A}, \\ y_L(s_0) &:= G_{i_0}(s_0) + BI_L(s_0), & \mathcal{J}_L(i_0, k, h) &:= \mathbb{E}[y_L^2 + \lambda_{cb}\pi_L^2]. \end{aligned}$$

The fixed- α limit above corresponds to $k = 1/\alpha$ and $h = -b/\alpha$. The values $k = A$ and $k = 0$ will account, respectively, for the singular- α (i.e., as α approaches α^*) and infinite- α limits.

Continuity at finite policies. We first show that the central bank's loss function is a continuous function of policy p at $\alpha \notin \{0, \alpha^*\}$. Since at $\alpha = 0$, T is undefined, we then use a separate argument for that case. Consider any sequence of feasible policies $p_n \rightarrow p$.

For $\gamma > 0$, q is continuous at stabilizing policies. For de-stabilizing policies, possible discontinuities occur only when $T(s_0; p) = \pm\gamma$, a probability zero even because T is a nondegenerate Gaussian random variable (see the argument in part (iii) below³³). For $\gamma = 0$, convergence follows directly from $|q(s_0; p_n)| \leq \gamma_n \rightarrow 0$. The equilibrium expressions above thus imply that current outcomes and the losses in both periods converge for almost every shock realization.

Since $p_n \rightarrow p$, the policy parameters are uniformly bounded, and $|1 - \alpha_n A|$ stays bounded away from zero for sufficiently large n . Together with $|q(s_0; p_n)| \leq \gamma_n$, the equilibrium formulas give a common linear bound in the shocks for current outcomes. The central bank's (quadratic) $t = 0$ loss and, by the third inequality in (20), the $t = 1$ loss thus have a common integrable quadratic bound. Dominated convergence then implies

$$\tilde{\mathcal{J}}(p_n) \rightarrow \tilde{\mathcal{J}}(p),$$

which shows that the objective function is continuous at every finite policy with $\alpha \notin \{0, \alpha^*\}$.

Now, consider a sequence of feasible policies $p_n \rightarrow p = (i_0, 0, b, \gamma)$. Since $D_n \rightarrow 1$ and $|q_n| \leq \gamma_n$, the equilibrium formulas again give a common linear bound in the shocks for current outcomes. In particular, inflation remains bounded for each fixed shock realization, so

$$C_n = \alpha_n \pi_{0,n} + b_n \rightarrow b.$$

For any realization with $P \neq b$, $C_n - P$ therefore eventually has the same sign as $b - P$. Thus,

$$I_n = C_n + \gamma_n \text{sign}(C_n - P) \rightarrow b + \gamma \text{sign}(b - P),$$

which is the aggregate rate expectation at p . Since P is a nondegenerate Gaussian

³³That argument applies to general $\alpha \in \mathbb{R} \setminus \alpha^*$ and thus requires [condition 8](#). In the case of de-stabilizing policies, this does not require the condition.

variable, $\mathbb{P}(P = b) = 0$, so this convergence holds almost surely.

The equilibrium formulas then imply convergence of output, inflation, and the losses in both periods for almost every shock realization. The same integrable quadratic bound used above allows dominated convergence to give

$$\tilde{\mathcal{J}}(p_n) \longrightarrow \tilde{\mathcal{J}}(p),$$

establishing continuity also at policies with $\alpha = 0$.

Bounds along a minimizing sequence. Let $p_n = (i_{0,n}, \alpha_n, b_n, \gamma_n)$ be any sequence of finite policies with bounded central bank's losses. Since total loss includes $\mathbb{E}[y_{0,n}^2] + \lambda_{cb} \mathbb{E}[\pi_{0,n}^2]$ and the remaining loss is nonnegative, the fact that losses are bounded implies the second moments of output and inflation are uniformly bounded. Equivalently, $\|y_{0,n}\|_2$ and $\|\pi_{0,n}\|_2$ are bounded, where $\|X\|_2 := \sqrt{\mathbb{E}[X^2]}$. Furthermore,

$$\begin{pmatrix} y_{0,n} - G_0 \\ \pi_{0,n} - F_0 \end{pmatrix} = \begin{pmatrix} -\sigma & B \\ -\sigma\kappa & A \end{pmatrix} \begin{pmatrix} i_{0,n} \\ I_n \end{pmatrix},$$

and the coefficient matrix has determinant

$$\sigma(\kappa B - A) = \sigma^2 \kappa \beta > 0.$$

Its inverse therefore bounds $|i_{0,n}|$ and $\|I_n\|_2$. The equilibrium inequality gives

$$\|C_n\|_2 \leq \|I_n\|_2 + 2\|P\|_2, \quad \frac{|b_n|}{1 + |\alpha_n|} \leq \frac{\|C_n\|_2 + |\alpha_n| \|\pi_{0,n}\|_2}{1 + |\alpha_n|}.$$

Consequently the normalized intercept is bounded.

Consider a convergent subsequence such with limit $p = (i_0, \alpha, b, \gamma)$, $\alpha \in \mathbb{R} \cup \{\pm\infty\}$ and $\gamma \in [0, \infty]$. By the inequality above, if the data dependence limit is finite, the intercept is bounded and can also be taken to converge, say to b . There are the following possibilities.

First, suppose $\alpha \in \mathbb{R} \setminus \{0, \alpha^*\}$. If $\gamma < \infty$, the limit point is a finite policy. The continuity of the objective function shown above thus implies loss convergences to that policy's loss.

If instead $\gamma_n \rightarrow \infty$, $T_n(s_0)$ converges to a finite value $T(s_0)$ for each fixed shock realization: T does not depend on γ , and its other policy arguments converge away

from 0 and α^* . Consequently, each realization eventually satisfies $|T_n(s_0)| \leq \gamma_n$. There, $C_n = P$, so

$$\pi_{0,n} = \frac{P - b_n}{\alpha_n} \longrightarrow \frac{P - b}{\alpha}.$$

The coefficients of T_n remain uniformly bounded, and D_n stays bounded away from zero. Thus, $|q_n| \leq |T_n|$ yields the same integrable bounds used in the fixed-parameter argument above. Dominated convergence then gives convergence of expected $t = 0$ loss to the loss at these limit outcomes, while the expected $t = 1$ loss tends to zero.

Second, suppose $\alpha = 0$. Bounded losses imply $C_n \rightarrow b$ in L^2 . Thus,

$$\mathbb{P}(C_n = P) \leq \mathbb{P}(|P - b| \leq \varepsilon) + \mathbb{P}(|C_n - b| > \varepsilon) \longrightarrow 0$$

taking first $n \rightarrow \infty$ and then $\varepsilon \downarrow 0$. For $\gamma > 0$, $|\vartheta_n| = 1$ whenever $C_n \neq P$, so $\|\vartheta_n\|_2 \rightarrow 1$. Consequently, the fact that

$$\gamma_n \|\vartheta_n\|_2 = \|I_n - C_n\|_2$$

implies the limiting vagueness γ is finite. By continuity at $\alpha = 0$, expected loss converges to the loss of this finite policy.

Third, suppose $\alpha = \alpha^*$. In this case, a finite limit γ would contradict bounded losses. To see this, suppose $\gamma < \infty$ and define $U_* \equiv \alpha^* F_{i_0} + b$. This is a nondegenerate affine Gaussian variable, so $|U_*| > \gamma + \varepsilon$ with positive probability for any $\varepsilon > 0$. On this event, $\alpha_n F_{i_{0,n}} + b_n \rightarrow U_*$ and $\gamma_n \rightarrow \gamma$, so

$$|\alpha_n F_{i_{0,n}} + b_n| - \gamma_n \longrightarrow |U_*| - \gamma > \varepsilon.$$

Since $|q_n| \leq \gamma_n$, for sufficiently large n , the numerator of I_n satisfies

$$|\alpha_n F_{i_{0,n}} + b_n + q_n| \geq |\alpha_n F_{i_{0,n}} + b_n| - \gamma_n > \varepsilon/2$$

Furthermore, its denominator D_n tends to zero. Consequently, $|I_n| \rightarrow \infty$ (on this event).

Since the event has positive probability, Fatou's lemma gives

$$\liminf_n \mathbb{E}[I_n^2] \geq \mathbb{E}[\liminf_n I_n^2] = \infty,$$

contradicting the bound on $\|I_n\|_2$ established above. Thus, it has to be the case

$\gamma_n \rightarrow \infty$.

Even though $D_n \rightarrow 0$, the sequence of thresholds T_n has a finite limit for each shock realization: its denominator $\alpha_n A \rightarrow \alpha^* A = 1$, and its numerator also converges. Since $\gamma_n \rightarrow \infty$, each realization therefore eventually satisfies $|T_n| \leq \gamma_n$. The selected equilibrium is then intermediate, with $C_n = P$, so

$$\pi_{0,n} = \frac{P - b_n}{\alpha_n} \rightarrow A(P - b), \quad I_n = \frac{\pi_{0,n} - F_{i_{0,n}}}{A} \rightarrow P - b - \frac{F_{i_0}}{A}.$$

Output converges accordingly through $y_{0,n} = G_{i_{0,n}} + BI_n$. Since limiting inflation is $AP - Ab$, the $t = 0$ loss associated with these limiting outcomes is $\mathcal{J}_L(i_0, A, -Ab)$.

Furthermore, since $t = 1$ loss is nonnegative, total loss is at least $t = 0$ loss. Applying Fatou's lemma to $t = 0$ loss consequently implies

$$\liminf_n \tilde{\mathcal{J}}(p_n) \geq \mathbb{E} \left[\liminf_n \left(y_{0,n}^2 + \lambda_{cb} \pi_{0,n}^2 \right) \right] = \mathcal{J}_L(i_0, A, -Ab).$$

Thus, every sequence approaching the singular slope has a lower limit of expected loss at least as large as the minimum over the limiting family.

