Forward Guidance for Skeptical Markets

Vicente Jimenez-Gimpel[†]

MIT

March 2, 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 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: uncertainty about both demand and cost-push shocks increases optimal vagueness, but only demand-side uncertainty increases the optimal degree of data dependence. Second, strict rules are costly: fully precise guidance entails sizable welfare losses that grow with the steepness of the Phillips curve. 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

[†]MIT Department of Economics. email: jimenezv@mit.edu. I am deeply grateful to Ricardo J. Caballero and Christian K. Wolf for their continual support and guidance. I thank Tomás Caravello for many insightful suggestions. I also thank Martin Beraja, Luigi Bocola, Andrés Drenik, Joao Guerreiro, Stephen Morris, Aleksei Oskolkov, Karthik Sastry, Luminita Stevens, Jean Tirole, Rafael Veiel and Iván Werning for helpful comments.

1 Introduction

Central banks devote enormous effort to communicating properly, and for good reason: Gurkaynak et al. (2005) find that up to 90% of the variation in long-term Treasury yields following FOMC announcements is attributable to statements rather than actions. Managing expectations is widely regarded as a crucial element of successful monetary policy (see, e.g., Yellen (2013), Woodford (2005) and Bernanke (2020)). Yet the macroeconomic models at the core of policy analysis have remarkably little to say about how this communication should be designed. Standard frameworks assume the central bank can announce—and commit to—a fully state-contingent policy rule, perfectly understood by the private sector (see, e.g., Woodford (2003)), reducing the design of forward guidance to the design of optimal policy rules. In practice, forward guidance is far coarser: it conditions on a narrow subset of observables, leaves room for discretion, and faces a private sector concerned about policy mistakes. Under what circumstances should the central bank make its announcements more or less precise? And how strongly should it tie future policy to imperfect incoming data?

I study these questions in a New Keynesian model where the central bank commits to a signal-contingent *set* of interest rates rather than a single rate for each state of the world. Optimal announcements are shown to take the form of signal-contingent intervals, characterized by three design margins: *vagueness*—the width of the interval, governing how much ex-post discretion the central bank retains—*data dependence*—how strongly the interval shifts with publicly observable signals—and the *center* of the interval. The private sector is ambiguity averse regarding future policy rates: its expectations reflect the least favorable rate consistent with the announced set. This gives rise to the core trade-off governing optimal communication: anchoring expectations through precise, data-responsive announcements must be weighed against the cost of constraining the central bank's ability to respond to information that cannot be built into the announcement.

When the conditionable signal is an endogenous variable, announcements shape the very signal they are anchored to, creating a feedback loop between expectations and inflation whose direction—stabilizing or destabilizing—is governed by the strength of data dependence. With inflation as the conditionable signal, three main lessons emerge: uncertainty affects communication differently depending on its source—uncertainty about both demand and cost-push shocks increases optimal vagueness, but only demand-side uncertainty increases the optimal degree of data dependence; strict rules entail sizable welfare costs; and optimal forward guidance differs sharply from simply communicating the best forecast of the future policy rate.

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 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 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

[1]While this form of commitment is stylized, the idea that central banks are constrained by their previous announcements has both conceptual appeal (a private sector concerned about policy mistakes would have little reason to put weight on announcements if they did not somehow shape future actions) and has been repeatedly acknowledged by policymakers as a downside of forward guidance (see, e.g., Barkin (2021) and Bowman (2022), and Romer and Romer (2024) for evidence that the Fed's forward guidance delayed its response to the 2021–2022 inflation surge). Furthermore, it is the natural extension to the environment considered here of the standard commitment assumption usually made in the literature of optimal monetary policy.

[2]Yellen (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 imprecise: "...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.*"

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 models featuring rational expectations, 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 vector, 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, in this benchmark, the familiar distinction between "Odyssean" forward guidance (policy commitments) and "Delphic" guidance (information disclosure) (Campbell et al., 2012) collapses, since both coincide by construction—a feature that is difficult to reconcile with the analytical relevance this distinction has proven to have for understanding central bank communication.

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 miscentering of the band, making equilibrium unique and self-correcting. Perhaps more surprisingly, the same stabilizing logic applies when data dependence is sufficiently negative, as the feedback through rate expectations is then strong enough to reverse the direct effect on inflation. For moderately negative data dependence, however, the feedback is destabilizing: extreme beliefs become self-confirming, giving rise to multiple equilibria. 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 no longer just widens the set of feasible rates but also shifts the band's center endogenously, as equilibrium inflation depends on the band's width. 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. Intuitively, cost-push shocks drive a wedge between inflation and output stabilization, making inflation a noisier guide to the full-information optimal policy rate. Conditioning policy tightly on inflation in such environments would induce overreaction to supply disturbances. Thus, the central bank optimally weakens the extent to which its guidance is anchored to inflation. This asymmetry is a robust feature of the numerical analysis, holding across a wide range of calibrations.

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 welfare losses relative to optimally vague guidance (roughly 6% under the baseline calibration). These losses increase with the steepness of the Phillips curve and the weight in the central bank's loss function on inflation stabilization. Both patterns are robust to sensible changes in the parameters. 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 forward guidance is not Delphic. Optimal data dependence generically differs from the "best-forecast" benchmark obtained by projecting the ex-post optimal rate onto the conditionable signal, since optimal guidance internalizes its stabilizing effects on expectations. Quantitatively, optimal data dependence substantially exceeds this benchmark—widening with demand uncertainty and narrowing with cost-push uncertainty—a pattern that is robust across sensible parameterizations. 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 draws standard New Keynesian parameters from Rotemberg and Woodford (1997) and Woodford (2003), with the volatility of natural rate shocks taken from the Federal Reserve Bank of New York's DSGE model to reflect current estimates of demand-side uncertainty. Robustness is assessed by varying each key parameter—the volatilities of demand and cost-push shocks, the slope of the Phillips curve, and the weight on inflation stabilization—individually over a wide range around this baseline.

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)). In cheap-talk models, the focus is on equilibrium information transmission, while in information-design models, it is on the sender-optimal disclosure rule. The private information may pertain to the state of the economy, the central bank's preferences, or both. Bassetto (2019), for instance, builds on a Barro and Gordon (1983) framework to study when forward guidance can enlarge the set of sustainable equilibria. 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, because 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.

The current paper takes a different angle: it studies the design of optimal forward guidance when fully state-contingent rules are infeasible and the private sector is worried about the possibility of policy misjudgment. Consequently, it is related 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 literature, 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 ex-ante encoded in a 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 partial 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, rather than as the joint design of policy and communication.

Finally, by considering ambiguity averse agents whose beliefs are not fully subjective but partially shaped by features of the environment (in this case, the central bank's announcement), the analysis builds on the decision-theoretic framework of Gajdos et al. (2008).

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, comparative statics, and robustness analysis 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. Additional results are presented in Appendix C (second-order conditions and partial comparative statics of the optimal announcement policy). 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 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 $(r_t^n, u_t) = (\overline{r}^n, 0)$ for $t \geq 2$, where $\overline{r}^n$ 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 Curve[3]

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

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's preferences are represented by a standard "dual mandate" over output gap and inflation volatility. As it is well known, this functional form with a specific weight on inflation can be obtained as a second-order approximation of the welfare of the representative household in a New Keynesian economy (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

$$\begin{aligned} \tilde{V}_0 = \min_{y_0, \pi_0, i_0, \mathcal{A}_1} & \mathbb{E}_0[y_0^2 + \lambda_{cb}\pi_0^2 + \beta V(\mathcal{A}_1)] \\ \text{s.t.} \quad & y_0 = \tilde{\mathbb{E}}_0[y_1|\mathcal{A}_1] - \sigma(i_0 - \tilde{\mathbb{E}}_0[\pi_1|\mathcal{A}_1] - r_0^n) \\ & \pi_0 = \kappa y_0 + \beta\tilde{\mathbb{E}}_0[\pi_1|\mathcal{A}_1] + u_0, \end{aligned}$$

where $\lambda_{cb} > 0$ captures the relative importance of inflation stability relative to output sta-

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

[4]In 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.

[5]The 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.

bility for the central bank and

$$
\begin{aligned}
V(\mathcal{A}_t) = \min_{y_t,\pi_t,i_t,\mathcal{A}_{t+1}} \mathbb{E}_0 \left[y_t^2 + \lambda_{cb}\pi_t^2 + \beta V_{t+1}(\mathcal{A}_{t+1}) \right] & \\
\text{s.t.} \quad & y_t = \tilde{\mathbb{E}}_t[y_{t+1}|\mathcal{A}_{t+1}] - \sigma(i_t - \tilde{\mathbb{E}}_t[\pi_{t+1}|\mathcal{A}_{t+1}] - r_t^n) \\
& \pi_t = \kappa y_t + \beta\tilde{\mathbb{E}}_t[\pi_{t+1}|\mathcal{A}_{t+1}] + u_t \\
& i_t \in \mathcal{A}_t
\end{aligned}
$$

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, we take $\tilde{s}_1$ 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 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).

We restrict feasible announcements 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*

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

The above is a translation-invariance property (shifting the signal by δ translates the entire announced set by $\alpha\delta$), that essentially amounts to imposing the conditioning on $\tilde{s}_t$ is linear. 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 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]

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

[7] It's straightforward to verify that if an announcement satisfies equivariance for some α, there is no $\alpha' \neq \alpha$

2.3 Beliefs

This subsection describes the beliefs of the central bank and the private sector. The key non-standard ingredients 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

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

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),

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

where (ϵ_r, ϵ_u) are independent random variables with 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 again $(\tilde{\epsilon}, \hat{\epsilon})$ are independent random variables with mean zero and variances $\tilde{\sigma}^2$ and $\hat{\sigma}^2$ and $\tilde{\omega}_r, \tilde{\omega}_u, \hat{\omega}_r, \hat{\omega}_u > 0$ are exogenous weights on demand and supply shocks (all random variables are jointly normal). As mentioned above, the public signal $\tilde{s}_1$ 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.

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.

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 it 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 is *ambiguity averse* as in the multiple priors representation of Gilboa and Schmeidler (1989) with one twist: instead of the set of priors being part of the preferences, it is shaped by central bank announcements. Thus, the private sector preferences regarding policy rates fall conceptually within the scope of the axiomatization in Gajdos et al. (2008), where there is both imprecise "objective information" (a set of possible probabilities, the ones supported on the central bank announcement) and an attitude towards it. In our case, we study preferences that feature the simplest possible attitude towards it: the set of possible beliefs (the ambiguity set) is precisely the objectively given set, i.e., the announcement (in their notation $\varphi = \text{id}$).

The private sector's worst-case rate is taken with respect to the second-order approximation of the utility function, thus mirroring the loss of the central bank[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 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: market (joint) beliefs are aligned with what it knows about the distribution of the exogenous variables. The second condition reflects that the central bank observes contemporaneous shocks (r_0^n, u_0) and the conditionable signal $\tilde{s}_1$—but not the underlying future shocks (r_1^n, u_1)—when choosing i_1, so the private sector's ambiguity concerns only the policy choice conditional on that information. Finally, the third condition follows from the market understanding that the central bank is bound by its announcements. Denoting this admissible class by $\mathcal{P}(s_0)$, the private sector worst-case

[8]As shown by Benigno and Woodford (2005), the Lagrangian corresponding to the household's utility function with the equilibrium conditions as constraints in the New Keynesian model is locally concave and thus it is minimized at an extreme point provided a parametric condition holds (and this is always the case if the steady state is efficient as in our setting). Thus, the global worst-case scenario always lies on the boundary of the announced set. Consequently, and because we are working with log-linearized private sector behavior, taking the worst-case with respect to a second-order approximation of welfare is a natural assumption.

beliefs solve[9]

$$\sup_{p \in \mathcal{P}(s_0)} \mathbb{E}_p \left[y_1^2 + \lambda_{ps} \pi_1^2 \right] \tag{1}$$

$$\begin{aligned} \text{s.t.} \quad & y_1 = -\sigma(i_1 - r_1^n) \\ & \pi_1 = \kappa y_1 + u_1 \\ & \tilde{s}_1 = \tilde{\omega}_r r_1^n + \tilde{\omega}_u u_1 + \tilde{\epsilon}_1, \end{aligned}$$

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 $t \geq 2$ as they have perfect foresight about the future state of the economy and because cost-push shocks are set to its mean, zero, the central bank can achieve full stabilization.

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 minimizes the private sector's utility 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 maximize its (i.e., the private sector's) utility. Despite this standard assumption, 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, e.g., Baker et al. (2016); Bauer et al. (2022)).[11]. Third, it abstracts from the precise source of concerns about future policy choices,

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

[10] The 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).

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

while remaining consistent with familiar interpretations (for instance, as an approximation to expected-utility behavior under uncertainty about central bank's preferences λ_{cb} and high risk aversion).

The importance of ambiguity aversion lies in its consequences for the way announcements map into the private sector's behavior. A necessary condition for communication policy to matter is that private-sector beliefs about future policy respond to announcements. If the private sector correctly anticipates the ex-post optimal policy choice of the central bank, or, more generally, holds beliefs that are invariant to announcements, the latter play no role and the optimal communication policy problem becomes moot.

In addition, one of the paper's two core mechanisms—the commitment-flexibility trade-off—requires that private-sector expectations after an announcement 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, endpoint sensitivity requires moving beyond a linearized setting: in models with linearized private-sector behavior only conditional means matter, so endpoints 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, we restrict the space of feasible central bank announcements 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 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 place different probability weight 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. In contrast, imposing an affine center forces the announcement to respond with the same marginal sensitivity at all signal realizations, precluding such state-dependent reweighting.

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 need not be 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 optimal central bank's policy. I proceed in three steps. First, I solve for the optimal rate policy for an arbitrary announcement: at $t = 1$ the bank sets the closest feasible rate to its best prediction of the full-information target, and at $t = 0$ it chooses i_0 to stabilize the predictable components of shocks and announcement-driven expectations fluctuations 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, it 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 in the more transparent $\bar{r}^n = 0$ case. Finally, I provide a characterization of optimal announcements in the general case.

3.1 Optimal Rate Policy

By assumption, there is no uncertainty from period $t = 2$ onward, so the central bank at $t \geq 2$ achieves full stabilization by letting $i_t = \bar{r}^n$, $\mathcal{A}_t = \{\bar{r}^n\}$ for $t \geq 2$ and this 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*

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

where $i_1^{\mathrm{FI}}(\lambda) := r_1^n + \Gamma(\lambda) u_1$ 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].

[12]That 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

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

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

Proposition 1. *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*

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

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

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

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 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]

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

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.

[13]This 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.

that include i_1^e, with a strictly positive coefficient on $(i_1^e)^2$. Widening the band (raising γ) increases the size of i_1^e and thus generally pushes $t = 0$ losses upward, so the communication trade-off would persist in a different form. 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 set the stage regarding the optimal rate policy, the following proposition describes the form of optimal announcements.