Finally, suppose $|\alpha_n| \rightarrow \infty$. The bound on $|b_n|/(1 + |\alpha_n|)$ established above implies that $-b_n/\alpha_n$ is bounded. We can therefore take a subsequence along which

$$-\frac{b_n}{\alpha_n} \rightarrow h$$

for some finite constant h . Rearranging $C_n = \alpha_n \pi_{0,n} + b_n$ gives

$$\pi_{0,n} = -\frac{b_n}{\alpha_n} + \frac{C_n}{\alpha_n}.$$

Since $\|C_n\|_2$ is bounded and $|\alpha_n| \rightarrow \infty$, the second term converges to zero in L^2 . Inflation therefore converges in L^2 to the constant h . The equilibrium formulas also give convergence of aggregate rate expectations and output in L^2 , with

$$I_n \rightarrow \frac{h - F_{i_0}}{A}.$$

Expected current loss consequently converges to $\mathcal{J}_L(i_0, 0, h)$: these limiting outcomes have constant inflation, corresponding to $k = 0$ in the family defined above. Since the

$t = 1$ loss is nonnegative,

$$\liminf_n \tilde{\mathcal{J}}(p_n) \geq \mathcal{J}_L(i_0, 0, h).$$

Approximating limiting outcomes with finite policies. Fix (i_0, k, h) . We now construct sequences of finite policies whose current outcomes approach those associated with $\pi_L = kP + h$ and whose expected losses approach $\mathcal{J}_L(i_0, k, h)$.

Choose $k_n \rightarrow k$ with $k_n \notin \{0, A\}$, and set

$$\alpha_n = \frac{1}{k_n}, \quad b_n = -\frac{h}{k_n}.$$

Excluding $k_n = 0$ ensures that α_n is finite, while excluding $k_n = A$ ensures that $\alpha_n \neq \alpha^*$. Thus, along the sequence of policies, we have feasible finite data dependence and intercepts, even if the limiting outcome has $k = 0$ or $k = A$.

For each fixed n , the result above for letting vagueness become arbitrarily large while holding the other parameters fixed entails a limiting inflation of

$$\frac{P - b_n}{\alpha_n} = k_n P + h,$$

plus convergence of current outcomes in L^2 and convergence of expected loss to $\mathcal{J}_L(i_0, k_n, h)$.

Thus, we can choose a sequence with finite vagueness $\gamma_n \geq n$ large enough that expected loss differs from $\mathcal{J}_L(i_0, k_n, h)$ by at most $1/n$, and that output and inflation are each within $1/n$ in L^2 of the outcomes associated with $k_n P + h$.

As $k_n \rightarrow k$, these converge in L^2 to the outcomes associated with $kP + h$, and their expected current losses converge to $\mathcal{J}_L(i_0, k, h)$. Consequently, the finite policies $p_n = (i_0, \alpha_n, b_n, \gamma_n)$ satisfy

$$\pi_{0,n} \rightarrow \pi_L, \quad y_{0,n} \rightarrow y_L \quad \text{in } L^2, \quad \tilde{\mathcal{J}}(p_n) \rightarrow \mathcal{J}_L(i_0, k, h).$$

Moreover, $\gamma_n \rightarrow \infty$ because $\gamma_n \geq n$. Since expected current loss and total loss converge to the same value, the expected future constraint loss tends to zero.

Thus, every loss in the limiting family defined above can be approached by feasible finite policies. Since V is the infimum over feasible finite policies, $V \leq \tilde{\mathcal{J}}(p_n)$ for every n . Taking limits therefore gives $V \leq \mathcal{J}_L(i_0, k, h)$.

Attainment and the boundary benchmark. Finally, let

$$V_L := \inf_{i_0, k, h \in \mathbb{R}} \mathcal{J}_L(i_0, k, h).$$

We show that this infimum is attained and compute its value. Let

$$H := B/A, \quad Z := G_0 - HF_0, \quad c_0 := -\sigma + H\sigma\kappa = -\frac{\sigma\beta}{1 + \sigma\kappa + \beta} \neq 0.$$

Since $y_L = Z + c_0 i_0 + H(kP + h)$, we can write

$$\begin{aligned} \mathcal{J}_L(i_0, k, h) = & \text{Var}(Z) + 2Hk \text{Cov}(Z, P) + (H^2 + \lambda_{cb})k^2 \text{Var}(P) \\ & + [\mathbb{E} Z + c_0 i_0 + H(k \mathbb{E} P + h)]^2 + \lambda_{cb}(k \mathbb{E} P + h)^2. \end{aligned}$$

Since (h, i_0) do not affect the variance terms, for each k , it is optimal to use them to set mean inflation and mean output to zero. This requires

$$h = -k \mathbb{E} P, \quad i_0 = -\mathbb{E} Z / c_0.$$

The remaining loss is a quadratic function of k with a strictly positive coefficient on k^2 , so it attains its minimum at a finite k . Hence, some triple (i_0, k, h) attains V_L .

Now, recall that V is the infimum of expected losses over feasible finite policies, while V_L is the minimum over the limiting family. Since finite policies can approach the limiting outcome, we have $V \leq V_L$.

If $V = V_L$, that limiting outcome attains the infimum. The construction above provides finite policies whose vagueness tend to infinity and whose outcomes and losses approach this optimum.

If $V < V_L$, choose a sequence of feasible finite policies p_n such that $\widetilde{\mathcal{J}}(p_n) \rightarrow V$. The bounds and analysis above show that, after taking a subsequence, either the policies converge to a feasible finite policy or their losses satisfy

$$\liminf_n \widetilde{\mathcal{J}}(p_n) \geq V_L.$$

The second possibility is incompatible with convergence of losses to $V < V_L$. Thus, the subsequence converges to a feasible finite policy p . By continuity,

$$\widetilde{\mathcal{J}}(p) = \lim_n \widetilde{\mathcal{J}}(p_n) = V,$$

so p is a global minimizer.

Consequently, the infimum is attained either by a finite policy or by a limiting outcome approached through increasingly vague finite announcements. In particular, if any finite policy has loss below V_L , then $V < V_L$ and a finite global minimizer exists. When $V = V_L$, a finite policy may also attain the same loss, but this is of course not guaranteed.

- (iii) *Differentiability and first-order conditions.* Fix a finite policy $p = (i_0, \alpha, b, \gamma)$ with $\alpha \in (-\infty, \alpha^*) \cup (0, \infty)$ (i.e., a stabilizing policy), and suppose [condition \(8\)](#) holds. First, note that the latter ensures that T is nondegenerate: using $\Gamma(\lambda_{ps}) = \lambda_{ps}\kappa/[\sigma(1+\lambda_{ps}\kappa^2)]$, the coefficients of r_0^n and u_0 in the numerator of T simplify to

$$\begin{aligned} L_r(\alpha) &:= \alpha\sigma\kappa - \rho_r, \\ L_u(\alpha) &:= \alpha \left(1 - \rho_u + \frac{(1 + \sigma\kappa + \beta)\rho_u}{1 + \lambda_{ps}\kappa^2} \right) - \Gamma(\lambda_{ps})\rho_u. \end{aligned}$$

Since the fundamental shocks are independent, this implies

$$\text{Var}(T(s_0; p)) = \frac{L_r(\alpha)^2\sigma_r^2 + L_u(\alpha)^2\sigma_u^2}{(\alpha A_\pi^{ie})^2}.$$

Since $\sigma_r^2, \sigma_u^2 > 0$, this variance equals zero only if both coefficients vanish. The first vanishes only at $\alpha = \rho_r/(\sigma\kappa)$, in which case the second equals

$$L_u\left(\frac{\rho_r}{\sigma\kappa}\right) = \frac{\rho_r[1 + (\sigma\kappa + \beta)\rho_u] - \lambda_{ps}\kappa^2[\rho_u - \rho_r(1 - \rho_u)]}{\sigma\kappa(1 + \lambda_{ps}\kappa^2)} \neq 0,$$

where the inequality follows directly from [condition \(8\)](#). Hence T is nondegenerate at every stabilizing policy.

Now, just as in part (ii) above, let $q := \gamma\vartheta_p$ for $\gamma > 0$ and $q := 0$ for $\gamma = 0$. In the stabilizing regions,

$$q = \min\{\gamma, \max\{-\gamma, T(s_0; p)\}\}.$$

Given the expression for $T(s_0; p)$ in [Proposition 4](#), q is continuous and locally Lipschitz in p . Thus, the numerator in the expression for equilibrium inflation in [Proposition 4](#) is also continuous and locally Lipschitz in policy. Since $\alpha \neq \alpha^*$, the denominator is in turn bounded away from zero on a sufficiently small neighborhood around the policy under consideration. Consequently, since the denominator is also affine in p , its reciprocal is locally Lipschitz. Since inflation is just the product of this reciprocal and

the locally Lipschitz numerator, it is locally Lipschitz too.

Further, for $\gamma > 0$, both $q(s_0; p)$ and $\pi_0(s_0; p)$ are continuously differentiable with respect to (i_0, α, b, γ) whenever $|T(s_0; p)| \neq \gamma$: away from these kinks, q equals one of the smooth expressions $-\gamma$, $T(s_0; p)$, or γ . Since s_0 and T is a nondegenerate affine function of s_0 , these kinks have probability zero.

On a sufficiently small compact neighborhood within a stabilizing region, T and its policy derivatives grow at most linearly in the shocks (uniformly over the neighborhood) because its affine coefficients and their policy derivatives are bounded. Since the minimum and maximum of locally Lipschitz functions are locally Lipschitz, q also has a local Lipschitz bound that also grows at most linearly in the shocks. The equilibrium formulas above map these bounds to output and inflation, which themselves also grow at most linearly in the shocks. The central bank's quadratic loss has thus a local Lipschitz bound that grows at most quadratically in the shocks. This bound is integrable and, together with almost-sure differentiability, justifies differentiation under the expectation by dominated convergence.

For $\gamma > 0$, the continuation-loss term $h(z, \gamma) := (|z| - \gamma)_+^2$ is continuously differentiable in (z, γ) . The expressions for equilibrium inflation above imply that z and its local Lipschitz bound grow at most linearly in the shocks. Taking absolute values and positive parts preserves these bounds, and squaring then gives a local Lipschitz bound for the loss that grows at most quadratically in the shocks. This bound is integrable, so dominated convergence again allows differentiation under the expectation.

The resulting derivatives are continuous in policy. As policies approach p , the expressions inside the expectations converge except at shock realizations s.t. $|T(s_0; p)| = \gamma$, a null event. Thus, we can use the same integrable bound to apply dominated convergence and get that expectations converge. We first give the expressions for $\gamma > 0$ and then specify their right limits at $\gamma = 0$.

Let

$$\tilde{G}_p(s_0) := 2A_y^{i^e} y_0 + 2\lambda_{cb} A_\pi^{i^e} \pi_0.$$

By the chain rule,

$$\begin{aligned} \frac{\partial \tilde{\Psi}}{\partial j} &= \mathbb{E} \left[\tilde{G}_p \frac{\partial i_1^e}{\partial j} \right], \quad j \in \{\alpha, b, \gamma\}, \\ \frac{\partial \tilde{\Psi}}{\partial i_0} &= \mathbb{E} \left[\tilde{G}_p \frac{\partial i_1^e}{\partial i_0} \right] - 2\sigma \mathbb{E} [y_0 + \lambda_{cb} \kappa \pi_0]. \end{aligned}$$

From the equilibrium inflation equation in [Proposition 4](#),

$$\pi_0 = \frac{C_\pi + A_\pi^r r_0^n + A_\pi^u u_0 - \sigma \kappa i_0 + A_\pi^{i^e} (b + \gamma \vartheta_p)}{1 - \alpha A_\pi^{i^e}},$$

we get

$$\begin{aligned} \frac{\partial \pi_0}{\partial i_0} &= \frac{-\sigma \kappa + A_\pi^{i^e} \gamma \frac{\partial \vartheta_p}{\partial i_0}}{1 - \alpha A_\pi^{i^e}}, & \frac{\partial \pi_0}{\partial \alpha} &= \frac{A_\pi^{i^e}}{1 - \alpha A_\pi^{i^e}} \left(\pi_0 + \gamma \frac{\partial \vartheta_p}{\partial \alpha} \right), \\ \frac{\partial \pi_0}{\partial b} &= \frac{A_\pi^{i^e}}{1 - \alpha A_\pi^{i^e}} \left(1 + \gamma \frac{\partial \vartheta_p}{\partial b} \right), & \frac{\partial \pi_0}{\partial \gamma} &= \frac{A_\pi^{i^e}}{1 - \alpha A_\pi^{i^e}} \left(\vartheta_p + \gamma \frac{\partial \vartheta_p}{\partial \gamma} \right). \end{aligned}$$

Thus, since $i_1^e = \alpha \pi_0 + b + \gamma \vartheta_p$,

$$\begin{aligned} \frac{\partial i_1^e}{\partial i_0} &= \frac{1}{1 - \alpha A_\pi^{i^e}} \left(-\alpha \sigma \kappa + \gamma \frac{\partial \vartheta_p}{\partial i_0} \right), \\ \frac{\partial i_1^e}{\partial \alpha} &= \frac{1}{1 - \alpha A_\pi^{i^e}} \left(\pi_0 + \gamma \frac{\partial \vartheta_p}{\partial \alpha} \right), \\ \frac{\partial i_1^e}{\partial b} &= \frac{1}{1 - \alpha A_\pi^{i^e}} \left(1 + \gamma \frac{\partial \vartheta_p}{\partial b} \right), \\ \frac{\partial i_1^e}{\partial \gamma} &= \frac{1}{1 - \alpha A_\pi^{i^e}} \left(\vartheta_p + \gamma \frac{\partial \vartheta_p}{\partial \gamma} \right). \end{aligned}$$

For the $t = 1$ term, let $h(z, \gamma) := (|z| - \gamma)_+^2$. For $j \in \{i_0, \alpha, b\}$,

$$\frac{\partial}{\partial j} \mathbb{E}[h(z_p, \gamma)] = 2 \mathbb{E} \left[(|z_p| - \gamma)_+ \text{sign}(z_p) \frac{\partial z_p}{\partial j} \right],$$

and

$$\frac{\partial}{\partial \gamma} \mathbb{E}[h(z_p, \gamma)] = 2 \mathbb{E} \left[(|z_p| - \gamma)_+ \text{sign}(z_p) \frac{\partial z_p}{\partial \gamma} \right] - 2 \mathbb{E}[(|z_p| - \gamma)_+].$$

Since $\mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0]$ is policy-invariant,

$$\begin{aligned} \frac{\partial z_p}{\partial \alpha} &= \pi_0 + \alpha \frac{\partial \pi_0}{\partial \alpha}, & \frac{\partial z_p}{\partial b} &= 1 + \alpha \frac{\partial \pi_0}{\partial b}, \\ \frac{\partial z_p}{\partial \gamma} &= \alpha \frac{\partial \pi_0}{\partial \gamma}, & \frac{\partial z_p}{\partial i_0} &= \alpha \frac{\partial \pi_0}{\partial i_0}. \end{aligned}$$

Substituting the partial derivatives above into these expressions yields the derivatives of the objective function. At a finite local minimizer in the stabilizing regions under [condition \(8\)](#), the necessary first-order conditions take the following expanded form

(we omit the p subscript for notational simplicity),

$$\begin{aligned}
0 &= \frac{1}{1 - \alpha A_\pi^{ie}} \mathbb{E} \left[\tilde{G} \left(-\alpha \sigma \kappa + \gamma \frac{\partial \vartheta}{\partial i_0} \right) \right] - 2\sigma \mathbb{E} [y_0 + \lambda_{cb} \kappa \pi_0] \\
&\quad + \frac{2W_1 \alpha}{1 - \alpha A_\pi^{ie}} \mathbb{E} \left[(|z| - \gamma)_+ \text{sign}(z) \left(-\sigma \kappa + A_\pi^{ie} \gamma \frac{\partial \vartheta}{\partial i_0} \right) \right], \\
0 &= \frac{1}{1 - \alpha A_\pi^{ie}} \mathbb{E} \left[\tilde{G} \left(\pi_0 + \gamma \frac{\partial \vartheta}{\partial \alpha} \right) \right] \\
&\quad + 2W_1 \mathbb{E} \left[(|z| - \gamma)_+ \text{sign}(z) \left(\pi_0 + \alpha \frac{\partial \pi_0}{\partial \alpha} \right) \right], \\
0 &= \frac{1}{1 - \alpha A_\pi^{ie}} \mathbb{E} \left[\tilde{G} \left(1 + \gamma \frac{\partial \vartheta}{\partial b} \right) \right] \\
&\quad + 2W_1 \mathbb{E} \left[(|z| - \gamma)_+ \text{sign}(z) \left(1 + \alpha \frac{\partial \pi_0}{\partial b} \right) \right], \\
0 &\leq \frac{1}{1 - \alpha A_\pi^{ie}} \mathbb{E} \left[\tilde{G} \left(\vartheta + \gamma \frac{\partial \vartheta}{\partial \gamma} \right) \right] \\
&\quad + 2W_1 \mathbb{E} \left[(|z| - \gamma)_+ \text{sign}(z) \alpha \frac{\partial \pi_0}{\partial \gamma} \right] - 2W_1 \mathbb{E} [(|z| - \gamma)_+].
\end{aligned} \tag{21}$$

For the $\gamma = 0$ case, let $\gamma \downarrow 0$ while holding the other policy coordinates fixed. For any shock realization with $T(s_0; p) \neq 0$, at sufficiently small positive γ we have

$$q(s_0; p) = \gamma \text{sign}(T(s_0; p)).$$

Since the sign of T is locally constant away from zero,

$$\frac{\partial q}{\partial j} = \gamma \frac{\partial \vartheta}{\partial j} \longrightarrow 0, \quad j \in \{i_0, \alpha, b\},$$

while

$$\frac{\partial q}{\partial \gamma} = \vartheta + \gamma \frac{\partial \vartheta}{\partial \gamma} \longrightarrow \text{sign}(T).$$

Thus, at $\gamma = 0$,

$$\frac{\partial q}{\partial i_0} = \frac{\partial q}{\partial \alpha} = \frac{\partial q}{\partial b} = 0, \quad \frac{\partial^+ q}{\partial \gamma} = \text{sign}(T(s_0; \alpha, b)).$$

These identities hold almost surely since $\mathbb{P}(T = 0) = 0$. The same limits hold when the other policy coordinates also converge to their values at a policy in the same stabilizing

region: continuity of T and its nonzero limiting value imply that eventually $|T| > \gamma$ and its sign remains unchanged.

Moreover, writing z_0 for the value of z at the limiting policy,

$$(|z| - \gamma)_+ \text{sign}(z) \longrightarrow z_0, \quad (|z| - \gamma)_+ \longrightarrow |z_0|.$$

The uniform bounds above justify taking these limits under the expectation. This proves the continuous extension of the first derivatives of the central bank's loss to $\gamma = 0$, with the γ -derivative interpreted from the right.

Finally, at a local minimum, the derivative of the loss with respect to γ is zero when $\gamma > 0$, and its right derivative is nonnegative when $\gamma = 0$, as required by the constraint $\gamma \geq 0$ (complementary slackness conditions).

□

C Second-Order Conditions: Exogenous Signals

This section provides easily checkable conditions that guarantee candidates that satisfy the first-order conditions in [Proposition 3](#) are local minima and thus the study of the optimal policy via the first-order characterization is well grounded. The proofs of the lemmas used in this section can be found in [Appendix D](#).

Throughout this section, assume that $\sigma_r^2, \sigma_u^2, \tilde{\sigma}^2, \hat{\sigma}^2 > 0$ and $\rho_r \rho_u \neq 0$.

Lemma 6. *Let $\mathcal{D} := \{(\alpha, b, \gamma) \in \mathbb{R}^3 : \gamma > 0, \alpha \neq a_{ps}\}$ and*

$$\mathcal{J}(\alpha, b, \gamma) = \Psi(\alpha, b, \gamma) + W_1 \mathbb{E} \left[(|z_{\alpha,b}| - \gamma)_+^2 \right], \quad z_{\alpha,b} := (\alpha - a_{cb})\tilde{s}_1 + (b - a_{cb}^0),$$

with Ψ and all constants as in [Propositions 1 and 3](#), and $i_1^e = \alpha\mu_{\tilde{s}} + b + \gamma\theta_{\alpha,b}$ as in [Lemma 2](#). Then, $\mathcal{J} \in C^2(\mathcal{D})$.