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*

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

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 $\mathcal{A} = [L, \underline{a}] \cup [\overline{a}, U]$: 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: ex post, after the state is realized, the central bank would like to implement a rate i_1^*, but it is constrained to pick $i_1 \in \mathcal{A}$. With a non-interval $\mathcal{A}$, states in which $i_1^*(\omega) \in (\underline{a}, \overline{a})$ force the bank to choose one of the boundary rates $\underline{a}$ or $\overline{a}$, creating an avoidable distortion. Figure 1b highlights this "hole": whenever the preferred rate falls in the shaded region, the non-interval announcement prevents the bank from implementing it. Announcing the convex hull $[L, U]$ instead strictly expands the set of feasible ex-post actions: relative to $\mathcal{A}$, it adds the entire hole $(\underline{a}, \overline{a})$ while leaving the extreme points L and U unchanged. This extra flexibility can only (weakly) improve welfare: for each realized state, the central bank can always replicate its choice under $\mathcal{A}$, and in states in which $i_1^* \in (\underline{a}, \overline{a})$ it can choose a rate closer to i_1^* than the boundary points $\underline{a}$ or $\overline{a}$. 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]

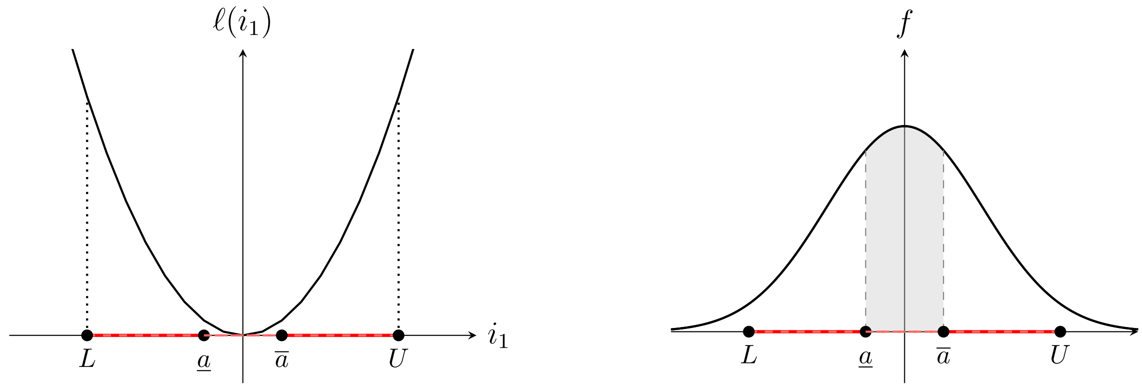

(a) The private sector's expected losses as a function of next-period's policy rate.

(b) The central bank's ex-post preferred rate distribution.

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.

Given the above, an (optimal) announcement is fully pinned down by three scalars: a slope α, a centering variable 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 its main parameters: the slope α, the center b and the width γ. In panel (2a), α is held fixed while γ varies: larger γ widens the admissible band at each signal realization (more vagueness) keeping the

[14] Figure 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.

slope fixed. In panel (2b), γ is held fixed while α varies: larger $|\alpha|$ rotates/tilts the band more strongly with the signal (greater data dependence) keeping the width fixed.

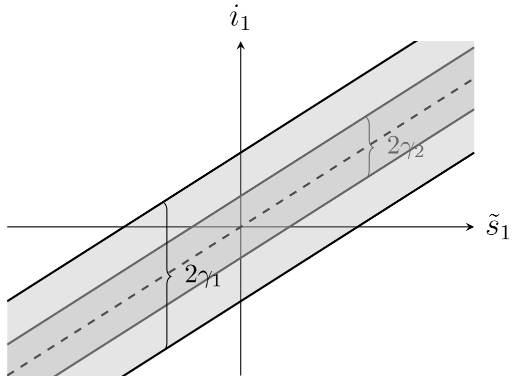

(a) Vagueness: $\gamma_1 > \gamma_2$

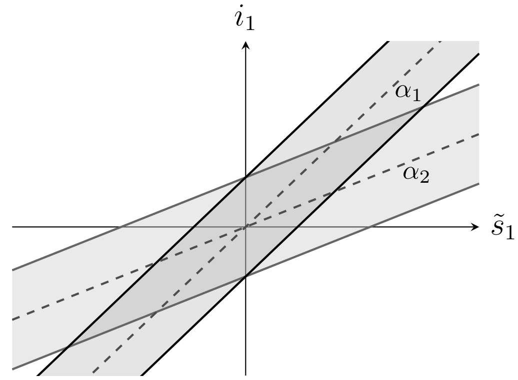

(b) Data dependence: $\alpha_1 > \alpha_2$

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

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

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

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, 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. If the band is exactly centered on the private sector's conditional expectation of its preferred rate, the two endpoints are equally likely to be the worst one and so $\theta_{\alpha,b} = 0$ and the expected worst-case rate equals the interval's midpoint. If, on the other hand, the band is "miscentered" relative to the private sector's information and preferences, one endpoint is more likely to be worse and expectations tilt toward it. Moreover, larger $|\alpha - a_{ps}|$ and a noisier conditionable signal raise the market's uncertainty regarding which endpoint is worse, decreasing its predictability and shrinking $\theta_{\alpha,b}(s_0)$ toward 0. As the denominator shrinks, the private sector's uncertainty decreases. In the limit as it goes to zero, the private sector perfectly knows which endpoint is worse, and $\theta_{\alpha,b}(s_0)$ becomes a step function that jumps from -1 to 1 at the point where the band is centered on the private sector indifference point. By contrast, as $\sigma_{\tilde{s}} \to \infty$ the worst endpoint becomes fully unpredictable, $\theta_{\alpha,b}(s_0) \to 0$, and the expected rate collapses to the band's midpoint.

Finally, note that the central bank can at most set the band's center to be *on average* equal to the private sector's expected optimal rate, as it does not observe the contemporaneous shocks r_0^n and u_0 when making its announcement. Further, this is generically not optimal as it is costly in terms of period-1 stabilization (it centers the feasible set at the "wrong" place).

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 creates a fundamental tension between $t = 0$ and $t = 1$ objectives: *data dependence* α governs how expectations and the feasible set 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, with the relative weights depending on signal quality, shock volatility, and the alignment between central bank and private sector preferences.

3.3 Optimal Announcements: The $\overline{r}^n = 0$ Benchmark

To build intuition, I first analyze the special case $\overline{r}^n = 0$. This case delivers a tractable benchmark in which the core (α, γ) trade-offs are easiest to see, while the role of the centering term b can be characterized cleanly. In particular, as the following proposition shows, symmetry implies that $b^* = 0$ is optimal, so there is no mean misalignment between the central bank's and the private sector's preferred rates in this benchmark.

Proposition 3. *If $\overline{r}^n = 0$, the optimal central bank policy features $i_0^* = b^* = 0$. Moreover, optimal data dependence and vagueness is pinned down by*[15]

$$\xi_i \operatorname{Var}(\mu_{\tilde{s}})(\alpha - \tilde{\alpha}_0) + \xi_i \gamma \left(\tilde{\omega}_r \rho_r \sigma_{r\theta} + \tilde{\omega}_u \rho_u \sigma_{u\theta}\right) + \gamma h(\alpha - a_{ps}; \gamma) + 4W_1(\alpha - \tilde{\alpha}_1) \operatorname{Var}(\tilde{s}_1) \left(1 - \Phi(\gamma/\sigma_z)\right) = 0,$$

$$(\xi_r + \alpha \xi_i \tilde{\omega}_r \rho_r)\sigma_{r\theta} + (\xi_u + \alpha \xi_i \tilde{\omega}_u \rho_u)\sigma_{u\theta} + \xi_i \sigma_\theta^2 \gamma + 4W_1 \left[\gamma \left(1 - \Phi(\gamma/\sigma_z)\right) - \sigma_z \phi(\gamma/\sigma_z)\right] \geq 0,$$

[15] See Appendix C.1 for easily verifiable conditions to ensure the second-order conditions hold around a candidate solution.

with equality unless $\gamma = 0$. Here, $\sigma_{r\theta}$ denotes the covariance between r_0^n and $\theta_{\alpha,0}$, $\sigma_{u\theta}$ the covariance between u_0 and $\theta_{\alpha,0}$, σ_z the s.d. of $z_{\alpha,0}$ and σ_θ the s.d. of $\theta_{\alpha,0}$. Further, ξ_r, ξ_u, ξ_i are constants,

$$h(\alpha;\gamma) := \mathbb{E}\left[((\xi_r + \alpha\xi_i\tilde{\omega}_r\rho_r)r_0^n + (\xi_u + \alpha\xi_i\tilde{\omega}_u\rho_u)u_0 + \xi_i\gamma\theta_{\alpha,0})\frac{\partial\theta_{\alpha,0}}{\partial\alpha}\right],$$

$$\tilde{\alpha}_0 := -\frac{\xi_r\tilde{\omega}_r\rho_r\sigma_r^2 + \xi_u\tilde{\omega}_u\rho_u\sigma_u^2}{\xi_i \operatorname{Var}(\mu_{\tilde{s}})} > 0, \tag{3}$$

and

$$\tilde{\alpha}_1 = a_{cb} + \frac{\operatorname{Cov}(\tilde{s}_1, a_{cb}^0(s_0, \hat{s}_1))}{\operatorname{Var}(\tilde{s}_1)}. \tag{4}$$

Centering and mean misalignment. The fact that it is optimal to set $b = 0$ is a consequence of the symmetry of the environment (losses being squared deviations and the primitive random variables being zero-mean Gaussian). In particular, because all the components of both the private sector's preferred policy rate $i_1^{\mathrm{FI}}(\lambda_{ps})$ and the central bank's preferred conditional rate $\mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}, \hat{s}]$ are mean zero, they are both mean zero and thus choosing an interval centered in zero on average achieves both period's objectives simultaneously (i.e., there is no *mean misalignment* between the central bank's preferred midpoint and the average private sector's preferred midpoint from the central bank's $t = 0$ perspective). Note, however, that while the *expected* band center is zero, the *volatility* of the band's location via $\operatorname{Var}(\alpha\tilde{s}_1)$ remains relevant.

Furthermore, the optimality of $i_0 = 0$ follows from the fact that with zero-mean shocks, there is no predictable component of y_0 and π_0 from the central bank's $t = 0$ information set (neither directly nor through the private sector's policy expectations, as $\mathbb{E}\, i_1^e = 0$). This, of course, depends directly on the optimality of $b^* = 0$. For $b \neq 0$, it is not optimal to set $i_0 = 0$, even if $\bar{r}^n = 0$, as there is a predictable source of inflation and output instability via rate expectations (both directly and via the worst-case selector, which does not have mean zero if $b \neq 0$).

Data dependence: balancing the market's and central bank's targets. 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. The first corresponds to the term

$$\xi_i \operatorname{Var}(\mu_{\tilde{s}})(\alpha - \tilde{\alpha}_0),$$

which shows that α is pulled toward the covariance—implied target $\tilde{\alpha}_0$. This target minimizes the variance of the "linear" part of the expected rate, since it offsets the direct impact of r_0^n and u_0 on inflation and output through their covariances with $\mu_{\tilde{s}}$.

The second force corresponds to the terms

$$\gamma \left(\xi_i \left(\tilde{\omega}_r \rho_r \sigma_{r\theta} + \tilde{\omega}_u \rho_u \sigma_{u\theta} \right) + h(\alpha - a_{ps}; \gamma) \right).$$

Both reflect the volatility introduced by the private sector's worst-case selector. The closer α is to a_{ps}, the weaker this volatility channel becomes. The covariance terms $\sigma_{r\theta}, \sigma_{u\theta}$ shrink when the worst-case selector is less correlated with the conditional mean, and the $h(\cdot)$ term reflects that the marginal effect of α on the selector is minimized at $\alpha = a_{ps}$ (because selection is then pinned by the intercept a_{ps}^0 rather than flipping across states). Far from a_{ps}, in contrast, many states (r_0^n, u_0) lie on the knife-edge where the worst-case flips with small slope changes, amplifying volatility.

In sum, $t = 0$ incentives balance being close to $\tilde{\alpha}_0$ to mitigate volatility from the signal-driven band location with being close to a_{ps} to mitigate volatility driven by the worst-case selection

Regarding $t = 1$ incentives, note first that because $\mathbb{E}[\tilde{s}_1] = 0$, the mean term in the $t = 1$ component of the α-FOC vanishes: as the signal is mean-zero, α is of no use to correct average misalignments (further, there are no average misalignment because as argued before $b^* = 0$ perfectly aligns $t = 1$ average bands under both central bank and private sector preferences and beliefs). Thus, the $t = 1$ component of the α-FOC only features the covariance term, which pulls α toward $\tilde{\alpha}_1$. The former term in $\tilde{\alpha}_1$, a_{cb}, corresponds to the ex-post preferred weight of the central bank on the conditionable signal. The additional adjustment comes from the fact that at $t = 1$ the central bank wants to use its other information to set the policy rate too: the realizations of lagged shocks and the non-conditionable signal. Because ex-ante is not allowed to condition on them, the central bank optimally projects this "missing information" $(s_0, \hat{s}_1)$ onto $\tilde{s}_1$. This is precisely what the covariance term represents. In the natural case, where both signals load positively on demand and supply shocks and the shocks are positively autocorrelated over time, the covariance term is positive so $\tilde{\alpha}_1 > a_{cb}$, i.e., the central bank announces a larger weight on the conditionable signal than what it would like to do ex-post. The remaining components of the $t = 1$ term in the α-FOC represent the overall importance of the $t = 1$ target versus the $t = 0$ "targets". As expected, this is increasing in σ_z, the volatility of the misalignment of the mean band and the ex-post central bank preferred rate.

Vagueness: calming markets vs. retaining flexibility. Note that the $t = 0$ component is affine in γ, and its coefficient is the (scaled) volatility of the worst-case rate selector. This clearly shows that the higher the central bank's uncertainty regarding the selection of the worst-case rate, the more costly is vagueness (the weight ξ_i is a sum of the squared coefficients of rate expectations on inflation and output).

The $t = 1$ component, on the other hand, highlights the benefit of vagueness, that is, decreasing the expected losses derived from not being able to properly set the policy rate to stabilize the predictable component of $t = 1$ shocks. In particular, it is not difficult to see that the term multiplying $4W_1$ equals the expected $t = 1$ deviations from the unconstrained optimal rate conditional on the lower bound of the announced interval binding times the probability of this event. Because the loss is quadratic and all random variables have symmetric distributions, the term associated with the upper bound binding is analogous. See Figure 3 for a graphical illustration of the role vagueness plays.

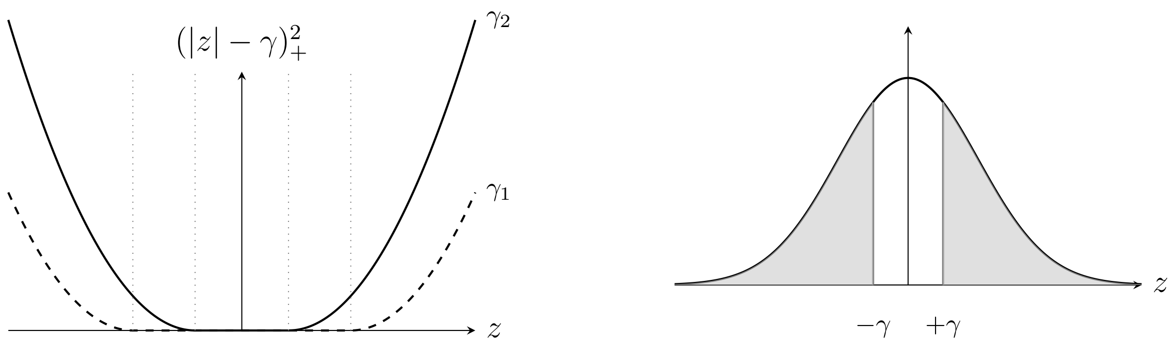

(a) $t = 1$ losses.

(b) Binding region: $p(z, \gamma) = P(|z| > \gamma)$.

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 preferred rate minus the midpoint of the announcement) 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.

3.4 Optimal Announcements: Beyond The Symmetric Benchmark

Recall from the benchmark $\bar{r}^n = 0$ case that optimal announcements balance two key trade-offs. First, *data dependence* α governs how the announcement responds to signals, trading off $t = 0$ stabilization (via how the market's policy expectations are determined) against $t = 1$ alignment (ensuring the announced feasible set is well-positioned relative to the ex-post optimal rate). Second, *vagueness* γ controls the width of the announced interval. More precise announcements reduce $t = 0$ volatility by calming the markets and thus making their rate expectations less uncertain, but vaguer ones preserve $t = 1$ flexibility to adjust policy in response to shocks. The worst-case selector $\theta_{\alpha,b}$ plays a central role: its volatility and covariances with shocks determine how expected policy rates respond to current shocks, directly affecting $t = 0$ output and inflation.