Moreover, the second derivatives of the $t=1$ term admit the closed forms

$$\begin{aligned} \frac{\partial^2}{\partial \alpha^2} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= 2W_1 \mathbb{E} \left[\mathbf{1}_{\{|z| > \gamma\}} \tilde{s}_1^2 \right], \\ \frac{\partial^2}{\partial b^2} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= 2W_1 \mathbb{E} \left[\mathbf{1}_{\{|z| > \gamma\}} \right], \\ \frac{\partial^2}{\partial \alpha \partial b} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= 2W_1 \mathbb{E} \left[\mathbf{1}_{\{|z| > \gamma\}} \tilde{s}_1 \right], \\ \frac{\partial^2}{\partial \gamma^2} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= 2W_1 \mathbb{P}(|z| > \gamma), \end{aligned}$$

$$\begin{aligned}\frac{\partial^2}{\partial \alpha \partial \gamma} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= -2W_1 \mathbb{E}[\mathbf{1}_{\{|z| > \gamma\}} \text{sign}(z) \tilde{s}_1], \\ \frac{\partial^2}{\partial b \partial \gamma} W_1 \mathbb{E}[(|z| - \gamma)_+^2] &= -2W_1 \mathbb{E}[\mathbf{1}_{\{|z| > \gamma\}} \text{sign}(z)],\end{aligned}$$

and the mixed partials are symmetric.

For the $t = 0$ term, let $I := \alpha \mu_{\tilde{s}} + b + \gamma \theta_{\alpha, b}$, $\xi_i := 2 \left((A_y^{i^e})^2 + \lambda_{cb} (A_\pi^{i^e})^2 \right)$ and let $G_{\alpha, b, \gamma}$ be defined as in [Lemma 5](#). For any $p, q \in \{\alpha, b, \gamma\}$

$$\frac{\partial^2 \Psi}{\partial p \partial q} = \xi_i \mathbb{E} \left[\frac{\partial I}{\partial p} \frac{\partial I}{\partial q} \right] - \frac{2\delta_{i^e}^2}{1 + \lambda_{cb} \kappa^2} \mathbb{E} \left[\frac{\partial I}{\partial p} \right] \mathbb{E} \left[\frac{\partial I}{\partial q} \right] + \mathbb{E} \left[G_{\alpha, b, \gamma}(s_0) \frac{\partial^2 I}{\partial p \partial q} \right], \quad (22)$$

where

$$\frac{\partial I}{\partial \alpha} = \mu_{\tilde{s}} + \gamma \frac{\partial \theta_{\alpha, b}}{\partial \alpha}, \quad \frac{\partial I}{\partial b} = 1 + \gamma \frac{\partial \theta_{\alpha, b}}{\partial b}, \quad \frac{\partial I}{\partial \gamma} = \theta_{\alpha, b},$$

and

$$\begin{aligned}\frac{\partial^2 I}{\partial p \partial q} &= \gamma \frac{\partial^2 \theta_{\alpha, b}}{\partial p \partial q}, & p, q \in \{\alpha, b\}, \\ \frac{\partial^2 I}{\partial p \partial \gamma} &= \frac{\partial^2 I}{\partial \gamma \partial p} = \frac{\partial \theta_{\alpha, b}}{\partial p}, & p \in \{\alpha, b\}, \\ \frac{\partial^2 I}{\partial \gamma^2} &= 0.\end{aligned}$$

with

$$\theta_{\alpha, b}(s_0) = 2\Phi(\nu(s_0)) - 1, \quad \nu(s_0) := \frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + b - a_{ps}^0(s_0)}{|\alpha - a_{ps}|\sigma_{\tilde{s}}},$$

and

$$\frac{\partial \theta_{\alpha, b}}{\partial p} = 2\phi(\nu) \frac{\partial \nu}{\partial p}, \quad \frac{\partial^2 \theta_{\alpha, b}}{\partial p \partial q} = 2 \left[-\nu \phi(\nu) \frac{\partial \nu}{\partial p} \frac{\partial \nu}{\partial q} + \phi(\nu) \frac{\partial^2 \nu}{\partial p \partial q} \right].$$

All second derivatives above are continuous on $\mathcal{D}$.

Lemma 7. Let $\overline{\mathcal{D}} := \{(\alpha, b, \gamma) : \gamma \geq 0, \alpha \neq a_{ps}\}$ and define $\mathcal{J}$ as in [Lemma 5](#). The objective function and its first derivatives are jointly continuous on $\overline{\mathcal{D}}$ (interpreting derivatives with respect to γ from the right at $\gamma = 0$). Furthermore, all second derivatives in [Lemma 6](#) admit jointly continuous extensions to $\overline{\mathcal{D}}$. At $\gamma = 0$, these extensions equal the corresponding right derivatives. In particular,

$$\frac{\partial_+^2}{\partial \gamma^2} \mathcal{J}(\alpha, b, 0) = \lim_{\gamma \downarrow 0} \frac{\partial^2 \mathcal{J}}{\partial \gamma^2}(\alpha, b, \gamma) = \xi_i \sigma_\theta^2 + \frac{2\lambda_{cb}}{1 + \lambda_{cb} \kappa^2} (\beta \sigma \kappa)^2 \mathbb{E}[\theta_{\alpha, b}]^2 + 2W_1 > 0,$$

where $\sigma_\theta^2 := \text{Var}(\theta_{\alpha, b})$. Similarly, the mixed right second derivatives with respect to γ exist

and equal the limits of the interior mixed second derivatives,

$$\begin{aligned}\frac{\partial_+^2 \mathcal{J}(\alpha, b, 0)}{\partial \alpha \partial \gamma} &= \lim_{\gamma \downarrow 0} \frac{\partial^2 \mathcal{J}}{\partial \alpha \partial \gamma}(\alpha, b, \gamma) \\ &= \xi_i \mathbb{E}[\mu_{\tilde{s}} \theta_{\alpha, b}] - \frac{2\delta_{ie}^2}{1 + \lambda_{cb} \kappa^2} \mathbb{E}[\mu_{\tilde{s}}] \mathbb{E}[\theta_{\alpha, b}] + \mathbb{E}\left[G_0 \frac{\partial \theta_{\alpha, b}}{\partial \alpha}\right] - 2W_1 \mathbb{E}[\text{sign}(z_{\alpha, b}) \tilde{s}_1], \\ \frac{\partial_+^2 \mathcal{J}(\alpha, b, 0)}{\partial b \partial \gamma} &= \lim_{\gamma \downarrow 0} \frac{\partial^2 \mathcal{J}}{\partial b \partial \gamma}(\alpha, b, \gamma) \\ &= \xi_i \mathbb{E}[\theta_{\alpha, b}] - \frac{2\delta_{ie}^2}{1 + \lambda_{cb} \kappa^2} \mathbb{E}[\theta_{\alpha, b}] + \mathbb{E}\left[G_0 \frac{\partial \theta_{\alpha, b}}{\partial b}\right] - 2W_1 \mathbb{E}[\text{sign}(z_{\alpha, b})],\end{aligned}$$

where $G_0 = G_{\alpha, b, 0}$ and all the other terms are defined as in [Lemma 6](#). Finally, the mixed right derivatives are symmetric.

Proposition 6. Let $\overline{\mathcal{D}} := \{(\alpha, b, \gamma) : \gamma \geq 0, \alpha \neq a_{ps}\}$ and suppose $\mathcal{J}$ is as in [Lemma 5](#). Let $(\alpha^*, b^*, \gamma^*) \in \overline{\mathcal{D}}$ be a local candidate. Consider the following two cases.

*I. **Interior** ($\gamma^* > 0$).* Assume the interior first-order conditions $\nabla \mathcal{J}(\alpha^*, b^*, \gamma^*) = 0$ hold and write the Hessian at $(\alpha^*, b^*, \gamma^*)$ as

$$H = \nabla^2 \mathcal{J} = \begin{pmatrix} H_{xx} & H_{x\gamma} \\ H_{\gamma x} & H_{\gamma\gamma} \end{pmatrix}, \quad x := (\alpha, b).$$

Let $c_\gamma := H_{\gamma\gamma} > 0$ and define the Schur complement

$$S := H_{xx} - H_{x\gamma} c_\gamma^{-1} H_{\gamma x}.$$

Then,

$$H \succ 0 \iff c_\gamma > 0 \text{ and } S \succ 0.$$

These equivalent curvature conditions imply that the candidate $(\alpha^*, b^*, \gamma^*)$ is an isolated strict local minimizer.

*II. **Boundary** ($\gamma^* = 0$).* Assume the KKT conditions in [Proposition 3](#) with $\gamma^* = 0$ hold.

Let

$$R := \nabla_{(\alpha, b)}^2 \mathcal{J}(\alpha^*, b^*, 0), \quad c_\gamma := \frac{\partial_+^2 \mathcal{J}(\alpha^*, b^*, 0)}{\partial \gamma^2} > 0,$$

and

$$H_{x\gamma}^+ := \begin{pmatrix} \frac{\partial_+^2 \mathcal{J}}{\partial \alpha \partial \gamma} \\ \frac{\partial_+^2 \mathcal{J}}{\partial b \partial \gamma} \end{pmatrix}(\alpha^*, b^*, 0).$$

(a) **Strict complementarity.** If $\partial\mathcal{J}/\partial\gamma > 0$ at $(\alpha^*, b^*, 0)$, then $(\alpha^*, b^*, 0)$ is an isolated strict local minimizer.

(b) **Degenerate corner.** If $\partial\mathcal{J}/\partial\gamma = 0$ at $(\alpha^*, b^*, 0)$, then $(\alpha^*, b^*, 0)$ is an isolated strict local minimizer of the constrained problem if

$$c_\gamma > (H_{x\gamma}^+)^\top R^{-1} H_{x\gamma}^+.$$

Proof.

I. Interior.

By Lemma 6, $\mathcal{J} \in C^2(\mathcal{D})$ and H is symmetric. For any block symmetric matrix with scalar bottom-right block $c_\gamma > 0$, positive definiteness is equivalent to positive definiteness of the Schur complement, i.e.,

$$H \succ 0 \iff c_\gamma > 0 \text{ and } S := H_{xx} - H_{x\gamma} c_\gamma^{-1} H_{\gamma x} \succ 0.$$

Further, from Lemma 6 and using that $A_\pi^{i^e} - \kappa A_y^{i^e} = -\beta\sigma\kappa$, we have

$$H_{\gamma\gamma} = \frac{\partial^2 \mathcal{J}}{\partial \gamma^2} = \xi_i \text{Var}(\theta_{\alpha,b}) + \frac{2\lambda_{cb}}{1 + \lambda_{cb}\kappa^2} (\beta\sigma\kappa)^2 \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma) > 0.$$

Hence, $S \succ 0$ implies $H \succ 0$ so the stated inequality guarantees $H \succ 0$. This implies that the candidate $(\alpha^*, b^*, \gamma^*)$ is a strict local minimum (standard KKT sufficiency).

II. Boundary.

Note first that at the boundary with $\gamma = 0$, $I = \alpha\mu_{\tilde{s}} + b$ is affine in (α, b) . Using (22) and the $t = 1$ loss term $W_1 \mathbb{E}[z_{\alpha,b}^2]$, for any nonzero $u = (u_1, u_2)^\top$ we have

$$u^\top R u = \xi_i u_1^2 \text{Var}(\mu_{\tilde{s}}) + \frac{2\lambda_{cb}(\beta\sigma\kappa)^2}{1 + \lambda_{cb}\kappa^2} (u_1 \mathbb{E}[\mu_{\tilde{s}}] + u_2)^2 + 2W_1 \mathbb{E}[(u_1 \tilde{s}_1 + u_2)^2] > 0.$$

The first two terms are nonnegative, and the last is strictly positive because $\text{Var}(\tilde{s}_1) \geq \tilde{\sigma}^2 > 0$ and $W_1 > 0$. Hence, $R \succ 0$.

Now, fix a feasible perturbation $(u, t) \in \mathbb{R}^2 \times [0, \infty)$, let $x^* := (\alpha^*, b^*)$, $h := H_{x\gamma}^+$, and $\psi(s) := \mathcal{J}(x^* + su, st)$. By Lemma 7, the right second derivative exists and equals

$$\psi''(0^+) = u^\top R u + 2tu^\top h + t^2 c_\gamma.$$

(a) Under strict complementarity, let $g := \mathcal{J}_\gamma(x^*, 0) > 0$. Since $\nabla_x \mathcal{J}(x^*, 0) = 0$ (i.e., the first-order conditions for unconstrained α and b hold), we have $\psi'(0^+) = gt$.

Thus, perturbations with $t > 0$ raise losses to first order. More precisely, continuity of $\mathcal{J}_\gamma$ implies $\mathcal{J}_\gamma(x, \gamma) \geq g/2$ in a sufficiently small feasible neighborhood of $(x^*, 0)$. Thus,

$$\psi(1) - \psi(0) \geq \mathcal{J}(x^* + u, 0) - \mathcal{J}(x^*, 0) + \frac{g}{2}t.$$

Since $\nabla_x \mathcal{J}(x^*, 0) = 0$ and $R \succ 0$, there exists $c > 0$ such that, for sufficiently small u ,

$$\mathcal{J}(x^* + u, 0) - \mathcal{J}(x^*, 0) \geq c\|u\|^2.$$

Consequently,

$$\psi(1) - \psi(0) \geq c\|u\|^2 + \frac{g}{2}t > 0$$

for every sufficiently small nonzero feasible perturbation. The boundary restriction is strictly convex locally, and $\mathcal{J}_\gamma > 0$ excludes nearby interior local minima, so this local minimizer is isolated.

- (b) In this case, all first-order derivatives vanish. Since $R \succ 0$, completing the square gives

$$\psi''(0^+) = (u + tR^{-1}h)^\top R(u + tR^{-1}h) + t^2(c_\gamma - h^\top R^{-1}h).$$

Thus, if $c_\gamma > h^\top R^{-1}h$, the quadratic form above is positive definite. In particular, there exists $m > 0$ such that

$$\psi''(0^+) \geq m(\|u\|^2 + t^2).$$

Using the continuous extension of the second derivatives to $\gamma = 0$, the second-order Taylor expansion for feasible directions of perturbations gives

$$\psi(1) = \psi(0) + \frac{1}{2}\psi''(0^+) + o(\|u\|^2 + t^2).$$

for $(u, t) \rightarrow (0, 0)$, $t \geq 0$. Consequently, for every sufficiently small nonzero feasible perturbation,

$$\mathcal{J}(x^* + u, t) - \mathcal{J}(x^*, 0) \geq \frac{m}{4}(\|u\|^2 + t^2) > 0.$$

Moreover, continuity of the Hessian preserves positive definiteness in a sufficiently small feasible neighborhood, so the local minimizer is isolated. This proves the stated sufficient condition.

□

D Proofs of Additional Technical Lemmas

D.1 Proof of Lemma 3

Proof. Let $x := \begin{pmatrix} r_1^n & u_1 \end{pmatrix}^T$. Conditional on (r_0^n, u_0) , we have

$$x \mid (r_0^n, u_0) \sim \mathcal{N}(\mu_{|0}, P_{|0}), \quad \mu_{|0} := \begin{pmatrix} (1 - \rho_r)\bar{r}^n + \rho_r r_0^n \\ \rho_u u_0 \end{pmatrix}, \quad P_{|0} := \text{diag}(\sigma_{\epsilon_r}^2, \sigma_{\epsilon_u}^2),$$

Define

$$\tilde{\Omega} := \begin{pmatrix} \tilde{\omega}_r & \tilde{\omega}_u \end{pmatrix}, \quad \hat{\Omega} := \begin{pmatrix} \hat{\omega}_r & \hat{\omega}_u \end{pmatrix},$$

and

$$V := \begin{pmatrix} \sigma_{\tilde{\epsilon}}^2 & 0 \\ 0 & \sigma_{\hat{\epsilon}}^2 \end{pmatrix}.$$

Conditionable signal only. We observe $\tilde{s}_1 = \tilde{\Omega}x + \tilde{\epsilon}$. For jointly Gaussian $(x, \tilde{s}_1)$, the conditional mean is

$$\mathbb{E}[x \mid \tilde{s}_1] = \mu_{|0} + \Sigma_{x\tilde{s}}\Sigma_{\tilde{s}\tilde{s}}^{-1}(\tilde{s}_1 - \mathbb{E}[\tilde{s}_1]),$$

where $\Sigma_{x\tilde{s}} = P_{|0}\tilde{\Omega}^\top$ and $\Sigma_{\tilde{s}\tilde{s}} = \tilde{\Omega}P_{|0}\tilde{\Omega}^\top + \sigma_{\tilde{\epsilon}}^2$. Hence,

$$\mathbb{E}[r_1^n \mid r_0^n, u_0, \tilde{s}_1] = \frac{\tilde{\omega}_r \sigma_{\epsilon_r}^2}{\Sigma_{\tilde{s}\tilde{s}}} \tilde{s}_1 + \mu_{r|0} - \frac{\tilde{\omega}_r \sigma_{\epsilon_r}^2}{\Sigma_{\tilde{s}\tilde{s}}} \tilde{\Omega} \mu_{|0}, \quad \mathbb{E}[u_1 \mid r_0^n, u_0, \tilde{s}_1] = \frac{\tilde{\omega}_u \sigma_{\epsilon_u}^2}{\Sigma_{\tilde{s}\tilde{s}}} \tilde{s}_1 + \mu_{u|0} - \frac{\tilde{\omega}_u \sigma_{\epsilon_u}^2}{\Sigma_{\tilde{s}\tilde{s}}} \tilde{\Omega} \mu_{|0}.$$

Consequently,

$$\mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid \tilde{s}_1, r_0^n, u_0] = a_{ps} \tilde{s}_1 + a_{ps}^0(r_0^n, u_0),$$

where

$$a_{ps} = \frac{\tilde{\omega}_r \sigma_{\epsilon_r}^2 + \Gamma(\lambda_{ps}) \tilde{\omega}_u \sigma_{\epsilon_u}^2}{\Sigma_{\tilde{s}\tilde{s}}}$$

and

$$a_{ps}^0(r_0^n, u_0) = \mu_{r|0} + \Gamma(\lambda_{ps}) \mu_{u|0} - \frac{\tilde{\omega}_r \sigma_{\epsilon_r}^2 + \Gamma(\lambda_{ps}) \tilde{\omega}_u \sigma_{\epsilon_u}^2}{\Sigma_{\tilde{s}\tilde{s}}} \tilde{\Omega} \mu_{|0},$$

where the latter is evidently an affine function of (r_0^n, u_0) .

Both signals. Let $y := \begin{pmatrix} \tilde{s}_1 \\ \hat{s}_1 \end{pmatrix} = Hx + \varepsilon$, with $H := \begin{pmatrix} \tilde{\Omega} \\ \hat{\Omega} \end{pmatrix}$ and $\varepsilon \sim \mathcal{N}(0, V)$. Again, by standard results given the Gaussian structure,

$$\mathbb{E}[x \mid y] = \mu_{|0} + K(y - H\mu_{|0}), \quad K := P_{|0}H^\top (HP_{|0}H^\top + V)^{-1}.$$

Let $S := HP_0H^\top + V$ with entries

$$\begin{aligned} S_{11} &= \tilde{\omega}_r^2 \sigma_{\epsilon_r}^2 + \tilde{\omega}_u^2 \sigma_{\epsilon_u}^2 + \sigma_{\epsilon}^2, \\ S_{22} &= \hat{\omega}_r^2 \sigma_{\epsilon_r}^2 + \hat{\omega}_u^2 \sigma_{\epsilon_u}^2 + \sigma_{\epsilon}^2, \\ S_{12} &= \tilde{\omega}_r \hat{\omega}_r \sigma_{\epsilon_r}^2 + \tilde{\omega}_u \hat{\omega}_u \sigma_{\epsilon_u}^2, \end{aligned}$$

and $\Delta := S_{11}S_{22} - S_{12}^2 > 0$. Then

$$S^{-1} = \frac{1}{\Delta} \begin{pmatrix} S_{22} & -S_{12} \\ -S_{12} & S_{11} \end{pmatrix}, \quad P_0H^\top = \begin{pmatrix} \sigma_{\epsilon_r}^2 \tilde{\omega}_r & \sigma_{\epsilon_r}^2 \hat{\omega}_r \\ \sigma_{\epsilon_u}^2 \tilde{\omega}_u & \sigma_{\epsilon_u}^2 \hat{\omega}_u \end{pmatrix}.$$