While the $\bar{r}^n = 0$ case provides clear intuition, the general case introduces additional subtleties. In particular, when $\bar{r}^n \neq 0$, *mean misalignment* considerations create additional trade-offs in the choice of b and interactions between all three announcement dimensions become more complex. Unlike the special case where the expected band center is zero (eliminating average misalignment), here the central bank's and private sector's preferred average policy rates differ, making the choice of b non-trivial. The following proposition characterizes the optimal announcement in the general case.

Proposition 4. *Let $W_1 := \beta\sigma^2(1 + \lambda_{cb}\kappa^2)$. The central bank optimal announcement minimizes*

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

where $\Psi(\alpha, b, \gamma) := \Psi(i_1^e(\alpha, b, \gamma; s_0))$, with $\Psi(\cdot)$ as in Proposition 1 and $i_1^e(\alpha, b, \gamma; s_0)$ as in Lemma 2,

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

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

There exists a solution to the central bank problem. Further, the map $(\alpha, b, \gamma) \mapsto \mathcal{J}(\alpha, b, \gamma)$ is continuously differentiable on $\mathcal{D} := \{(\alpha, b, \gamma) \in \mathbb{R}^3 : \gamma > 0, \alpha \neq a_{ps}\}$.

Moreover, all local minimizers satisfy the following first–order conditions:[16]

$$\mathbb{E}\left[G_{\alpha,b,\gamma}\left(\mu_{\tilde{s}} + \gamma\frac{\partial\theta_{\alpha,b}}{\partial\alpha}\right)\right] + 2W_1\left(\mathbb{E}[\mu_{\tilde{s}}]\,\mathbb{E}\left[(|z_{\alpha,b}| - \gamma)_+ \operatorname{sign}(z_{\alpha,b})\right] + \operatorname{Cov}(\tilde{s}_1, z_{\alpha,b})p(z_{\alpha,b}, \gamma)\right) = 0,$$

$$\mathbb{E}\left[G_{\alpha,b,\gamma}\left(1 + \gamma\frac{\partial\theta_{\alpha,b}}{\partial b}\right)\right] + 2W_1\,\mathbb{E}\left[(|z_{\alpha,b}| - \gamma)_+ \operatorname{sign}(z_{\alpha,b})\right] = 0,$$

$$\mathbb{E}\left[G_{\alpha,b,\gamma}\theta_{\alpha,b}\right] - 2W_1\,\mathbb{E}\left[(|z_{\alpha,b}| - \gamma)_+\right] \geq 0,$$

where $G_{\alpha,b,\gamma}(s_0)$ is defined in Lemma 5, $p(z_{\alpha,b}, \gamma)$ denotes the probability that the announcement is binding, and the latter inequality holds without slack unless $\gamma = 0$. Further, all three derivatives extend continuously to points where $\gamma = 0$, interpreting $\partial\mathcal{J}/\partial\gamma$ as a right derivative there.

[16] See Appendix C.1 for easily verifiable conditions to ensure the second-order conditions hold around a candidate solution.

Disagreement over band center. Relative to the benchmark in Section 3.3, the key change when $\overline{r}^n \neq 0$ is that the central bank's and the market's preferred band centers no longer coincide in expectation from the $t = 0$ perspective. In the benchmark, symmetry implies that setting $b = 0$ simultaneously aligns the expected midpoint with both sides' preferred centers, so centering is not a contested margin. Once $\overline{r}^n \neq 0$, this coincidence breaks and b becomes a genuine policy choice: centering closer to the market reduces worst-case selection volatility at $t = 0$, while centering closer to the central bank reduces expected misalignment losses at $t = 1$ by lowering the frequency and severity with which the announcement binds (see Figure 3b). Furthermore, this midpoint margin reshapes the optimal degree of data dependence and vagueness as discussed below.

Data dependence. When $\overline{r}^n \neq 0$, the signal and the market's expectations of it no longer have zero mean from the central bank's perspective, which introduces additional destabilizing forces in periods $t = 0$ and $t = 1$ (in the former even after accounting for an optimally chosen i_0, see footnote 12). Hence, the extent of data dependence shapes both the ex-post mean location of the band and the market's average policy expectation. Consequently, in addition to the volatility considerations highlighted in the benchmark $\overline{r}^n = 0$, optimally choosing α also involves balancing the average disagreement between the central bank and the market regarding their preferred set of future policy rates. While this consideration also shapes the choice of b, as discussed above, these are not fully equivalent because of the random nature of $\tilde{s}_1$ and its effect on the market's worst-case selection.

Vagueness. When $\overline{r}^n \neq 0$, the choice of optimal vagueness interacts with the band's center consideration. Concretely, because of the disagreement, in this case $z_{\alpha,b}$ will generically not be centered at zero at the optimal (α, b). Thus, vagueness becomes more valuable because this increases the probability of the announcement being binding at $t = 1$, which in turn increases the marginal value of a wider band. However, in this case it is also the case that generically at the optimal choice of (α, b) the band will be "miscentered" from the private sector's perspective on average. This may increase the "market's fear" cost of vagueness, because it makes the market's worst-case evaluation "one-sided" (i.e., it increases the likelihood that one end of the announced interval is worse than the other), so making a vaguer announcement raises the mean worst-case expectation by more (which, again, is optimally only partly offset by i_0).

Appendix C.2 further explores the mechanics of the model by studying how each announcement dimension responds individually to shifts in signal quality, holding the other fixed. While these partial effects do not pin down the joint response of optimal announcements, they help clarify the forces shaping each dimension and thus aid interpretation of the numerical exercises in Section 5 and thus the overall lessons the framework offers.

In the next section I modify the model to let the central bank condition on an endogenous signal: inflation. As will become apparent shortly, this variant 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 5 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 jointly characterize the optimal rate and communication policy (Proposition 6) and contrast it with the exogenous-signals baseline. Three differences stand out. First, fully 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 shifts 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 rather than on an exogenous signal; 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.

On the other hand, in the endogenous-signals model, removing shock persistence has two effects. First, the response of market expectations to $t = 0$ policy and announcements becomes fully predictable (as in the exogenous-signals model). Second, π_0 becomes uninformative about the relevant $t = 1$ state. The second effect implies that from the perspective of $t = 1$ losses alone, the central bank should 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

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

[18]The 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.

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, we restrict the analysis 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 their worst-case rate expectation, while that expectation is in turn anchored by the announcement through $\alpha\pi_0$. Expectations and inflation are therefore jointly determined, creating a feedback loop that does not arise with exogenous signals. As standard in macroeconomics, I take the representative agent to stand for an aggregation of non-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 5. *For any policy rate i_0 and announcement (α, b, γ) with $\alpha \neq \alpha^\star := \frac{1}{A_\pi^{i^e}}$, let $\Theta_{\alpha,b,\gamma}(\pi, s_0) \subseteq [-1, 1]$ 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*

$$\theta_0 \in \Theta_{\alpha,b,\gamma}(\pi_0, s_0) \quad \textit{and} \quad \pi_0 = \frac{A_\pi^r r_0^n + A_\pi^u u_0 - \sigma\kappa i_0 + A_\pi^{i^e}(b + \gamma\theta_0)}{1 - \alpha A_\pi^{i^e}},$$

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

$$\mathcal{E}_{\alpha,b,\gamma}(s_0) := \left\{\theta \in [-1,1] : \exists \pi \ s.t.\ \theta \in \Theta_{\alpha,b,\gamma}(\pi, s_0) \ and \ \pi = \frac{A^r_\pi r^n_0 + A^u_\pi u_0 - \sigma\kappa i_0 + A^{i^e}_\pi(b+\gamma\theta)}{1-\alpha A^{i^e}_\pi}\right\}.$$

Then (π_0, θ_0) is an equilibrium if and only if $\theta_0 \in \mathcal{E}_{\alpha,b,\gamma}(s_0)$ and π_0 is given by the above display.

Let

$$T(s_0;\alpha,b) := -\frac{1}{\alpha A^{i^e}_\pi}\Big(\Big(\alpha A^r_\pi - (1-\alpha A^{i^e}_\pi)\rho_r\Big) r^n_0 + \Big(\alpha A^u_\pi - (1-\alpha A^{i^e}_\pi)\Gamma(\lambda_{ps})\rho_u\Big) u_0 \\ + b - (\alpha\sigma\kappa) i_0 - (1-\alpha A^{i^e}_\pi)(1-\rho_r)\bar{r}^n\Big).$$

The equilibrium selection set is characterized as follows.

- *If $\alpha > 0$ or $\alpha < \alpha^\star$, then the equilibrium selection set is a singleton $\mathcal{E}_{\alpha,b,\gamma}(s_0) = \{\vartheta_{\alpha,b,\gamma}(s_0)\}$ with*

$$\vartheta_{\alpha,b,\gamma}(s_0) = \begin{cases} -1 & \text{if } T(s_0;\alpha,b) < -\gamma \\ \frac{T(s_0;\alpha,b)}{\gamma} & \text{if } -\gamma \le T(s_0;\alpha,b) \le \gamma \\ +1 & \text{if } T(s_0;\alpha,b) > \gamma. \end{cases}$$

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

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

- *If $\alpha = 0$, then*

$$\mathcal{E}_{0,b,\gamma}(s_0) = \begin{cases} \{-1\} & \text{if } b < E[i^{\mathrm{FI}}_1(\lambda_{ps})|s_0] \\ [-1,+1] & \text{if } b = E[i^{\mathrm{FI}}_1(\lambda_{ps})|s_0] \\ \{+1\} & \text{if } b > E[i^{\mathrm{FI}}_1(\lambda_{ps})|s_0]. \end{cases}$$

The object $T(s_0;\alpha,b)$ 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 $\alpha A^{i^e}_\pi$) and how this response is amplified or dampened by the feedback loop (through $1-\alpha A^{i^e}_\pi$). 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 $T \ll 0$ 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 baseline model, the market's expected worst-case selection $\theta_{\alpha,b}$ is not a smooth function of its perceived misalignment. 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 worst-case selection is locally insensitive to changes in vagueness. This follows from the fact that marginal change in the announcement's vagueness marginally increases the share of the private sector expecting the upper bound to be the worse and the share expecting the lower bound to be the worse by the same amount, netting out. Nevertheless, all three announcement's dimensions affect the probability that the equilibrium worst-case selection falls in each region.

Stabilizing vs. Destabilizing Announcements. When $\alpha > 0$ or $\alpha < \alpha^\star$, 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 miscentering 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 $\alpha < \alpha^\star$, the logic is similar: if expectations lean toward the upper endpoint of the band, inflation increases, and thus, because $\alpha^\star < 0$, the miscentering 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 $\alpha \in (\alpha^\star, 0)$, 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,

[19]Note 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 the same-period of endogenous variables (chiefly among them, inflation).

so the equilibrium reduces to a simple comparison of the band's center b with the private sector's preferred rate.

Note, finally, that in all three cases there are some shock realizations where there are "intermediate" equilibria. In these, the private sector is indifferent between the two corners and thus a fraction of it chooses one corner and the rest the other. Importantly, there is a unique split between the two corners that can be sustained as an equilibrium, due to the feedback into the evaluation of the worst-case (except in the $\alpha = 0$ case, as then there is no such feedback). In these cases, T/γ represents the fraction of agents choosing the upper endpoint minus the fraction of agents choosing the lower endpoint.

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.

Because the $t = 0$ equilibrium worst-case selection need not be unique (Proposition 5), the optimal policy problem studied below requires an equilibrium selection rule. We therefore fix a measurable selection throughout. Two remarks simplify the role of this convention. First, under the Gaussian specification for s_0 and the affine structure of $\mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$, the knife-edge event $b = \mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$ in the $\alpha = 0$ case occurs with probability zero for any fixed b, so the tie-break at equality is immaterial for expected losses. Second, numerical exercises indicate that the qualitative and quantitative results reported below are insensitive to alternative selections in the multiplicity region: varying the selection rule for $\alpha \in (\alpha^\star, 0)$ leads to essentially identical optimal announcements and welfare values under the preferred calibration. Furthermore, in the baseline calibration the computed optimum features $\alpha > 0$.

Selection convention. Proposition 5 implies that for each $(i_0, \alpha, b, \gamma, s_0)$ the set of $t = 0$ equilibrium worst-case selections is $\mathcal{E}_{\alpha,b,\gamma}(s_0) \subseteq [-1, 1]$. For the remainder of the paper, fix a measurable selection

$$\vartheta_{\alpha,b,\gamma}(s_0) \in \mathcal{E}_{\alpha,b,\gamma}(s_0)$$

as follows: when $\mathcal{E}_{\alpha,b,\gamma}(s_0)$ is a singleton, $\vartheta_{\alpha,b,\gamma}(s_0)$ is its unique element; when $\alpha \in (\alpha^\star, 0)$ and $|T(s_0; \alpha, b)| \leq \gamma$, select the intermediate equilibrium $\vartheta_{\alpha,b,\gamma}(s_0) = T(s_0; \alpha, b)/\gamma$; and when $\alpha = 0$ and $b = \mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$, set $\vartheta_{0,b,\gamma}(s_0) = 0$ (otherwise $\vartheta_{0,b,\gamma}(s_0) = \pm 1$ as in Proposition 5).

With this selection convention in place, the equilibrium objects $(\pi_0, \vartheta_{\alpha,b,\gamma})$ are single-valued functions of $(i_0, \alpha, b, \gamma, s_0)$, which allows us to define a well-posed ex-ante loss for the central bank. We can therefore characterize optimal policy using first-order conditions, taking into account that inflation-anchored announcements feed back into contemporaneous inflation and thus affect equilibrium losses.

4.3 Optimal Inflation-Anchored Announcements

In what follows, I present and discuss the first-order conditions that pin down 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, we characterize its first-order condition jointly with the optimal announcements'.