Writing $\eta := y - H\mu_{|0} = \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix}$, with $\eta_1 = \tilde{s}_1 - \tilde{\omega}_r \mu_{r|0} - \tilde{\omega}_u \mu_{u|0}$, $\eta_2 = \hat{s}_1 - \hat{\omega}_r \mu_{r|0} - \hat{\omega}_u \mu_{u|0}$, the first component of $\mu_{|0} + K\eta$ is

$$\mathbb{E}[r_1^n \mid r_0^n, u_0, \tilde{s}_1, \hat{s}_1] = \mu_{r|0} + \frac{\sigma_{\epsilon_r}^2}{\Delta} [(\tilde{\omega}_r S_{22} - \hat{\omega}_r S_{12})\eta_1 + (-\tilde{\omega}_r S_{12} + \hat{\omega}_r S_{11})\eta_2],$$

and the second component is

$$\mathbb{E}[u_1 \mid r_0^n, u_0, \tilde{s}_1, \hat{s}_1] = \mu_{u|0} + \frac{\sigma_{\epsilon_u}^2}{\Delta} [(\tilde{\omega}_u S_{22} - \hat{\omega}_u S_{12})\eta_1 + (-\tilde{\omega}_u S_{12} + \hat{\omega}_u S_{11})\eta_2].$$

Consequently,

$$\mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid \hat{s}_1, \tilde{s}_1, r_0^n, u_0] = a_{cb} \tilde{s}_1 + a_{cb}^0(\hat{s}_1, r_0^n, u_0),$$

where

$$a_{cb} = \frac{\sigma_{\epsilon_r}^2 (\tilde{\omega}_r S_{22} - \hat{\omega}_r S_{12}) + \Gamma(\lambda_{cb}) \sigma_{\epsilon_u}^2 (\tilde{\omega}_u S_{22} - \hat{\omega}_u S_{12})}{\Delta}$$

and an affine $a_{cb}^0(\hat{s}_1, r_0^n, u_0)$. □

D.2 Proof of Lemma 4

Proof. Fix a realization s of $\tilde{s}_1$. Using $y_1 = -\sigma(i_1 - r_1^n)$ and $\pi_1 = \kappa y_1 + u_1$, the quadratic loss inside (1) can be written as

$$y_1^2 + \lambda_{ps} \pi_1^2 = \sigma^2 (1 + \lambda_{ps} \kappa^2) (i_1 - i_1^{\text{FI}}(s_1, \lambda_{ps}))^2 + \text{const},$$

where $i_1^{\text{FI}}(s_1, \lambda_{ps}) := r_1^n + \Gamma(\lambda_{ps})u_1$ and the second term does not depend on i_1 . Hence, for each $\tilde{s}_1 = s$, the worst-case choice of $a \in [L(s), U(s)]$ maximizes the expected squared

deviation conditional on $(s_0, \tilde{s}_1)$:

$$a \mapsto \mathbb{E} \left[(a - i_1^{\text{FI}}(s_1, \lambda_{ps}))^2 \mid s_0, \tilde{s}_1 = s \right] = (a - m_{ps}(s, s_0))^2 + \text{Var}(i_1^{\text{FI}} \mid s_0, \tilde{s}_1 = s),$$

where $m_{ps}(s, s_0) := \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0, \tilde{s}_1 = s] = a_{ps}s + a_{ps}^0(s_0)$. Since the variance does not depend on a , the maximizer is the endpoint farthest from m_{ps} . Writing $c(s) := \frac{L(s)+U(s)}{2}$ and $\gamma(s) := \frac{U(s)-L(s)}{2} \geq 0$,

$$i_1^{\text{w}}(s, s_0) = c(s) + \gamma(s) \text{sign} \left(c(s) - a_{ps}s - a_{ps}^0(s_0) \right),$$

with the convention that if $a_{ps}s + a_{ps}^0(s_0) = c(s)$ the upper endpoint is selected. Taking conditional expectations over $\tilde{s}_1$ given s_0 then yields

$$i_1^e(s_0) = \mathbb{E}[i_1^{\text{w}} \mid s_0] = \mathbb{E}[c(\tilde{s}_1) \mid s_0] + \mathbb{E}[\gamma(\tilde{s}_1) \text{sign} (c(\tilde{s}_1) - a_{ps}\tilde{s}_1 - a_{ps}^0(s_0)) \mid s_0],$$

as stated. $\square$

D.3 Proof of Lemma 5

Proof. Fix $(\alpha_0, b_0, \gamma_0) \in \mathcal{D}$. Further, let $s_0 = (r_0^n, u_0)$ and set $\varepsilon := |\alpha_0 - a_{ps}|/2 > 0$ and choose a neighborhood

$$\mathcal{N} = \{(\alpha, b, \gamma) : |\alpha - \alpha_0| \leq \varepsilon, |b - b_0| \leq 1, \gamma \in [\gamma_0/2, 2\gamma_0]\}.$$

Then, $|\alpha - a_{ps}| \geq \varepsilon$ and $\gamma \geq \gamma_0/2$ for all $(\alpha, b, \gamma) \in \mathcal{N}$.

For clarity, we consider the $t = 1$ and $t = 0$ components of the objective function separately.

(i) $t = 1$ component. Let $h(z, \gamma) := (|z| - \gamma)_+^2$. Clearly $h \in C^1(\mathbb{R} \times (0, \infty))$ with

$$\frac{\partial h(z, \gamma)}{\partial z} = 2(|z| - \gamma)_+ \text{sign } z, \quad \frac{\partial h(z, \gamma)}{\partial \gamma} = -2(|z| - \gamma)_+.$$

These formulas hold for every z , including $z = \pm\gamma$. Here, $z = (\alpha - a_{cb})\tilde{s}_1 + (b - a_{cb}^0)$ is an affine function of jointly Gaussian variables and thus has a finite second moment.

For $(\alpha, b, \gamma) \in \mathcal{N}$, we have $(|z| - \gamma)_+ \leq |z|$ and consequently

$$|(|z| - \gamma)_+ \tilde{s}_1 \text{sign } z| \leq |z| |\tilde{s}_1|,$$

Further, using the triangle inequality and the definition of $\mathcal{N}$, there exists a C_1 such

that

$$|z| \leq |\alpha - a_{cb}| |\tilde{s}_1| + |b| + |a_{cb}^0| \leq C_1 (1 + |\tilde{s}_1| + |a_{cb}^0|).$$

Moreover, since a_{cb}^0 is affine in $(r_0^n, u_0, \hat{s}_1)$, there exists C_2 such that

$$|a_{cb}^0| \leq C_2 (1 + |r_0^n| + |u_0| + |\hat{s}_1|).$$

Note that both C_1 and C_2 depend on $\mathcal{N}$ (and thus $(\alpha_0, b_0, \gamma_0)$) only. Consequently, using the Cauchy-Schwarz inequality,

$$|(|z| - \gamma)_+ \text{sign } z \tilde{s}_1| \leq C (1 + |\tilde{s}_1|^2 + |\hat{s}_1|^2 + |r_0^n|^2 + |u_0|^2),$$

$$(|z| - \gamma)_+ \leq C (1 + |\tilde{s}_1| + |\hat{s}_1| + |r_0^n| + |u_0|),$$

for some constant C that does not depend on the particular (α, b, γ) picked (it depends only on the neighborhood and thus on its center $(\alpha_0, b_0, \gamma_0)$). Because all random variables are Gaussian, the right hand sides are integrable. Consequently, by the Dominated Convergence Theorem and the derivatives of h computed above,

$$\frac{\partial}{\partial \alpha} [W_1 \mathbb{E} h(z, \gamma)] = 2W_1 \mathbb{E} [(|z| - \gamma)_+ \text{sign } z \tilde{s}_1],$$

$$\frac{\partial}{\partial b} [W_1 \mathbb{E} h(z, \gamma)] = 2W_1 \mathbb{E} [(|z| - \gamma)_+ \text{sign } z],$$

$$\frac{\partial}{\partial \gamma} [W_1 \mathbb{E} h(z, \gamma)] = -2W_1 \mathbb{E} [(|z| - \gamma)_+],$$

and these are continuous on $\mathcal{N}$ (again by dominated convergence given $h \in C^1$).

- (ii) $t = 0$ component. Note that plugging the form of i_1^e for an arbitrary affine interval announcement into the $t = 0$ loss in [Proposition 1](#) yields

$$\Psi = \mathbb{E} \left[QI^2 + 2HI - \frac{(\delta_r \bar{r}^n + \delta_c + \delta_{i^e} m_I)^2}{1 + \lambda_{cb} \kappa^2} \right] + \text{t.i.p.},$$

where $I = \alpha \mu_{\tilde{s}} + b + \gamma \theta_{\alpha, b}$, $m_I := \mathbb{E}[I]$,

$$Q := (A_y^{i^e})^2 + \lambda_{cb} (A_\pi^{i^e})^2 \quad H(s_0) := A_y^{i^e} (C_y + A_y^r r_0^n + A_y^u u_0) + \lambda_{cb} A_\pi^{i^e} (C_\pi + A_\pi^r r_0^n + A_\pi^u u_0).$$

Further, it's straightforward to verify that

$$\theta_{\alpha, b}(s_0) = 2\Phi(\nu(s_0)) - 1, \quad \nu(s_0) := \frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + (b - a_{ps}^0(s_0))}{|\alpha - a_{ps}|\sigma_{\tilde{s}}},$$

with constant $\sigma_{\tilde{s}}^2 := \text{Var}(\tilde{s}_1 \mid s_0) > 0$. Clearly, I is continuously differentiable w.r.t (α, b, γ) for all s_0 , and we have

$$\frac{\partial I}{\partial \alpha} = \mu_{\tilde{s}} + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha}, \quad \frac{\partial I}{\partial b} = 1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b}, \quad \frac{\partial I}{\partial \gamma} = \theta_{\alpha,b} \quad (23)$$

with

$$\begin{aligned} \frac{\partial \theta_{\alpha,b}}{\partial b} &= \frac{2}{|\alpha - a_{ps}| \sigma_{\tilde{s}}} \phi(\nu), \\ \frac{\partial \theta_{\alpha,b}}{\partial \alpha} &= 2\phi(\nu(s_0)) \left(\frac{(b - a_{ps}^0(s_0)) \text{sign}(a_{ps} - \alpha)}{(\alpha - a_{ps})^2 \sigma_{\tilde{s}}} \right). \end{aligned}$$