Proposition 6. *Let* $W_1 := \beta\sigma^2(1 + \lambda_{cb}\kappa^2)$*. Fix the equilibrium selection* $\vartheta_{\alpha,b,\gamma}(s_0)$ *as in the Selection Convention following Proposition 5. For each* (i_0, α, b, γ) *with* $\gamma \geq 0$ *and* $\alpha \neq \alpha^\star := 1/A_\pi^{ie}$*, let* $\pi_0(i_0, \alpha, b, \gamma; s_0)$ *denote the associated* $t = 0$ *equilibrium inflation under this selection as characterized in Proposition 5. Define*

$$i_1^e(i_0, \alpha, b, \gamma; s_0) := \alpha\pi_0(i_0, \alpha, b, \gamma; s_0) + b + \gamma\vartheta_{\alpha,b,\gamma}(s_0).$$

Let

$$\tilde{\Psi}(i_0, \alpha, b, \gamma) := \mathbb{E}\left[(A_y^r r_0^n + A_y^u u_0 - \sigma i_0 + A_y^{ie} i_1^e)^2 + \lambda_{cb}(A_\pi^r r_0^n + A_\pi^u u_0 - \sigma\kappa i_0 + A_\pi^{ie} i_1^e)^2\right],$$

and

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

where the coefficients are as defined in Proposition 1. The central bank solves

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

over (i_0, α, b, γ) *with* $\alpha \neq \alpha^\star$ *and* $\gamma \geq 0$*. There exists a minimizer* $(i_0^*, \alpha^*, b^*, \gamma^*) \in \mathbb{R} \times \overline{\mathbb{R}} \setminus \{\alpha^\star\} \times \mathbb{R} \times \overline{\mathbb{R}}_+$ *(i.e., the minimum may entail* $|\alpha| = \infty$ *and/or* $\gamma = +\infty$*). Moreover,* $(i_0, \alpha, b, \gamma) \mapsto \tilde{\mathcal{J}}(i_0, \alpha, b, \gamma)$ *is generically continuously differentiable on* $\mathbb{R} \times \mathbb{R} \setminus \{\alpha^\star\} \times \mathbb{R} \times \mathbb{R}_+$ *and extends continuously to* $\gamma = 0$ *using the right derivative.*[20] *If the minimum is attained at*

[20] The objective is piecewise smooth because the equilibrium selection $\vartheta_{\alpha,b,\gamma}(s_0)$ is a projection of the affine threshold $T(s_0; \alpha, b)$ and the $t = 1$-loss term involves $(|z_{\alpha,b,\gamma}| - \gamma)_+$. Kinks occur at the events $|T(s_0; \alpha, b)| = \gamma$ and $|z_{\alpha,b,\gamma}(s_0)| = \gamma$. Because s_0 is Gaussian, these events have probability zero whenever $T(s_0; \alpha, b)$ and $z_{\alpha,b,\gamma}(s_0)$ are nondegenerate (i.e., have positive variance), which fails only on knife-edge parameterizations. Even in those cases $\tilde{\mathcal{J}}$, is well-defined and everything goes through using right derivatives.

a finite α and finite γ, it satisfies the following first–order conditions (if not, the boundary conditions hold):

$$\begin{aligned}
&\frac{1}{1-\alpha A_\pi^{i^e}}\,\mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma}(s_0)\left(-\alpha\sigma\kappa+\gamma\frac{\partial\vartheta_{\alpha,b,\gamma}}{\partial i_0}\right)\right]-2\sigma\,\mathbb{E}\left[y_0+\lambda_{cb}\kappa\pi_0\right]\\
&+\frac{2}{1-\alpha A_\pi^{i^e}}W_1\,\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+\operatorname{sign}(z_{\alpha,b,\gamma})\alpha\left(-\sigma\kappa+A_\pi^{i^e}\gamma\frac{\partial\vartheta_{\alpha,b,\gamma}}{\partial i_0}\right)\right]=0,\\
&\frac{1}{1-\alpha A_\pi^{i^e}}\,\mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma}\left(\pi_0+\gamma\tfrac{\partial\vartheta_{\alpha,b,\gamma}}{\partial\alpha}\right)\right]+2W_1\,\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+\operatorname{sign}(z_{\alpha,b,\gamma})\left(\pi_0+\alpha\tfrac{\partial\pi_0}{\partial\alpha}\right)\right]=0,\\
&\frac{1}{1-\alpha A_\pi^{i^e}}\,\mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma}\left(1+\gamma\frac{\partial\vartheta_{\alpha,b,\gamma}}{\partial b}\right)\right]+2W_1\,\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+\operatorname{sign}(z_{\alpha,b,\gamma})\left(1+\alpha\frac{\partial\pi_0}{\partial b}\right)\right]=0,\\
&\frac{1}{1-\alpha A_\pi^{i^e}}\,\mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma}\left(\vartheta_{\alpha,b,\gamma}+\gamma\frac{\partial\vartheta_{\alpha,b,\gamma}}{\partial\gamma}\right)\right]+2W_1\,\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+\operatorname{sign}(z_{\alpha,b,\gamma})\alpha\frac{\partial\pi_0}{\partial\gamma}\right]\\
&-2W_1\,\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+\right]\geq 0,
\end{aligned}$$

where

$$\tilde{G}_{\alpha,b,\gamma}(s_0):=2A_y^{i^e}\left(A_y^r r_0^n+A_y^u u_0-\sigma i_0+A_y^{i^e}i_1^e\right)+2\lambda_{cb}A_\pi^{i^e}\left(A_\pi^r r_0^n+A_\pi^u u_0-\sigma\kappa i_0+A_\pi^{i^e}i_1^e\right).$$

Here π_0 is the equilibrium inflation from Proposition 5 (i.e., it accounts for the feedback effect through beliefs). The last FOC corresponds to the complementary slackness condition associated with the constraint $\gamma\geq 0$. It holds with equality whenever $\gamma^*>0$, and as a right-derivative inequality when $\gamma^*=0$.

Key differences with the baseline model. Beyond the direct consequence of using π_0 as the conditionable variable (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.* Allowing for $\gamma=\infty$ captures the possibility of pure Delphic announcements, that is, the central bank issuing a statement that places no effective constraint on future policy. With exogenous signals this is always dominated, since widening the band indefinitely would make the private sector jump between arbitrarily high or arbitrarily low policy expectations, so $t=0$ volatility would diverge. With inflation-anchored announcements, by contrast, as the band widens, an extreme worst-case expectation for the private sector is not sustainable as an equilibrium unless the shock realizations are very large. This is because overly high or overly low rate expectations move inflation in the opposite direction, hence feeding back through the mean-band term to reduce the magnitude of those expectations. Hence, at the limit of purely Delphic announcements, an equilibrium with a non-degenerate aggregate

worst-case selection will arise almost surely. Further, this equilibrium is insensitive to γ, as widening the band moves both endpoints out symmetrically, and the fraction of agents at each endpoint rebalances in exactly the inverse proportion so that the average worst-case expectation doesn't move. Consequently, as $\gamma \to +\infty$, the central bank's overall loss converges to a finite constant, making it an admissible solution. However, this may not be optimal, since by picking fully vague announcements the worst-case selection $\gamma\vartheta_{\alpha,b,\gamma}$ is not truncated and thus induces a potentially high output and inflation volatility at $t = 0$.

2. *An additional effect of vagueness.* With exogenous signals, the extent of vagueness of the announcements only affects $t = 1$ losses by changing the flexibility to react if the ex-post optimal rate is not what was announced. With inflation-anchored announcements, vagueness also shifts the band center because inflation π_0 depends on the band's width. 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, worst-case rates for the private sector are too low ones, 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 stabilizing announcements ($\alpha > 0$ or $\alpha < \alpha^\star$), two effects work in opposite directions. Higher α dampens the impact of beliefs (the worst-case term $\gamma\vartheta_{\alpha,b,\gamma}$) 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 $\alpha > 0$ (and similarly for $\alpha < \alpha^\star$). 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 4).

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 welfare cost of strict rules (the additional losses from constraining the central bank to make fully precise announcements). Fourth, I compare optimal data dependence to the Delphic (best-forecast) benchmark and discuss the implications for optimal communication. Finally, I show that the qualitative patterns are robust to sensible parameter changes.

5.1 Calibration

The baseline parameterization draws from standard quarterly New Keynesian specifications. The Phillips curve slope is set to $\kappa \approx 0.02$, consistent with the calibration in Rotemberg and Woodford (1997), which implies a moderate degree of price stickiness. The discount factor is $\beta = 0.99$, corresponding to an annual real interest rate of approximately 4%. The intertemporal elasticity of substitution is set to $\sigma = 6.25$ following the calibration in Rotemberg and Woodford (1997) (admittedly on the upper end of the range of values reported in the literature).

The natural rate of interest has a mean of $\overline{r}^n = 0.65\%$ (quarterly) and follows an AR(1) process with persistence $\rho_r = 0.35$ and standard deviation $\sigma_r = 0.65\%$. The latter is taken from the New York Fed DSGE forecast of December 2024, ensuring the calibration reflects current estimates of demand-side uncertainty. Cost-push shocks have zero mean ($\overline{u} = 0\%$) and follow an AR(1) process with persistence $\rho_u = 0.8$ and standard deviation $\sigma_u = 0.25\%$. The higher persistence of cost-push shocks relative to natural rate shocks reflects the typically more persistent nature of supply-side disturbances (Woodford, 2003).

The central bank and private sector share the same loss function weights $\lambda_{cb} = \lambda_{ps} \approx 400$, matching the welfare-based relative weight on inflation stabilization implied by our New—Keynesian parameterization and normalization. For the private sector, adopting the welfare-based benchmark is the most natural choice. For the central bank, an alternative specification with substantially lower inflation weight (e.g., $\lambda_{cb} = 1$, a pure "dual mandate" central bank) would entail a sizable preference disagreement. To isolate the model's core mechanisms from such considerations, I set $\lambda_{cb} = \lambda_{ps}$. Exploring the consequences of heterogeneous preferences would be a natural extension in a richer quantitative analysis.

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

$$\alpha^* \approx 2.5$$

and

$$\gamma^* \approx 0.0034,$$

which corresponds to an interval width of about 68 basis points (roughly 272 basis points when annualized for comparison to policy-rate projections).

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 remarkably 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.

For the vagueness dimension, a rough empirical counterpart can be drawn from the range of policy-rate projections reported by FOMC participants in the December 2024 Summary of Economic Projections, which spans roughly 130 basis points (in annualized terms). Although this is far from an ideal analogue (it reflects individual projections rather than a collective commitment) it nonetheless provides a reasonable check that the model's magnitude of vagueness is not implausible. A market-based comparison can be made using the Kansas City Fed's *Measure of Policy Rate Uncertainty*, which roughly implies a 430 basis-point range for the 95% confidence interval of policy rates over the same horizon (also annualized). Both this measure and the SEP dispersion, however, reflect unconditional uncertainty: they incorporate variation in π_0 and other anchors for future rates. As such, they are not directly comparable to what the model terms *vagueness*. Moreover, because the market-based measure is derived from asset prices, it reflects risk-adjusted beliefs rather than the underlying probability assessments of future policy rates.

Having established that the model's quantitative implications are broadly reasonable, a natural next step is to study how optimal announcements vary with crucial structural features of the economy: demand and supply (cost-push) uncertainty.[21]

Standard deviation of the natural rate σ_r. When uncertainty about the natural rate increases, it raises $t = 0$ output and inflation volatility through both direct and indirect channels. The latter operates through worst-case policy expectations, since private-sector beliefs about the future rate depend on the realization of the contemporaneous natural rate. Greater uncertainty also increases the volatility of the gap between the mean of the announced interval and the conditional expectation of the full-information optimal rate. By increasing data dependence, the central bank simultaneously mitigates both effects: the feedback from beliefs dampens the response of $t = 0$ inflation to shocks, and the coefficient of the natural-rate component in the gap narrows. Figure 4a shows this upward adjustment of α.

The reaction of vagueness is more nuanced. Higher demand uncertainty raises both the benefit and the cost of making less precise announcements. On the one hand, higher vagueness provides flexibility to accommodate greater uncertainty regarding $t = 1$ outcomes; on the other, it amplifies current volatility because beliefs about the worst-case rate become more extreme. In general the marginal value of vagueness is non-monotonic in α (see the discussion below Proposition 6), but under this parameterization it is increasing. When σ_r is low, α is low and the $t = 0$ cost channel dominates, so greater uncertainty leads to more precise announcements. As σ_r rises, data dependence rises with it and eventually reverses the effect: at high σ_r, optimal vagueness increases because stronger data dependence raises the marginal value of flexibility. The result is a U-shaped reaction in vagueness as uncertainty grows, driven by the monotonic increase in α (Figure 4b).

In short, higher demand uncertainty calls for announcements that are more responsive to data and—once uncertainty is sufficiently large—also more vague.

[21]The discussion for the underlying mechanisms is supported by additional, unreported numerical exercises.

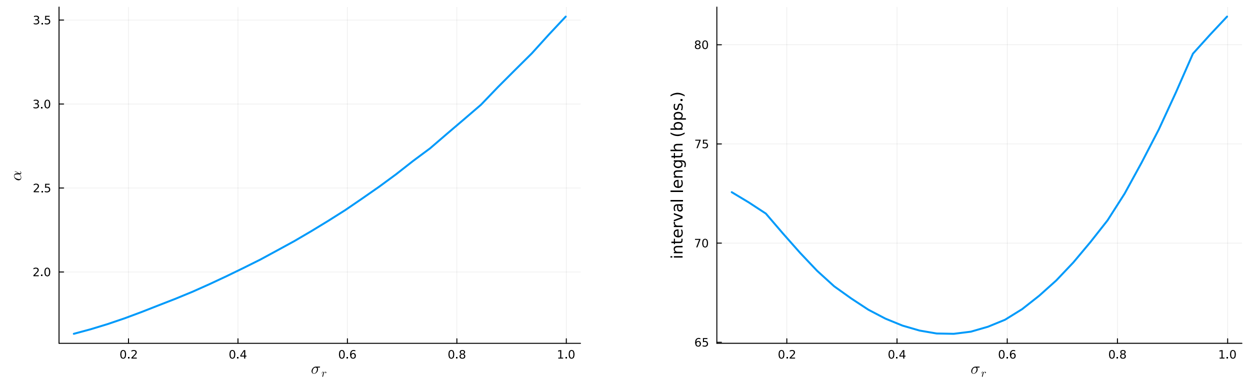

(a) Optimal Data Dependence

(b) Optimal Vagueness

Figure 4: The Optimal Communication Policy for different values of σ_r.

Standard deviation of cost-push shocks σ_u. When uncertainty about cost-push shocks increases, the logic of optimal data dependence policy partially reverses. Higher σ_u makes outcomes more volatile in both periods, as does higher σ_r, but it also reduces the usefulness of $t = 0$ inflation for predicting the ex-post optimal rate. In this environment, keeping policy highly data-dependent would cause the central bank to overreact to cost-push disturbances. Consequently, the central bank optimally reduces data dependence, lowering the sensitivity of the announced center of the interval to the public signal. While this adjustment improves expected outcomes at $t = 1$, it increases volatility at $t = 0$ by weakening the beliefs-feedback loop. Figure 5a shows this downward shift in α.

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 slightly. 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.

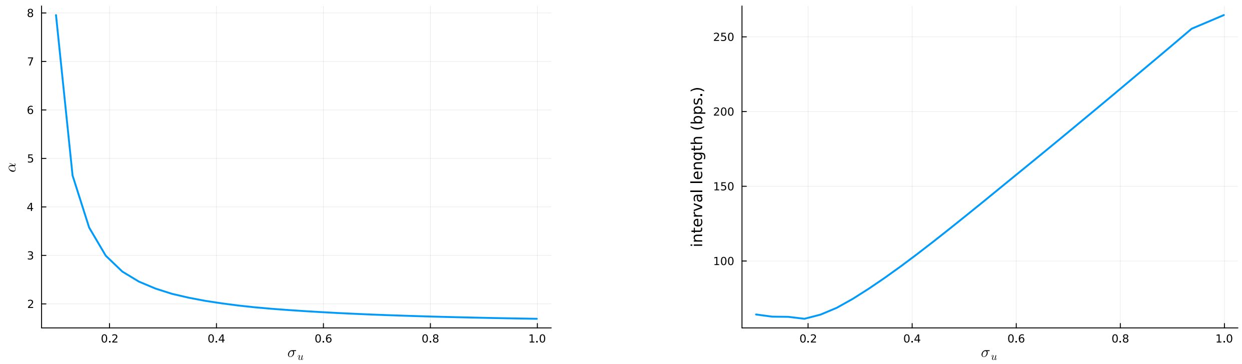

(a) Optimal Data Dependence

(b) Optimal Vagueness

Figure 5: The Optimal Communication Policy for different values of σ_u

Together, these results highlight that not all uncertainty calls for the same communication strategy. While very high uncertainty —whether from demand or supply— justifies less precise announcements that preserve flexibility to react, the optimal degree of data dependence 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. When uncertainty is primarily supply-driven, by contrast, optimal communication calls for downplaying the emphasis on inflation in its decisionmaking. More broadly, this can be interpreted as favoring a less data-point–specific form of guidance, one that conditions policy on the overall state of the economy rather than narrowly on realized inflation outcomes. Together with the associated increase in vagueness, this entails an overall less informative form of forward guidance.

A second lesson also emerges from these results. The two dimensions of communication —data dependence and vagueness— cannot 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 converse, however, is not true: to a first approximation, optimal data dependence is largely unaffected by 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 Rules

This section revisits the classic debate on rules versus discretion in monetary policy through the lens of the model. Since Friedman's early call for fixed, mechanical 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 modern literature, beginning with Kydland and Prescott (1977), formalized the argument that commitment to rules improves outcomes by allowing the policymaker to internalize how its actions shape private expectations. Subsequent contributions, notably Taylor (1993) and the large literature it inspired, made this discussion operational by expressing systematic policy as a rule in which the policy rate responds to deviations of inflation and output from their targets. Yet central bankers have often defended discretion, emphasizing the need for flexibility amid evolving data and unforeseen disturbances (see, e.g., Blinder (2002),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 that follows uses this setting to provide an illustrative assessment of how costly a "pure rules" policy can be, comparing welfare under optimally vague announcements with that under fully precise ones ($\gamma = 0$).

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 model for studying monetary policy. The parameter α captures the stabilizing role of systematic responses—rules that anchor expectations and smooth current outcomes—while γ formalizes the policymaker's ability to retain flexibility in the face of unforeseen developments. The optimal announcement thus combines these two dimensions, yielding a mix of commitment and discretion rather than either extreme.

Under the baseline parameterization, a regime of zero discretion raises expected losses by roughly 6 percent 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 modeling flexibility explicitly and motivates further quantitative explorations.

Beyond the baseline magnitude, it is informative to examine how the welfare 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.

The relative value of inflation stability λ. See Figure 6 below. When the policymaker places little weight on inflation relative to output (low λ), the optimal communication strategy is to make fully precise, binding announcements. In this case, the primary concern is to stabilize output in the face of volatile demand shocks, whose impact on output fluctuations is much larger than on inflation—both directly and through expectations. Fully precise announcements minimize the belief-induced volatility of output at $t = 0$ by "calming the markets", making a zero-discretion (rule-based) policy effectively optimal. Consequently, the welfare cost of rules is very small.

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 $t = 1$ fleixibility (see the discussion on the comparative static of the optimal policy with respect to demand uncertainty above). Consequently, the monotonicity in optimal vagueness (and, therefore, in the welfare cost of forbidding it) as λ increases operates through the response of data dependence. Indeed, additional non-reported exercise, show that varying λ while holding α fixed keeps optimal vagueness fairly constant.

Because λ spans a wide range of empirically and normatively relevant values, from the "dual mandate" case ($\lambda = 1$) to the welfare-grounded calibration ($\lambda \approx 400$), understanding how communication design depends on this parameter is central for interpreting the model's policy lessons.

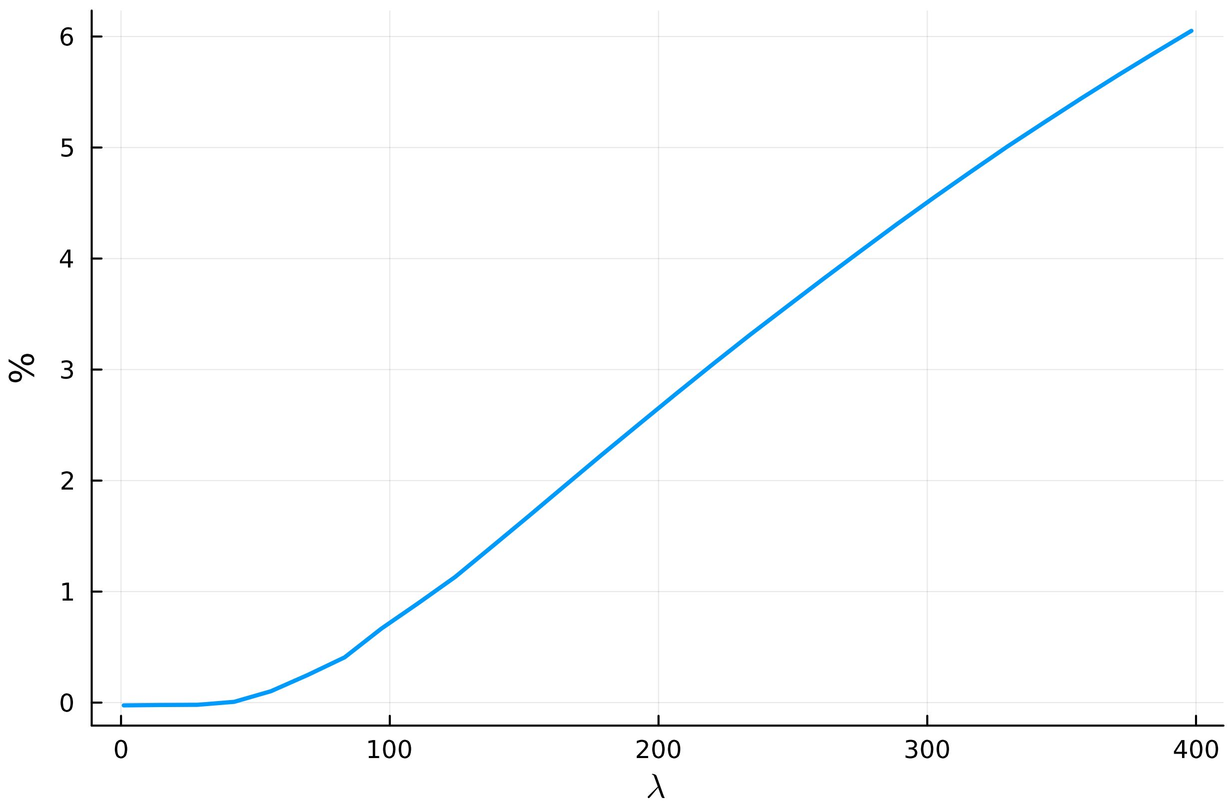

Figure 6: Additional losses under pure rules for different values of λ

The slope of the Phillips Curve κ. As the Phillips curve steepens, inflation becomes more sensitive to real activity in both periods. This amplifies two related forces that raise the value of retaining some flexibility in communication. First, at $t = 0$, the higher responsiveness of inflation to shocks makes expectations—and hence current inflation—more volatile. To contain these fluctuations, the central bank optimally increases data dependence, as the feedback it creates through beliefs helps dampen the sensitivity of inflation to shocks. 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: with a larger κ, small state realizations generate larger inflation movements, so the over- or under-tightening implied by fully precise announcements becomes more costly for a given volatility of shocks.[22] Both channels make precise, rule-like guidance increasingly costly as κ rises, so the welfare cost of rules increases with the slope of the Phillips curve. In our calibration

[22] Admittedly, this sort of "mechanical" effect of a change in κ would suggest adjusting the weight λ as κ varies (indeed, the welfare-grounded mapping features $\lambda \propto 1/\kappa$). However, here λ is held fixed to isolate the direct effect of κ on optimal communication design.

the latter is stronger, highlighting the importance of studying optimal communication across both dimensions jointly.

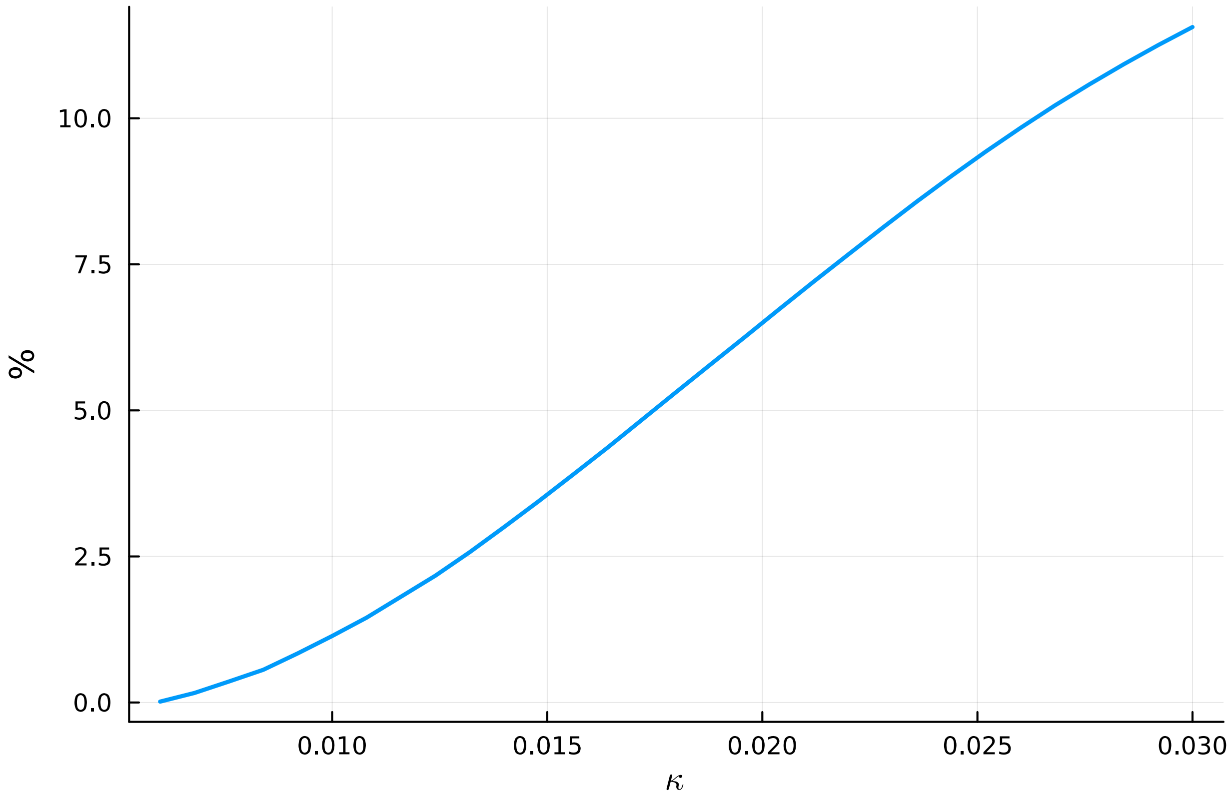

Figure 7: Additional losses under pure rules for different values of κ

5.4 On Odyssean and Delphic Forward Guidance

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 to influence expectations today, and Delphic guidance, which conveys the central bank's best forecast of its future policy given current information (Campbell et al., 2012). In standard New Keynesian frameworks with rational expectations, this distinction collapses: under commitment, the best prediction of the future rate coincides with the announced rate, while under discretion any announcement lacks force, as the private sector anticipates that it will be abandoned if ex post optimality so requires.

By contrast, the present framework makes the distinction both meaningful and quantifiable. Because announcements shape current outcomes through the expectations channel, the optimal announcement internalizes its stabilization effects on inflation and output today, whereas a purely Delphic projection does not. In equilibrium, this feedback between announced guidance and beliefs implies that the rate the central bank should announce generally differs from its best forecast of the ex-post optimal rate. Allowing announcements to be somewhat vague softens this divergence: by leaving room for adjustment within a range, the central bank can ex post bring the realized rate closer to its earlier internal projection while remaining consistent with its guidance. Within this setting, Odyssean guidance corresponds to a bounded form of commitment: a commitment to a range of future actions rather than a precise path.[23] For concreteness, the Delphic benchmark is defined as the conditional linear forecast of the ex-post optimal rate based on current inflation,

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

that is, the regression coefficient of the optimal $t = 1$ policy rate on today's inflation.[24]

This structure allows for a systematic analysis of when and why optimal guidance diverges from purely Delphic communication. Under the baseline parameterization, the central bank's optimal announcement differs sharply from its best projection of the ex-post optimal rate (as defined above). The latter entails $\alpha^d \approx 1.1$, while the optimal announcement features a substantially higher $\alpha^* \approx 2.5$. This gap reflects the stabilization motive embedded in optimal policy: by committing to a stronger response, the central bank shapes beliefs in a way that dampens current inflation and output volatility.

The following numerical exercises show how the Odyssean–Delphic gap responds to changes in demand and cost-push uncertainty.

Standard deviation of natural rate shocks σ_r. 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 always above α^d. Second, the gap between them widens as σ_r increases.

The fact that $\alpha^* > \alpha^d$ throughout reflects the stabilizing role of forward guidance in the model. The Delphic projection, α^d, by construction simply a best forecast, does not internalize how announcements feed back into current outcomes through expectations. By contrast, the optimal announcement α^* is chosen not only to match the preferred rate at $t = 1$ as closely as possible, but also to stabilize inflation and output at $t = 0$, taking into account that expectations of the future rate directly affect current demand and price-setting. Be-

[23] This interpretation aligns with the model's structure: as announcements are intervals, commitment should be understood as applying to that range rather than to a point forecast.

[24] Note 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.

cause a stronger response to inflation today anchors beliefs and dampens contemporaneous volatility, the optimal policy “leans into” the expectations channel more aggressively than the pure forecast. This stabilization motive pushes α^* above the purely predictive α^d. This connects directly to the earlier discussion regarding stabilizing vs. destabilizing announcements. While any announcement with positive data dependence is stabilizing, the extent of this stabilization mechanism depends on how large it is.

On the other hand, the widening gap between α^* and α^d as demand uncertainty rises follows from the amplification of volatility that the expectations channel entails. When natural-rate shocks become more volatile, both output and inflation at $t = 0$ respond more strongly to r_0^n, directly and through beliefs, increasing the benefit of using announcements to calm markets and thus strengthening the stabilization mechanism. This stabilization motive for raising α comes on top of the “projection motive” which is always there given announcements bind the feasible range of future rates. Consequently, the increase in α^* as demand uncertainty increases is greater than the increase in α^d, generating this widening gap. Finally, note that the rise in the projection coefficient itself reflects the partial nature of stabilization through announcements: if announcements fully offset the increased sensitivity of the economy to demand shocks, the projection coefficient would remain constant.

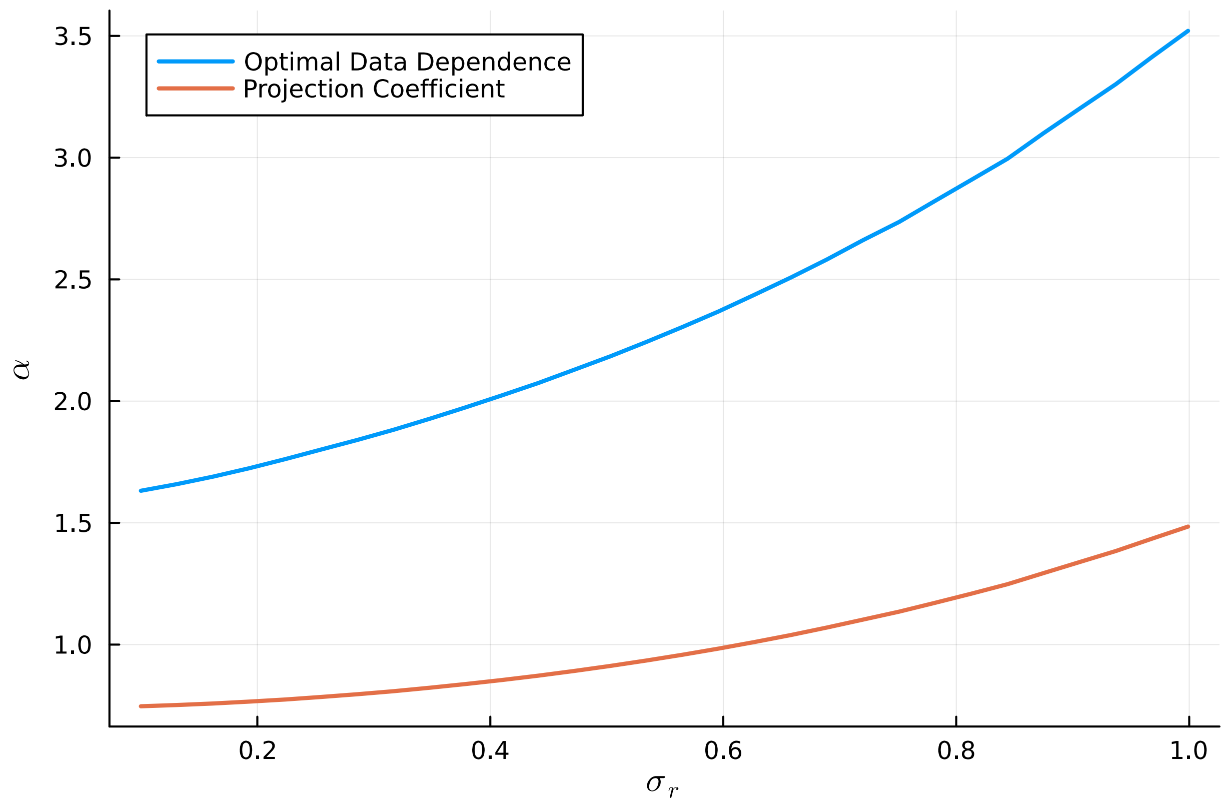

Figure 8: Optimal Data Dependence and Optimal Projection for different values of σ_r

Standard deviation of cost-push shocks σ_u. Figure 8 (right panel) 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. The decline in both coefficients reflects that higher cost-push volatility makes current inflation a noisier signal of the ex-post optimal rate: a larger share of π_0 now reflects transitory cost-push disturbances rather than information relevant for future policy, reducing the projection α^d. The stronger decline in α^* arises because the $t = 0$ stabilization motive weakens at the same time. When inflation becomes more driven by cost-push shocks, adjusting the degree of data dependence does less to stabilize π_0—the beliefs channel loses power, as π_0 is less responsive to α. Thus, the diminishing marginal benefit of the stabilization motive in optimal data dependence generates a narrowing of the gap between α^* and α^d.