Moreover, $QI^2 + 2HI$ is trivially continuously differentiable w.r.t. I for all s_0 . To verify that $(\alpha, b, \gamma) \mapsto m_I$ is continuously differentiable on $\mathcal{N}$ (and thus complete the proof that the expression inside the expectation in Ψ is continuously differentiable in (α, b, γ) pointwise), note first that I is of course continuously differentiable w.r.t. (α, b, γ) . Moreover, on $\mathcal{N}$, $\theta_{\alpha,b}$ is bounded by 1 and by definition of $\mathcal{N}$ and the fact that $\phi(x) \leq \phi(0) = 1/\sqrt{2\pi}$ for all $x \in \mathbb{R}$,

$$\left| \frac{\partial \theta_{\alpha,b}}{\partial b} \right| \leq \frac{2}{\varepsilon \sigma_{\tilde{s}} \sqrt{2\pi}} =: c_b,$$

$$\left| \frac{\partial \theta_{\alpha,b}}{\partial \alpha} \right| \leq \frac{2}{\sqrt{2\pi}} \left(\frac{1}{\varepsilon^2 \sigma_{\tilde{s}}} |b - a_{ps}^0(s_0)| \right) \leq c_\alpha (1 + |r_0^n| + |u_0|),$$

for constants c_b, c_α depending only on $\mathcal{N}$ and parameters (the rightmost inequality uses that a_{ps}^0 is affine in s_0). Thus, $\partial I / \partial b$ is bounded by a constant on $\mathcal{N}$ and $\partial I / \partial \alpha$ grows at most linearly in s_0 uniformly on $\mathcal{N}$. Finally, $\partial I / \partial \gamma = \theta_{\alpha,b}$ is bounded by 1 on $\mathcal{N}$. Thus, by the Dominated Convergence Theorem, $(\alpha, b, \gamma) \mapsto m_I$ is C^1 on $\mathcal{N}$ and

$$\begin{aligned} \frac{\partial m_I}{\partial \alpha} &= \mathbb{E} \left[\frac{\partial I}{\partial \alpha} \right] = \mathbb{E} [\mu_{\tilde{s}}] + \gamma \mathbb{E} \left[\frac{\partial \theta_{\alpha,b}}{\partial \alpha} \right], \\ \frac{\partial m_I}{\partial b} &= \mathbb{E} \left[\frac{\partial I}{\partial b} \right] = 1 + \gamma \mathbb{E} \left[\frac{\partial \theta_{\alpha,b}}{\partial b} \right], \\ \frac{\partial m_I}{\partial \gamma} &= \mathbb{E} \left[\frac{\partial I}{\partial \gamma} \right] = \mathbb{E} [\theta_{\alpha,b}]. \end{aligned}$$

Hence, the expression inside the expectation in Ψ is continuously differentiable in (α, b, γ) for all s_0 .

Furthermore, we already proved that (i) $\partial I / \partial b$ is bounded by a constant on $\mathcal{N}$; (ii)

$\partial I/\partial \alpha$ grows at most linearly in s_0 uniformly on $\mathcal{N}$; and (iii) $\partial I/\partial \gamma$ is bounded by 1 on $\mathcal{N}$. It is straightforward to verify that I grows at most linearly in s_0 on $\mathcal{N}$. Further, the derivative of the expression inside the expectation in Ψ w.r.t I is given by

$$G_{\alpha,b,\gamma} = 2QI + 2H - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2}(\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I)$$

and clearly grows at most linearly on $\mathcal{N}$. Thus, the products $G_{\alpha,b,\gamma} \partial I/\partial p$ for $p \in \{\alpha, b, \gamma\}$ admit bounds that are quadratic on s_0 and not dependent on the specific point $(\alpha, b, \gamma) \in \mathcal{N}$. Thus, again because the random variables are Gaussian, we can apply the Dominated Convergence Theorem to get

$$\frac{\partial \Psi(\alpha, b, \gamma)}{\partial p} = \mathbb{E} \left[(2QI + 2H) \frac{\partial I}{\partial p} \right] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2}(\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I) \mathbb{E} \left[\frac{\partial I}{\partial p} \right] = \mathbb{E} \left[G_{\alpha,b,\gamma} \frac{\partial I}{\partial p} \right],$$

and continuity of the above derivative follows from the same argument (recall the expression inside the expectation in Ψ was proven to be continuously differentiable in (α, b, γ)).

Now, since every $(\alpha_0, b_0, \gamma_0) \in \mathcal{D}$ admits such a neighborhood $\mathcal{N}$ and the formulas above hold on each $\mathcal{N}$, $\mathcal{J}$ is C^1 on $\mathcal{D}$.

To prove the continuous extension to $\gamma = 0$, fix (α, b) with $\alpha \neq a_{ps}$ and let $\gamma \downarrow 0$.

For the $t = 1$ component, note first that $(|z| - \gamma)_+^2 \rightarrow |z|^2$ monotonically. Thus, by the Monotone Convergence Theorem, $\mathbb{E}[(|z| - \gamma)_+^2] \rightarrow \mathbb{E}[|z|^2]$. For $\partial/\partial \alpha$ and $\partial/\partial b$, note that $(|z| - \gamma)_+ \text{sign } z \rightarrow |z| \text{sign } z = z$. Further, the bounds found above do not depend on γ near 0, so we can use those and apply the Dominated Convergence Theorem to get that as $\gamma \downarrow 0$,

$$\frac{\partial W_1 \mathbb{E}[(|z_{\alpha,b}| - \gamma)_+^2]}{\partial \alpha} \rightarrow 2W_1 \mathbb{E}[z_{\alpha,b} \tilde{s}_1],$$

$$\frac{\partial W_1 \mathbb{E}[(|z_{\alpha,b}| - \gamma)_+^2]}{\partial b} \rightarrow 2W_1 \mathbb{E}[z_{\alpha,b}],$$

proving the continuity at $\gamma = 0$ of the derivatives. Moreover, for $\partial/\partial \gamma$, $(|z| - \gamma)_+ \rightarrow |z|$ and thus again the bound applies irrespective of $\gamma > 0$ or not. So, by the Dominated Convergence Theorem, the right derivative exists and it is given by

$$\frac{\partial^+ W_1 \mathbb{E}[(|z_{\alpha,b}| - \gamma)_+^2]}{\partial \gamma} = -2W_1 \mathbb{E}[|z_{\alpha,b}|].$$

For the $t = 0$ component, note that pointwise as $\gamma \downarrow 0$,

$$I(\alpha, b, \gamma; s_0) \rightarrow I_0(s_0) = \alpha \mu_{\bar{s}}(s_0) + b, \quad m_I(\alpha, b, \gamma) \rightarrow m_{I,0} := \mathbb{E}[I_0],$$

the latter using again a dominated convergence argument as $\theta_{\alpha,b}$ does not depend on γ and $|\theta_{\alpha,b}| \leq 1$. Further, because $\mu_{\bar{s}}$ and a_{ps}^0 are affine in (r_0^n, u_0) and $|\theta_{\alpha,b}| \leq 1$, for any fixed $\gamma_1 > 0$,

$$|I(\alpha, b, \gamma; s_0)| \leq |I_0(s_0)| + \gamma_1 \leq C(1 + |r_0^n| + |u_0|), \quad \gamma \in [0, \gamma_1],$$

for some constant C . Hence, I^2 is dominated by $C'(1 + |r_0^n| + |u_0|)^2$ for some constant C' , which is integrable given s_0 is Gaussian. Thus, $QI^2 + 2HI$ can be bounded with a bound independent of $\gamma \in [0, \gamma_1]$ that is quadratic in s_0 and thus integrable. Further, $m_I(\alpha, b, \gamma) = m_{I,0} + \gamma \mathbb{E}[\theta_{\alpha,b}]$ is continuous in γ and bounded on $[0, \gamma_1]$. By the Dominated Convergence Theorem,

$$\mathbb{E}[QI^2 + 2HI] \rightarrow \mathbb{E}[QI_0^2 + 2HI_0],$$

and the other term in Ψ converges since $m_I(\alpha, b, \gamma) \rightarrow m_{I,0}$. Thus $\Psi(\alpha, b, \gamma) \rightarrow \Psi(\alpha, b, 0)$.