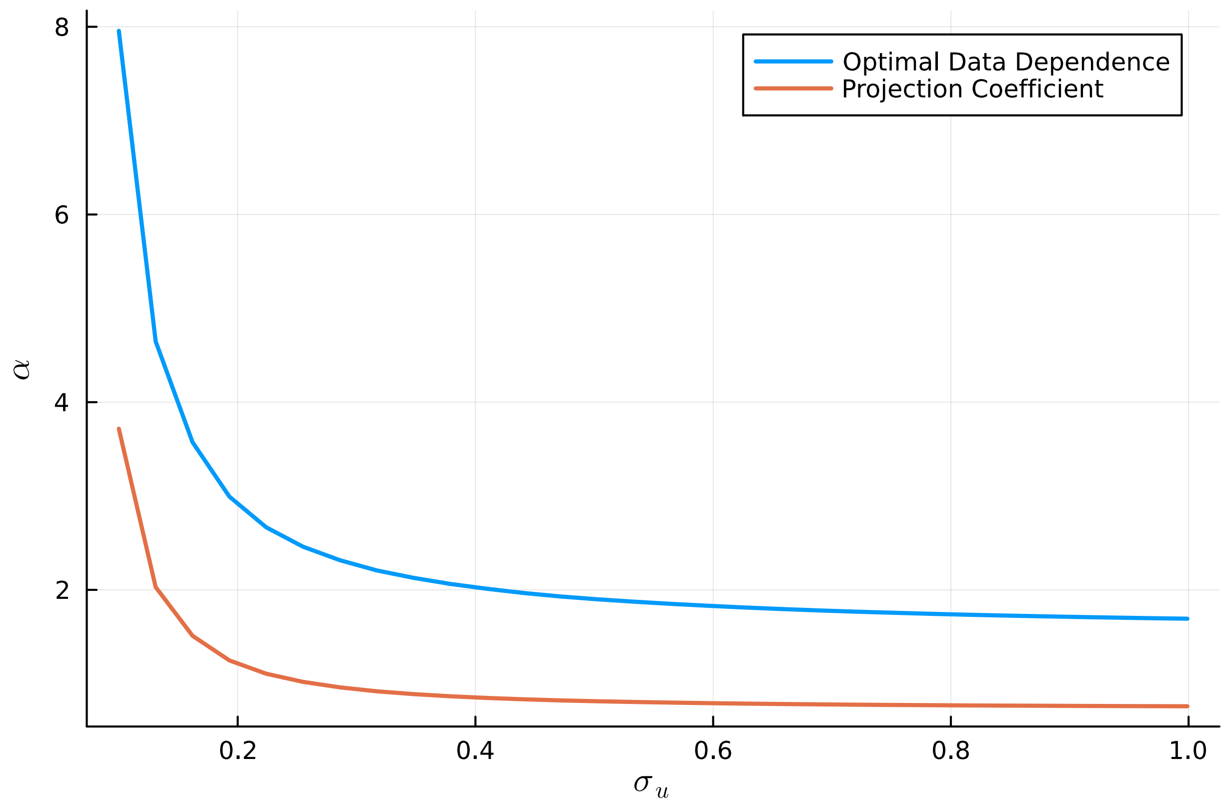

Figure 9: Optimal Data Dependence and Optimal Projection for different values of σ_u

Taken together, these results indicate that the relevance of policy expectations for economic stabilization —captured by the gap between Odyssean and Delphic guidance— depends crucially on 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, communication should convey a more measured responsiveness, closer to the central bank's own forecasts, as the informational content of inflation weakens and the benefits of aggressive commitment diminish.

5.5 Robustness

The numerical results presented above rely on a particular calibration. A natural question is whether the qualitative patterns (the signs of the comparative statics) are robust to alternative parameterizations. This subsection addresses this concern through a systematic sensitivity analysis.

The analysis proceeds as follows. Each parameter is varied one at a time over a wide neighborhood of the baseline—roughly 40% to 300% of the baseline value—with exceptions driven by (i) natural restrictions (positivity, $\beta < 1$, $\rho_r < 1$, $\rho_u < 1$) and (ii) uncertainty in the range of plausible calibrations reported in the literature.[25] The parameters varied include κ, the loss function weight on inflation λ (shared by the central bank and the private sector), the discount factor β, the intertemporal elasticity of substitution σ, the persistence of natural rate shocks ρ_r, the persistence of cost-push shocks ρ_u, the mean natural rate $\overline{r}^n$, and the standard deviations of both shock processes σ_r and σ_u.[26] For each configuration, the model is solved and the following relationships are evaluated:

- Optimal data dependence α^* is increasing in demand uncertainty σ_r
- Optimal data dependence α^* is decreasing in supply uncertainty σ_u
- The welfare cost of rules (prohibiting vagueness) is increasing in λ
- The welfare cost of rules is increasing in κ
- The Odyssean-Delphic gap $(\alpha^* - \alpha^d)$ is increasing in σ_r
- The Odyssean-Delphic gap is decreasing in σ_u
- Optimal vagueness γ^* is eventually increasing in σ_r[27]
- Optimal vagueness γ^* is eventually increasing in σ_u

Table 1 summarizes the results of the robustness exercise. Each cell reports the lowest/highest tested deviation as a percentage of the baseline at which the corresponding qualitative relationship fails ("R" indicates that no failure is detected on that side within the tested range). Overall, the six benchmark comparative statics are robust across most parameter configurations, with the few exceptions reported in the table.

[25]For example, the lowest value of the slope of the Phillips curve κ is below 40% of the baseline following Hazell et al. (2022).

[26]The resulting tested ranges in levels are: $\kappa \in [0.0062, 0.0577]$, $\lambda \in [1, 1000]$, $\beta \in [0.9, 0.999]$, $\sigma \in [1, 18.75]$, $\rho_r \in [0.14, 0.875]$, $\rho_u \in [0.32, 0.96]$, $\overline{r}^n \in [0.26\%, 1.95\%]$, $\sigma_r \in [0.26\%, 1.95\%]$, and $\sigma_u \in [0.10\%, 0.75\%]$.

[27]That is, there exists a σ_r^* such that $\gamma^*(\sigma_r) > \gamma^*(\sigma_r^*)$ for all $\sigma_r > \sigma_r^*$ (analogously for σ_u).

Param	$\alpha \uparrow \sigma_r$	$\alpha \downarrow \sigma_u$	$L \uparrow \lambda$	$L \uparrow \kappa$	$G \uparrow \sigma_r$	$G \downarrow \sigma_u$	$\gamma \uparrow \sigma_r$	$\gamma \uparrow \sigma_u$
κ	R/R	R/R	R/R	—	R/R	R/R	40%/R*	R/R
λ	R/R	R/R	—	R/R	R/R	R/R	R/R	R/250%*
β	R/R	R/R	R/R	R/R	R/R	R/R	R/R	R/R
σ	R/150%	R/R	R/R	R/R	R/150%	R/R	16%/150%*	16%/R*
ρ_r	R/200%	R/R	R/R	R/R	R/200%	R/R	R/R	R/R
ρ_u	R/R	R/R	R/R	R/R	R/R	R/R	R/R	R/R
$\bar{r}^n$	R/R	R/R	R/R	R/R	R/R	R/R	R/R	R/R
σ_r	—	R/R	R/R	R/R	—	R/R	—	R/R
σ_u	R/R	—	R/R	R/R	R/R	—	50%/R*	—

Table 1: Robustness threshold search results. Each entry is "low/high" where "low" (resp. "high") reports the lowest (resp. highest) tested deviation from baseline at which the relationship fails; "R" indicates robustness within the tested range; and "—" denotes a not-applicable case in which the parameter being perturbed coincides with the variable defining the comparative static. In the two rightmost columns, * denotes that the apparent failure is driven by non-fundamental numerical reasons (e.g. γ being weakly identified on the tested uncertainty grid), rather than a genuine reversal of the eventual increase property (see explanation below).

Note that the six reported lack of robustness of the "eventually increasing" relationships (two rightmost columns of the table) are not meaningful failures of the tested property. In all cases but $\kappa = 40\%$, these come from the optimal policy entailing essentially zero commitment (large γ^*). As the objective function is extremely flat at large γ (from some point onward, more flexibility is almost never used and doesn't move private sector's expectations by much), the relationship is rejected because the solver jumps between very large values of γ in a non-monotonic way. On the other hand, for $\kappa = 40\%$, the reported non-robustness is due to not considering a large enough range for σ_r. Once expanded, γ^* does become increasing at some large σ_r, confirming the robustness of the property.

Thus, the only real non-robust features are the monotonicity of data dependence α and the Odyssean-Delphic gap with respect to demand uncertainty at large values of σ and ρ_r respectively.[28] While the data dependence and the Odyssean-Delphic gap comparative statics fail at the same points, the comparative static of α^d (the optimal projection value) also breaks at those points (not reported). Thus, at these parameter values (very large σ and large ρ_r) the effects of demand uncertainty σ_r on both (i) the best linear projection of the ex-post optimal rate on π_0 and (ii) the contemporaneous stabilization value of announcements flip. Fortunately, in the case of the elasticity of intertemporal substitution σ this is not

[28] It is worth noting that at 40% of the baseline value of σ_u, data dependence α and the Odyssean-Delphic gap become flat in σ_r (because of Monte Carlo error, the *flat* label is applied whenever the relationship has overall a slope of less than 0.1). Similarly, at 300% of the baseline value of κ, the welfare loss from strict rule-like guidance becomes flat in λ. Thus, in their strict monotonicity form, these relationships also break at these values.

particularly worrying inasmuch as the baseline value, 6.25, is already at the upper end of the range of plausible values (if anything, it is reassuring that the comparative statics hold at values of σ much lower than the baseline).

Overall, these findings strongly suggest that the qualitative insights from the numerical analysis are not artifacts of the particular calibration chosen. The asymmetric effects of demand versus supply uncertainty on optimal communication, the substantial welfare costs of strict rule-like guidance, and the systematic divergence between optimal and Delphic announcements all emerge as robust features of the model across a wide range of plausible parameterizations.

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 signal 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 welfare losses under a standard calibration. Third, optimal forward guidance differs sharply from the best forecast of future rates (Delphic guidance), as the central bank internalizes the stabilizing role of expectations. These lessons are shown to be robust to sensible changes in the model's parameters. 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 both are optimized simultaneously.

The analysis also suggests natural directions for future work. The current model assumes that commitments are fully credible and binding, but in practice credibility evolves over time depending of the correspondence between past announcements and subsequent actions, and earlier guidance for distant horizons interacts with the design of new announcements as those horizons approach, creating additional trade-offs. 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.

Barro, R. J. and Gordon, D. B. (1983), ‘A positive theory of monetary policy in a natural rate model’, *Journal of Political Economy* **91**(4), 589–610.

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. (2002), *Central Banking in Theory and Practice*, The Lionel Robbins Lectures, MIT Press, Cambridge, MA.

Bowman, M. W. (2022), ‘Forward Guidance as a Monetary Policy Tool: Considerations for the Current Economic Environment’, Speech. Money Marketeers of New York University.

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.

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).

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.

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), ‘Monetary policy & anchored expectations—An endogenous gain learning model’, *Journal of Monetary Economics* **140**, S37–S47.

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), 39.

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* .

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. (2026), 'Disagreement about monetary policy', *American Economic Journal: Macroeconomics* . Forthcoming.

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.

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.

Appendices

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 $A : (\mathbb{R}, \mathcal{B}(\mathbb{R})) \to (\mathcal{K}(\mathbb{R}), \mathcal{E})$. 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 $L, U : \mathbb{R} \to \mathbb{R}$ 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 $A : \mathbb{R} \to \mathcal{K}(\mathbb{R})$, **the set of admissible private sector's 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.*

- $p(B \times C) = \int_C K(B \mid \tilde{s}_1) F_{(r_1^n, u_1, \tilde{s}_1)}\Big(dr_1^n, du_1, d\tilde{s}_1 \mid s_0\Big) \quad \forall B \in \mathcal{B}(\mathbb{R}),\ C \in \mathcal{B}(\mathbb{R}^3),$

- $\operatorname{supp} K(\cdot \mid \tilde{s}_1) \subseteq \mathcal{A}_1(\tilde{s}_1)$ *for $F_{\tilde{s}_1}(\cdot \mid s_0)$-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

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

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 $A(\tilde{s}_1)$,

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

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 $\mathcal{A}_1$

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

Similarly for the NKPC,

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

Thus, defining

$$A_y^r := \sigma\left(1 + (1+\sigma\kappa)\rho_r\right), \qquad A_y^u := \sigma\rho_u, \qquad A_y^{i^e} := -\sigma(1+\sigma\kappa),$$
$$A_\pi^r := \sigma\kappa(1 + (1+\sigma\kappa+\beta)\rho_r), \qquad A_\pi^u := 1 + (\sigma\kappa+\beta)\rho_u, \qquad A_\pi^{i^e} := -\sigma\kappa(1+\sigma\kappa+\beta).$$

we have

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

Define

$$\delta_r := A_y^r + \lambda_{cb}\kappa A_\pi^r, \quad \delta_u := A_y^u + \lambda_{cb}\kappa A_\pi^u, \quad \delta_{i^e} := A_y^{i^e} + \lambda_{cb}\kappa A_\pi^{i^e} = -\sigma\left[(1+\sigma\kappa) + \lambda_{cb}\kappa^2(1+\sigma\kappa+\beta)\right].$$

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

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

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

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

B.3 Proof of Proposition 2

The proof consists of three steps. First, it is shown that intervals weakly dominate any other shape of the announcement set. Second, it is proved that a constant (i.e., independent of $\tilde{s}_1$) interval width is optimal. Finally, it is shown that the center of the interval is always 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

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

The second term equals zero (use the Law of Iterated Expectations with inner expectation conditional on $t = 1$ information, Δ_i is measurable with respect to it and 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 $\tilde{s}_1$, $\tilde{A}_t(\tilde{s}_1) = \operatorname{co} A_t(\tilde{s}_1)$ (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 decreases under convexification (if the target random variable i_1^{unc} has full support and the convexification is non-trivial, it strictly decreases). 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 set[29], 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.

[29]As $A_t(\tilde{s}_1) \subseteq \mathbb{R}$, 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 $\tilde{s}_1$.

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

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

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

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

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

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

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. Therefore

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

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. *Consider an interval announcement with $L, U \in L^1$ such that $L(s) \le U(s)$ 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.,*

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

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

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

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*

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

Lemma 4. *There exist $a_{ps}, a_{cb} \in \mathbb{R}$ and affine maps a_{ps}^0, a_{cb}^0 such that*

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

Proof of Lemma 2. Applying Lemma 3 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

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

Let $Z := (\alpha - a_{ps})\tilde{s}_1 + b - a_{ps}^0(s_0)$. Clearly,

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

Thus,

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

Further, note that

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

and

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

This gives the expression in the statement. □

B.5 Proof of Proposition 3

For convenience, the proof of the special case $\bar{r}^n = 0$ relies on the proof of the general case even though the latter it is presented later in the paper (see B.6 for its proof). Thus, set $\bar{r}^n = 0$ and consider the FOCs in Proposition 4. Write $i_1^e = \alpha\mu_{\tilde{s}} + b + \gamma\theta_{\alpha,b}$, with $\mu_{\tilde{s}} := \mathbb{E}[\tilde{s}_1 \mid s_0]$ as in Lemma 2 and let

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

and $\sigma_z^2 := \mathrm{Var}(z_{\alpha,b})$ as in Proposition 4. Finally, as in Lemma 5, let

$$\begin{aligned} G_{\alpha,b,\gamma}(s_0) =& 2A_y^{i^e}\left(A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e(\alpha, b, \gamma; s_0)\right) + 2\lambda A_\pi^{i^e}\left(A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e(\alpha, b, \gamma; s_0)\right) \\ & - \frac{2\delta_{i^e}}{1+\lambda\kappa^2}\left(\delta_r \bar{r}^n + \delta_{i^e}\mathbb{E}[i_1^e]\right). \end{aligned}$$

(i) $b^* = 0$ *and* $i^* = 0$. The b–FOC from Proposition 4 reads

$$\mathbb{E}\left[G_{\alpha,b,\gamma}\left(1 + \gamma\frac{\partial\theta_{\alpha,b}}{\partial b}\right)\right] + 2W_1\mathbb{E}\left[(|z_{\alpha,b}| - \gamma)_+ \operatorname{sign}(z_{\alpha,b})\right] = 0.$$

Let $X := (\alpha - a_{cb})\tilde{s}_1 - a^0_{cb}(s_0, \hat{s}_1)$. If $\bar{r}^n = 0$, X is centered and symmetric. The map $g(x) := (|x| - \gamma)_+ \operatorname{sign}(x)$ is odd, so

$$\mathbb{E}\left[(|z_{\alpha,0}| - \gamma)_+ \operatorname{sign}(z_{\alpha,0})\right] = \mathbb{E}[g(X)] = 0,$$

hence the $t = 1$ term in the b–FOC vanishes at $b^* = 0$.

For the $t = 0$ term, note that if $\bar{r}^n = 0$, $\mathbb{E}[\mu_{\tilde{s}}] = 0$ and thus by symmetry $\mathbb{E}[\theta_{\alpha,0}] = 0$ if $b = 0$. Hence, $\mathbb{E}[i_1^e] = 0$ under $b = 0$. Therefore,

$$\mathbb{E}\left[G_{\alpha,0,\gamma}\right] = 0,$$

because the only terms that survive the expectation in $G_{\alpha,0,\gamma}$ are proportional to $\mathbb{E}[i_1^e]$. It remains to show that $\mathbb{E}\left[G_{\alpha,0,\gamma}\frac{\partial\theta_{\alpha,0}}{\partial b}\right] = 0$. Fix s_0 and let

$$\Delta_{\alpha,b} := (\alpha - a_{ps})\tilde{s}_1 + b - a^0_{ps}(s_0), \qquad m(s_0, b) := \mathbb{E}[\Delta_{\alpha,b} \mid s_0] = (\alpha - a_{ps})\mu_{\tilde{s}}(s_0) + b - a^0_{ps}(s_0),$$

$$v(s_0) := \mathrm{Var}(\Delta_{\alpha,b} \mid s_0) = (\alpha - a_{ps})^2\,\mathrm{Var}(\tilde{s}_1 \mid s_0).$$

Then,

$$\theta_{\alpha,b}(s_0) = \mathbb{E}\left[\text{sign}(\Delta_{\alpha,b}) \mid s_0\right] = 2\Phi\left(\frac{m(s_0,b)}{\sqrt{v(s_0)}}\right) - 1,$$

hence

$$\frac{\partial\theta_{\alpha,b}}{\partial b}(s_0) = \frac{2\phi\left(m(s_0,b)/\sqrt{v(s_0)}\right)}{\sqrt{v(s_0)}}.$$

In particular, the map $s_0 \mapsto \partial\theta_{\alpha,0}/\partial b$ is an *even* function of the scalar index $m(s_0,b)$. Further, r_0^n, u_0, $\mu_{\tilde{s}}(s_0)$ are linear (hence odd) in s_0. Finally,

$$\theta_{\alpha,0}(s_0) = 2\Phi\left(\frac{m(s_0,0)}{\sqrt{v(s_0)}}\right) - 1$$

is odd in $m(s_0,0)$, hence odd in s_0. Consequently, with $b=0$,

$$\mathbb{E}\left[r_0^n \frac{\partial\theta_{\alpha,0}}{\partial b}\right] = \mathbb{E}\left[u_0 \frac{\partial\theta_{\alpha,0}}{\partial b}\right] = \mathbb{E}\left[\mu_{\tilde{s}}(s_0) \frac{\partial\theta_{\alpha,0}}{\partial b}\right] = \mathbb{E}\left[\theta_{\alpha,0} \frac{\partial\theta_{\alpha,0}}{\partial b}\right] = 0,$$

because each integrand is an odd function of a centered Gaussian vector. Since $G_{\alpha,0,\gamma}(s_0)$ is a linear combination of r_0^n, u_0 and $i_1^e(\alpha,0,\gamma;s_0) = \alpha\mu_{\tilde{s}}(s_0) + \gamma\theta_{\alpha,0}(s_0)$, it follows that

$$\mathbb{E}\left[G_{\alpha,0,\gamma}\frac{\partial\theta_{\alpha,0}}{\partial b}\right] = 0.$$

Consequently, $b=0$ satisfies the $b=0$-FOC when $\bar{r}^n = 0$.

From Proposition 1 with $\bar{r}^n = 0$ and $\mathbb{E}[i_1^e] = 0$, we get $i_0^* = 0$.

(ii) *The α–FOC.* Define

$$\xi_r := 2\left(A_y^{ie}A_y^r + \lambda A_\pi^{ie}A_\pi^r\right), \qquad \xi_u := 2\left(A_y^{ie}A_y^u + \lambda A_\pi^{ie}A_\pi^u\right), \qquad \xi_i := 2\left((A_y^{ie})^2 + \lambda(A_\pi^{ie})^2\right),$$

so that

$$G_{\alpha,0,\gamma}(s_0) = \xi_r r_0^n + \xi_u u_0 + \xi_i(\alpha\mu_{\tilde{s}} + \gamma\theta_{\alpha,0}),$$

because the mean-term in G (the part proportional to $\bar{r}^n$ and $\mathbb{E}[i_1^e]$) vanishes at $(\bar{r}^n, b) = (0,0)$.

Let ρ_r, ρ_u be the AR(1) coefficients of (r_t^n, u_t). Note that

$$\mu_{\tilde{s}} = \tilde{\omega}_r\rho_r r_0^n + \tilde{\omega}_u\rho_u u_0, \qquad \text{Var}(\mu_{\tilde{s}}) = \tilde{\omega}_r^2\rho_r^2\sigma_r^2 + \tilde{\omega}_u^2\rho_u^2\sigma_u^2,$$

As in the statement of the Corollary, denote

$$\sigma_{r\theta} := \mathrm{Cov}(r_0^n, \theta_{\alpha,0}), \quad \sigma_{u\theta} := \mathrm{Cov}(u_0, \theta_{\alpha,0}), \quad \sigma_\theta^2 := \mathrm{Var}(\theta_{\alpha,0}).$$

Recall from Proposition 4, the α-FOC is given by

$$\mathbb{E}\left[G_{\alpha,0,\gamma}\left(\mu_{\tilde{s}} + \gamma\frac{\partial\theta_{\alpha,0}}{\partial\alpha}\right)\right] + 2W_1\left(\mathbb{E}[\mu_{\tilde{s}}]\,\mathbb{E}[(|z_{\alpha,0}| - \gamma)_+ \operatorname{sign} z_{\alpha,0}] + \mathrm{Cov}(\tilde{s}_1, z_{\alpha,0})p(z_{\alpha,0}, \gamma)\right) = 0.$$

Since $\mathbb{E}[\mu_{\tilde{s}}] = 0$, the mean term in the $t = 1$ part drops. Further, because $z_{\alpha,0}$ is mean zero, it's straightforward to get $p(z_{\alpha,0}, \gamma) = 2\left(1 - \Phi(\gamma/\sigma_z)\right)$. Thus, the $t = 1$ terms reduce to

$$4W_1\,\mathrm{Cov}(\tilde{s}_1, z_{\alpha,0})\left(1 - \Phi(\gamma/\sigma_z)\right).$$

Next, expanding the $t = 0$ part and noting that $\mathbb{E}[r_0^n] = \mathbb{E}[u_0] = \mathbb{E}[\theta_{\alpha,0}] = \mathbb{E}[\mu_{\tilde{s}}] = 0$, one gets

$$\begin{aligned}
\mathbb{E}\left[G_{\alpha,0,\gamma}\mu_{\tilde{s}}\right] &= \mathbb{E}\left[(\xi_r r_0^n + \xi_u u_0 + \xi_i \alpha\mu_{\tilde{s}} + \xi_i\gamma\theta_{\alpha,0})\mu_{\tilde{s}}\right] \\
&= \xi_r\,\mathrm{Cov}(r_0^n, \mu_{\tilde{s}}) + \xi_u\,\mathrm{Cov}(u_0, \mu_{\tilde{s}}) + \xi_i\alpha\,\mathrm{Var}(\mu_{\tilde{s}}) + \xi_i\gamma\,\mathrm{Cov}(\theta_{\alpha,0}, \mu_{\tilde{s}}) \\
&= \xi_i\,\mathrm{Var}(\mu_{\tilde{s}})(\alpha - \tilde{\alpha}_0) + \xi_i\gamma\left(\tilde{\omega}_r\rho_r\sigma_{r\theta} + \tilde{\omega}_u\rho_u\sigma_{u\theta}\right),
\end{aligned}$$

where

$$\tilde{\alpha}_0 := -\frac{\xi_r\tilde{\omega}_r\rho_r\sigma_r^2 + \xi_u\tilde{\omega}_u\rho_u\sigma_u^2}{\xi_i\,\mathrm{Var}(\mu_{\tilde{s}})}.$$

Finally, define

$$h(\alpha - a_{ps}; \gamma) := \mathbb{E}\left[(\xi_r r_0^n + \xi_u u_0 + \xi_i\alpha\mu_{\tilde{s}} + \xi_i\gamma\theta_{\alpha,0})\,\frac{\partial\theta_{\alpha,0}}{\partial\alpha}\right].$$

Collecting terms, the α–FOC becomes

$$\xi_i\,\mathrm{Var}(\mu_{\tilde{s}})(\alpha - \tilde{\alpha}_0) + \xi_i\gamma\left(\tilde{\omega}_r\rho_r\sigma_{r\theta} + \tilde{\omega}_u\rho_u\sigma_{u\theta}\right) + \gamma h(\alpha - a_{ps}; \gamma) + 4W_1\,\mathrm{Cov}(\tilde{s}_1, z_{\alpha,0})\left(1 - \Phi(\gamma/\sigma_z)\right) = 0.$$

Note that $z_{\alpha,0} = (\alpha - a_{cb})\tilde{s}_1 - a_{cb}^0(s_0, \hat{s}_1)$ under $b = 0$, so

$$\mathrm{Cov}(\tilde{s}_1, z_{\alpha,0}) = (\alpha - a_{cb})\,\mathrm{Var}(\tilde{s}_1) - \mathrm{Cov}\Big(\tilde{s}_1, a_{cb}^0(s_0, \hat{s}_1)\Big) = (\alpha - \tilde{\alpha}_1)\,\mathrm{Var}(\tilde{s}_1),$$

with $\tilde{\alpha}_1$ as defined in the statement.

(iii) *The γ–FOC.* Recall the γ–FOC in Proposition 4 is given by

$$\mathbb{E}\left[G_{\alpha,0,\gamma}\theta_{\alpha,0}\right] - 2W_1\,\mathbb{E}\left[(|z_{\alpha,0}| - \gamma)_+\right] \geq 0,$$

with equality unless $\gamma = 0$. Expanding the $t = 0$ term and using that all primitive random variables are mean zero, we get

$$\begin{aligned}\mathbb{E}\left[G_{\alpha,0,\gamma}\theta_{\alpha,0}\right] &= \mathbb{E}\left[(\xi_r r_0^n + \xi_u u_0 + \xi_i \alpha \mu_{\tilde{s}} + \xi_i \gamma \theta_{\alpha,0})\theta_{\alpha,0}\right] \\ &= (\xi_r + \alpha \xi_i \tilde{\omega}_r \rho_r)\sigma_{r\theta} + (\xi_u + \alpha \xi_i \tilde{\omega}_u \rho_u)\sigma_{u\theta} + \xi_i \sigma_\theta^2 \gamma,\end{aligned}$$

where the second equality follows from $\mathbb{E}[\mu_{\tilde{s}}\theta_{\alpha,0}] = \tilde{\omega}_r\rho_r\sigma_{r\theta} + \tilde{\omega}_u\rho_u\sigma_{u\theta}$. For the $t = 1$ term, note that $z_{\alpha,0} \sim \mathcal{N}(0, \sigma_z^2)$ and thus

$$\mathbb{E}\left[(|z_{\alpha,0}| - \gamma)_+\right] = 2\int_\gamma^\infty (z - \gamma)\frac{1}{\sigma_z}\phi\left(\frac{z}{\sigma_z}\right)dz = 2\left(\sigma_z\phi(\gamma/\sigma_z) - \gamma\left(1 - \Phi(\gamma/\sigma_z)\right)\right).$$

Hence the γ–FOC in the statement follows.

B.6 Proof of Proposition 4

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 4 and Ψ as defined in Proposition 1 evaluated at i_1^e as defined in Lemma 3. Then, $\mathcal{J}$ is continuously differentiable on $\mathcal{D}$ and the derivatives are given by

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

where

$$G_{\alpha,b,\gamma}(s_0) := 2A_y^{i^e}\left(A_y^r r_0^n + A_y^u u_0 + A_y^{i^e} i_1^e(s_0;\alpha,b,\gamma)\right) + 2\lambda_{cb}A_\pi^{i^e}\left(A_\pi^r r_0^n + A_\pi^u u_0 + A_\pi^{i^e} i_1^e(s_0;\alpha,b,\gamma)\right)$$
$$- \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \bar{r}^n + \delta_{i^e}\,\mathbb{E}[i_1^e(s_0;\alpha,b,\gamma)]\right),$$

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

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

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 4.

(i) *Objective Function.* The fact that $t = 0$ losses are given by $\Psi(\alpha, b, \gamma)$ follows directly from Proposition 1 and Lemma 2. To get the expression for $t = 1$ losses, note that by Lemma 1, these losses are given by

$$W_1\,\mathbb{E}\left[(i_1^* - i_1^{\text{unc}})^2\right] = W_1\,\mathbb{E}\left[(|i_1^{\text{unc}} - (\alpha\tilde{s}_1 + b)| - \gamma)_+^2\right]$$
$$= W_1\,\mathbb{E}\left[\left(|(\alpha - a_{cb})\tilde{s}_1 + (b - a_0^{(1)})| - \gamma\right)_+^2\right],$$

where the last equality follows from the form of the conditional expectation in Lemma 4.