Now, for $\gamma > 0$, we proved above that

$$\frac{\partial \Psi}{\partial \alpha} = \mathbb{E} \left[(2QI + 2H) \left(\mu_{\bar{s}} + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha} \right) \right] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I) \mathbb{E} \left[\mu_{\bar{s}} + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha} \right],$$

$$\frac{\partial \Psi}{\partial b} = \mathbb{E} \left[(2QI + 2H) \left(1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b} \right) \right] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I) \mathbb{E} \left[1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b} \right].$$

Further, note that the bounds for the derivatives of $\theta_{\alpha,b}$ found above do not depend on $\gamma > 0$, so for a small interval $\gamma \in [0, \gamma_1]$,

$$\left| \frac{\partial \theta_{\alpha,b}}{\partial b} \right| \leq c_b, \quad \left| \frac{\partial \theta_{\alpha,b}}{\partial \alpha} \right| \leq c_\alpha(1 + |r_0^n| + |u_0|),$$

for constants c_b, c_α depending on (α, b) and primitives. Hence, the integrands above are dominated by integrable expressions independent of $\gamma \in [0, \gamma_1]$. Given the pointwise limits $I \rightarrow I_0$, $m_I \rightarrow m_{I,0}$ and $\gamma \frac{\partial \theta_{\alpha,b}}{\partial p} \rightarrow 0$ for $p \in \{\alpha, b\}$, the Dominated Convergence Theorem implies

$$\frac{\partial \Psi}{\partial \alpha} \rightarrow \mathbb{E}[(2QI_0 + 2H)\mu_{\bar{s}}] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_{I,0}) \mathbb{E}[\mu_{\bar{s}}],$$

and

$$\frac{\partial \Psi}{\partial b} \rightarrow \mathbb{E}[(2QI_0 + 2H)] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_{I,0}).$$

Finally, using the results above, for $\gamma > 0$,

$$\frac{\partial \Psi}{\partial \gamma} = \mathbb{E}[(2QI + 2H)\theta_{\alpha,b}] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I) \mathbb{E}[\theta_{\alpha,b}].$$

The integrand is affine in I and bounded by $C(1 + |r_0^n| + |u_0|)$ uniformly in $\gamma \in [0, \gamma_1]$. Thus, given the pointwise limits $I \rightarrow I_0$ and $m_I \rightarrow m_{I,0}$ we can apply again the Dominated Convergence Theorem to obtain the limit of the right derivative,

$$\frac{\partial \Psi}{\partial \gamma} = \mathbb{E}[(2QI_0 + 2H)\theta_{\alpha,b}] - \frac{2\delta_{ie}}{1 + \lambda_{cb}\kappa^2} (\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_{I,0}) \mathbb{E}[\theta_{\alpha,b}],$$

thus completing the proof of the continuous extension of the derivatives of Ψ , and thus of $\mathcal{J}$ to $\gamma = 0$. $\square$

D.4 Proof of Lemma 6

Proof. First, we show that $\text{Var}(z_{\alpha,b}) > 0$ for every (α, b) . The positive measurement-error variances ensure that the conditional covariance matrix of the two signals is invertible and that $\sigma_s^2 > 0$. To economize on notation throughout the proof, let

$$X := \tilde{s}_1, \quad Y := r_1^n + \Gamma(\lambda_{cb})u_1, \quad T := \mathbb{E}[Y \mid s_0, X, \hat{s}_1].$$

If $\text{Var}(z_{\alpha,b}) = 0$, then $T = \alpha X + c$ almost surely for some constant c . Thus, conditioning on s_0 and comparing the coefficients of the independent initial shocks we have

$$\alpha \tilde{\omega}_r = 1, \quad \alpha \tilde{\omega}_u = \Gamma(\lambda_{cb}),$$

since by assumption $\sigma_r^2, \sigma_u^2 > 0$ and $\rho_r \rho_u \neq 0$. Thus, $Y = \alpha(X - \tilde{\epsilon})$, and $\mathbb{E}[Y] = \mathbb{E}[T]$ implies $c = 0$ so $T = \alpha X$. Since T conditions on X among others, conditional expectation errors are orthogonal to X . This, together with the above, implies

$$0 = \text{Cov}(Y - T, X) = -\alpha \text{Cov}(\tilde{\epsilon}, X) = -\alpha \tilde{\sigma}^2 \neq 0,$$

a contradiction. Hence, $z_{\alpha,b}$ is nondegenerate for every (α, b) .

Fix $(\alpha_0, b_0, \gamma_0) \in \mathcal{D}$ and choose a neighborhood on which $|\alpha - a_{ps}| \geq \varepsilon > 0$ and $\gamma \in [\gamma_0/2, 2\gamma_0]$ (as in Lemma 5). For the $t = 1$ part, set $h(z, \gamma) := (|z| - \gamma)_+^2$. For $z \neq \pm\gamma$, we clearly have

$$h_z = 2(|z| - \gamma)_+ \text{sgn} z, \quad h_{zz} = 2\mathbf{1}_{\{|z| > \gamma\}}, \quad h_{z\gamma} = -2\mathbf{1}_{\{|z| > \gamma\}} \text{sgn} z, \quad h_{\gamma\gamma} = 2\mathbf{1}_{\{|z| > \gamma\}}.$$

Since $z_{\alpha,b}$ is Gaussian with strictly positive variance, as established above, $\mathbb{P}(|z_{\alpha,b}| = \gamma) = 0$ and the above hold almost surely. Dominated convergence applies since the displayed derivatives are bounded by integrable functions independent of parameters on the neighborhood (see the proof of [Lemma 5](#) for details of an analogous argument). The chain rule with $\partial z/\partial \alpha = \tilde{s}_1$, $\partial z/\partial b = 1$, $\partial z/\partial \gamma = 0$ yields the six formulas and their continuity.

Regarding the $t = 0$ part, write $\Psi = \mathbb{E}[QI^2 + 2HI] - \frac{(\delta_r \bar{r}^n + \delta_c + \delta_{ie} m_I)^2}{1 + \lambda_{cb} \kappa^2}$ with $I = \alpha \mu_{\tilde{s}} + b + \gamma \theta_{\alpha,b}$. On $|\alpha - a_{ps}| \geq \varepsilon$, the map $(\alpha, b) \mapsto \theta_{\alpha,b} = 2\Phi(\nu) - 1$ is C^2 , with derivatives that are bounded by polynomials in $|r_0^n| + |u_0|$ times $\phi(\nu)$ so one can apply dominated convergence (again, see the proof of [Lemma 5](#) for details of an analogous argument). First derivatives of I are taken from [Lemma 5](#), second derivatives follow by differentiating again, using that $\partial \theta_{\alpha,b}/\partial \gamma = 0$ and that $\partial^2 \theta_{\alpha,b}/\partial p \partial q$ exists and is bounded as above.

Differentiating Ψ twice and using $\partial m_I/\partial q = \mathbb{E}[I_q]$ gives [\(22\)](#). The expressions for the derivatives of the worst-case selector imply the uniform bounds

$$|I| + |I_p| + |G_{\alpha,b,\gamma}| \leq C(1 + \|s_0\|), \quad |I_{pq}| \leq C(1 + \|s_0\|^2).$$

Thus, the Hessian integrands are bounded by integrable polynomials of degree at most three in $\|s_0\|$, and each term is continuous by dominated convergence. This implies $\Psi \in C^2$ locally. Since $(\alpha_0, b_0, \gamma_0)$ was arbitrary, $\Psi \in C^2(\mathcal{D})$.

The above together gives $\mathcal{J} \in C^2(\mathcal{D})$ and the stated formulas. Continuity of all second derivatives follows from the same dominated convergence arguments, using both the fact that $\mathbb{P}(|z| = \gamma) = 0$ and the ε -separation from $\alpha = a_{ps}$. $\square$

D.5 Proof of [Lemma 7](#)

Proof. Using the expressions for the derivatives in [Lemma 6](#), one gets that for $\gamma > 0$,

$$\begin{aligned} \frac{\partial^2 \mathcal{J}}{\partial \gamma^2} &= \xi_i \mathbb{E}[\theta_{\alpha,b}^2] - \frac{2\delta_{ie}^2}{1 + \lambda_{cb} \kappa^2} \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma) \\ &= \xi_i \sigma_\theta^2 + \left(\xi_i - \frac{2\delta_{ie}^2}{1 + \lambda_{cb} \kappa^2} \right) \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma) \\ &= \xi_i \sigma_\theta^2 + \frac{2\lambda_{cb}}{1 + \lambda_{cb} \kappa^2} (\beta \sigma \kappa)^2 \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma), \end{aligned}$$

where the last equality follows from the definitions of ξ_i and δ_{ie} and some algebra.

Note that the first two terms are independent of γ and as $\gamma \downarrow 0$, $\mathbb{P}(|z| > \gamma) \rightarrow 1$ since $z_{\alpha,b}$ has a continuous, non-degenerate Gaussian distribution. Thus, by a dominated convergence argument (see the proof of [Lemma 5](#) for details of an equivalent argument)

one can conclude. Analogously for the mixed second derivatives, one can take the limit as $\gamma \downarrow 0$ in the corresponding expressions from [Lemma 6](#) via dominated convergence, using that $\mathbf{1}_{\{|z|>\gamma\}} \rightarrow 1$ almost surely. This yields the stated expressions.

To show joint continuity, fix $(\alpha_0, b_0, 0) \in \overline{\mathcal{D}}$ and consider a sequence of announcements such that $(\alpha_n, b_n, \gamma_n) \rightarrow (\alpha_0, b_0, 0)$ and $\gamma_n \geq 0$. Let $z_n = z_{\alpha_n, b_n}$ and $z_0 = z_{\alpha_0, b_0}$. Thus, since z is a continuous function of (α, b) , we have $z_n \rightarrow z_0$ almost surely. Further, since $\gamma_n \rightarrow 0$ and, as shown in the proof of [Lemma 6](#), $\mathbb{P}(z_0 = 0) = 0$, it follows that

$$\mathbf{1}_{\{|z_n|>\gamma_n\}} \rightarrow 1, \quad \mathbf{1}_{\{|z_n|>\gamma_n\}} \text{sign}(z_n) \rightarrow \text{sign}(z_0)$$

almost surely. Now, as already shown, all six Hessian integrands of the $t = 1$ loss term are bounded by $C(1 + \tilde{s}_1^2)$, so dominated convergence gives their joint limits. These limits agree with the derivatives at the boundary itself: at $\gamma = 0$ the $t = 1$ loss term equals $W_1 \mathbb{E}[z_{\alpha, b}^2]$ and its right width derivative equals $-2W_1 \mathbb{E}[|z_{\alpha, b}|]$. Differentiation and right difference quotients using $\mathbb{P}(z_{\alpha, b} = 0) = 0$ give the same six derivative formulas. For the $t = 1$ loss itself and its first derivatives, the integrands also converge almost surely to their boundary values. On a bounded neighborhood of $(\alpha_0, b_0, 0)$, their absolute values are bounded by $C(1 + \tilde{s}_1^2 + z_0^2)$, which is clearly integrable. Hence, we can also use dominated convergence to establish their joint continuity at the boundary.

For the $t = 0$ loss, Ψ , the polynomial bounds in the preceding proof hold uniformly on bounded neighborhoods that include $\gamma = 0$ and remain separated from $\alpha = a_{ps}$, since the worst-case selector is independent of γ . Differentiation under the expectation and dominated convergence thus establish the joint continuity of Ψ and its first and second derivatives at the boundary. Combining the two loss components proves the joint continuity of $\mathcal{J}$ and its first and second derivatives as stated, interpreting vagueness derivatives from the right at the $\gamma = 0$ boundary. $\square$