(ii) *Existence of a solution.* I will prove that $\mathcal{J}$ is coercive. Because the feasible set $\mathcal{X} := \mathbb{R}^2 \times \mathbb{R}_+$ is closed and non-empty and $\mathcal{J}$ is clearly continuous, the sublevel sets $\{(\alpha, b, \gamma) \in \mathcal{X} : \mathcal{J}(\alpha, b, \gamma) \leq t\}$ are closed for all t. It is a known result that a function is coercive if and only if its sublevel sets are bounded. Thus, if $\mathcal{J}$ is coercive, its sublevel sets are compact and by Weierstrass' Theorem a minimizer exists in the feasible set. To verify that $\mathcal{J}$ is coercive, note that if $|\alpha|$ or $|b| \to +\infty$, the $t = 0$ term diverges (Ψ is quadratic in i_1^e, $i_1^e = \alpha\mu_{\tilde{s}} + b + \gamma\theta_{\alpha,b}$ and $\theta_{\alpha,b} \in [-1, 1]$) while the $t = 1$ term diverges too (as $|z| - \gamma > 0$ has positive probability for every finite γ and z is linear in (α, b)). On the other hand, if $\gamma \to +\infty$, the $t = 1$ term vanishes but the $t = 0$ term diverges (i_1^e is a linear function of $\gamma\theta_{\alpha,b}$ and Ψ is quadratic in i_1^e). Thus $\mathcal{J}$ is coercive and hence attains a minimum.

(iii) *Differentiability and first-order conditions.* The continuous differentiability of $\mathcal{J}$ on $\mathcal{D}$ and the continuous extension of the derivatives to the $\gamma = 0$ case follow from Lemma 5. As mentioned above, the feasible set $\mathcal{X}$ is closed. Further, it is clearly convex. Thus, the first-order conditions are necessary. The boundary condition for γ is the usual KKT sign condition. The expression for the first-order conditions also follow from Lemma 5.

□

B.7 Proof of Proposition 5

Fix (i_0, α, b, γ) with $\gamma \geq 0$ and $\alpha \neq \alpha^\star := 1/A_\pi^{i^e}$. The constraint $\alpha \neq \alpha^\star$ is imposed because when $\alpha = \alpha^\star$ the reduced-form fixed point that pins down inflation becomes ill-defined as the coefficient on π_0 vanishes. Indeed, substituting the equilibrium expectation $i_1^e = \alpha\pi_0 + b + \gamma\theta_0$ into

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

yields

$$(1 - \alpha A_\pi^{i^e})\,\pi_0 = A_\pi^r r_0^n + A_\pi^u u_0 - \sigma\kappa i_0 + A_\pi^{i^e}(b + \gamma\theta_0).$$

Hence, when $\alpha = \alpha^\star = 1/A_\pi^{i^e}$ the left-hand side is identically zero, and generically the right-hand side is nonzero, so no $t = 0$ equilibrium exists. Only on the knife-edge set of parameters/shocks for which the right-hand side is exactly zero would π_0 be indeterminate. We therefore exclude $\alpha^\star$.

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

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

At $t = 0$, an equilibrium is a pair (π_0, θ_0) such that $\theta_0 \in [-1, 1]$ and

$$i_1^e = \alpha\pi_0 + b + \gamma\theta_0$$

as the argument of Lemma 2 applies analogously, substituting the exogenous signal by π_0. Substituting and solving yields that, for any $\theta \in [-1, 1]$,

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

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

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

$$D(\theta; s_0) := \alpha\, \pi_0(\theta) + b - \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0].$$

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

$$|D(\theta; s_0) + \gamma| > |D(\theta; s_0) - \gamma| \quad \Longleftrightarrow \quad D(\theta; s_0) > 0,$$

and the lower end is strictly worse if and only if

$$|D(\theta; s_0) - \gamma| > |D(\theta; s_0) + \gamma| \quad \Longleftrightarrow \quad D(\theta; s_0) < 0.$$

Agents are indifferent between the two ends if and only if $D(\theta; s_0) = 0$.

Since $\pi_0(\theta)$ is affine in θ, so is $D(\theta; s_0)$. In particular, one can rewrite the indifference condition $D(\theta; s_0) = 0$ as

$$\theta = \frac{T(s_0; \alpha, b)}{\gamma},$$

where $T(s_0; \alpha, b)$ is as defined in the statement of the proposition. Therefore, the unique θ that supports indifference satisfies $\theta \in [-1, 1]$ if and only if $|T(s_0; \alpha, b)| \le \gamma$; equivalently, the "intermediate" equilibrium value is $\theta = T/\gamma$ whenever $|T| \le \gamma$.

Next, consider the strict-end equilibria. Evaluating $D(\theta; s_0)$ at $\theta = +1$ and $\theta = -1$ 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 $\alpha < \alpha^\star$. In this region $\alpha A_\pi^{i^e}/(1 - \alpha A_\pi^{i^e}) < 0$, and the strict best-response conditions imply

$$\theta_0 = +1 \text{ is an equilibrium} \iff T(s_0; \alpha, b) > \gamma,$$
$$\theta_0 = -1 \text{ is an equilibrium} \iff T(s_0; \alpha, b) < -\gamma.$$

When $|T(s_0; \alpha, b)| \le \gamma$, neither extreme is strictly worse, and the unique equilibrium is the indifference value $\theta_0 = T(s_0; \alpha, b)/\gamma$, with the interpretation that a share of the market picks the upper extreme and the rest pick the lower extreme. Hence $\mathcal{E}_{\alpha,b,\gamma}(s_0)$ is a singleton and equals the expression in the statement.

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

$$\theta_0 = -1 \text{ is an equilibrium } \iff T(s_0;\alpha,b) > \gamma,$$
$$\theta_0 = +1 \text{ is an equilibrium } \iff T(s_0;\alpha,b) < -\gamma.$$

Moreover, when $|T(s_0;\alpha,b)| \le \gamma$, the indifference value $\theta_0 = T(s_0;\alpha,b)/\gamma$ is also feasible and therefore 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

$$D(\theta; s_0) = b - \mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$$

is independent of θ. Hence $\theta_0 = +1$ is the unique equilibrium if $b > \mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$, $\theta_0 = -1$ is the unique equilibrium if $b < \mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0]$, and if equality holds then any $\theta_0 \in [-1,1]$ is an equilibrium.

Collecting the three cases yields the characterization of $\mathcal{E}_{\alpha,b,\gamma}(s_0)$ in the statement of the proposition. Given any $\theta_0 \in \mathcal{E}_{\alpha,b,\gamma}(s_0)$, equilibrium inflation is then given by $\pi_0(\theta_0)$ above.

B.8 Proof of Proposition 6

(i) *Objective.* Fix (i_0, α, b, γ) with $\gamma \ge 0$ and $\alpha \ne \alpha^\star$. Under the selection convention following Proposition 5, for each s_0 define

$$i_1^e(i_0,\alpha,b,\gamma;s_0) := \alpha\pi_0(i_0,\alpha,b,\gamma;s_0) + b + \gamma\vartheta_{\alpha,b,\gamma}(s_0),$$

where $\pi_0(\cdot)$ is the associated $t=0$ equilibrium inflation under the selection and $\vartheta_{\alpha,b,\gamma}(\cdot)$ 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

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

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 the $t = 1$ loss equals

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

Define

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

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 objective is not coercive in γ because $\tilde{\Psi}(i_0, \alpha, b, \gamma)$ can converge to a finite limit as $\gamma \to \infty$: under the selection, $b+\gamma\vartheta_{\alpha,b,\gamma}(s_0) \to b+T(s_0; \alpha, b)$ in probability, and the $t = 1$ term vanishes. We compactify γ by allowing $\gamma \in [0, \infty]$ and endowing the (positive) extended reals with the usual order topology. Correspondingly, we define the extension at $\gamma = \infty$ as

$$\widetilde{\mathcal{J}}(i_0, \alpha, b, \infty) := \lim_{\gamma \to \infty} \tilde{\mathcal{J}}(i_0, \alpha, b, \gamma)$$

whenever the limit exists.

Fix (i_0, α, b) with $\alpha \neq \alpha^\star$ and $\alpha \neq 0$. By Proposition 5, $T(s_0; \alpha, b)$ is affine in s_0 and hence Gaussian. For each $\gamma > 0$ let $A_\gamma := \{|T| \leq \gamma\}$ and $B_\gamma := \{|T| > \gamma\}$. On A_γ the selection implies $\vartheta_{\alpha,b,\gamma}(s_0) = T(s_0; \alpha, b)/\gamma$, hence

$$b + \gamma\vartheta_{\alpha,b,\gamma}(s_0) = b + T(s_0; \alpha, b)$$

on A_γ. On B_γ, $\vartheta_{\alpha,b,\gamma}(s_0) \in \{-1, +1\}$, so $|b+\gamma\vartheta_{\alpha,b,\gamma}(s_0)| \leq |b|+\gamma$. Since T is Gaussian, $\mathbb{P}(B_\gamma) \to 0$ as $\gamma \to \infty$.

Fix (i_0, α) and let $g(s_0, x)$ denote the integrand of $\tilde{\Psi}(i_0, \alpha, b, \gamma)$ when i_1^e is evaluated at $x = b + \gamma\vartheta_{\alpha,b,\gamma}(s_0)$, so that $\tilde{\Psi}(i_0, \alpha, b, \gamma) = \mathbb{E}[g(s_0, b + \gamma\vartheta_{\alpha,b,\gamma}(s_0))]$. Because g is a quadratic polynomial in (s_0, x), there exists $C < \infty$ such that

$$|g(s_0, x)| \leq C\left(1 + |s_0|^2 + x^2\right) \qquad \text{for all } (s_0, x).$$

Split the expectation over A_γ and B_γ. On A_γ,

$$g\left(s_0, b + \gamma\vartheta_{\alpha,b,\gamma}(s_0)\right) = g\left(s_0, b + T(s_0; \alpha, b)\right),$$

so $\mathbb{E}[g(\cdot)\mathbf{1}_{A_\gamma}] \to \mathbb{E}[g(s_0, b + T(s_0; \alpha, b))]$ by dominated convergence, using that $\mathbb{E}[1 +$

$|s_0|^2 + T(s_0;\alpha,b)^2] < \infty$. On B_γ, using the bound above and $|b+\gamma\vartheta| \le |b| + \gamma$,

$$|g(s_0, b+\gamma\vartheta_{\alpha,b,\gamma}(s_0))|\mathbf{1}_{B_\gamma} \le C\Big(1+|s_0|^2 + (|b|+\gamma)^2\Big)\mathbf{1}_{B_\gamma}.$$

Since s_0 is Gaussian, $\mathbb{E}[(1+|s_0|^2)^2] < \infty$. Hence, by Cauchy–Schwarz,

$$\mathbb{E}\left[\Big(1+|s_0|^2+(|b|+\gamma)^2\Big)\mathbf{1}_{B_\gamma}\right] \le \left(\mathbb{E}\Big[\Big(1+|s_0|^2+(|b|+\gamma)^2\Big)^2\Big]\right)^{1/2} \mathbb{P}(B_\gamma)^{1/2} \longrightarrow 0$$

as $\gamma \to \infty$ because $\mathbb{P}(B_\gamma)$ decays exponentially in γ. Therefore,

$$\tilde{\Psi}(i_0,\alpha,b,\gamma) \;\longrightarrow\; \tilde{\Psi}(i_0,\alpha,b,\infty) := \mathbb{E}\Big[g\Big(s_0, b+T(s_0;\alpha,b)\Big)\Big].$$

For the $t=1$ term, note that $z_{\alpha,b,\gamma}(s_0)$ is affine in s_0 on each of the regions A_γ and B_γ and has finite second moment. Moreover $(|z|-\gamma)_+^2 \le z^2\mathbf{1}_{\{|z|>\gamma\}}$ and $\mathbf{1}_{\{|z|>\gamma\}} \to 0$ pointwise as $\gamma \to \infty$. Since $\mathbb{E}[z_{\alpha,b,\gamma}(s_0)^2] < \infty$ for each fixed (i_0,α,b), dominated convergence yields

$$\mathbb{E}\left[(|z_{\alpha,b,\gamma}|-\gamma)_+^2\right] \longrightarrow 0.$$

Hence $\tilde{\mathcal{J}}(i_0,\alpha,b,\gamma)$ admits a finite limit as $\gamma\to\infty$, and we set $\widetilde{\mathcal{J}}(i_0,\alpha,b,\infty)$ equal to that limit.[30]

For the α-dimension, reparametrize by $m := \alpha/(1-\alpha A_\pi^{i^e})$, a bijection $\mathbb{R}\setminus\{\alpha^\star\} \to \mathbb{R}$ with inverse $\alpha = m/(1+A_\pi^{i^e} m)$. Under Proposition 5 and the Selection Convention, the maps $(i_0,m,b,\gamma,s_0) \mapsto \pi_0(i_0,m,b,\gamma;s_0)$ and $(m,b,\gamma,s_0)\mapsto \vartheta_{\alpha,b,\gamma}(s_0)$ are measurable and piecewise affine, so the extended objective can be written as

$$\widetilde{\mathcal{J}}(i_0,\alpha,b,\gamma) = \widehat{\mathcal{J}}(i_0,m,b,\gamma) \qquad \text{for } \gamma\in[0,\infty],$$

where $\widehat{\mathcal{J}}$ is jointly lower semicontinuous on $\mathbb{R}\times\mathbb{R}\times\mathbb{R}\times[0,\infty]$. Note that the inverse mapping is not defined at $m=-1/A_\pi^{i^e}$, which corresponds to $|\alpha|=\infty$. At that value, $1+mA_\pi^{i^e}=0$ so i_1^e remains finite and $\widehat{\mathcal{J}}$ extends lower semicontinuously. We therefore extend the feasible set to include the boundary $|\alpha|=\infty$.

Because i_1^e enters $\tilde{\Psi}$ quadratically with $\mathrm{Var}(r_0^n), \mathrm{Var}(u_0) > 0$ and the $t=1$ term is nonnegative, we have $\widehat{\mathcal{J}}(i_0,m,b,\gamma)\to\infty$ as $|i_0|\to\infty$, $|m|\to\infty$, or $|b|\to\infty$, uniformly in $\gamma\in[0,\infty]$. Hence the sublevel sets $\{(i_0,m,b,\gamma)\in\mathbb{R}^3\times[0,\infty] : \widehat{\mathcal{J}}\le c\}$ are bounded.

[30]When $\alpha=0$, $T(s_0;\alpha,b)$ is not defined (Proposition 5), but then $\vartheta_{0,b,\gamma}$ is independent of γ and $i_1^e = b+\gamma\vartheta_{0,b,\gamma}$ enters $\tilde{\Psi}$ quadratically, so $\tilde{\mathcal{J}}(i_0,0,b,\gamma)\to+\infty$ as $\gamma\to\infty$. Thus, the minimizer cannot feature $\gamma=\infty$ in that case.

Since $\widehat{\mathcal{J}}$ is jointly lower semicontinuous, they are also closed, and therefore compact. By the extreme value theorem, $\widehat{\mathcal{J}}$ attains a minimum at some $(i_0^*, m^*, b^*, \gamma^*) \in \mathbb{R}^3 \times [0, \infty]$. Mapping back to α yields a minimizer $(i_0^*, \alpha^*, b^*, \gamma^*) \in \mathbb{R} \times \overline{\mathbb{R}} \setminus \{\alpha^\star\} \times \mathbb{R} \times \overline{\mathbb{R}}_+$.

(iii) *Differentiability and first–order conditions.* On $\mathbb{R} \times (\mathbb{R} \setminus \{\alpha^\star\}) \times \mathbb{R} \times (0, \infty)$, both $\vartheta_{\alpha,b,\gamma}(s_0)$ and $\pi_0(i_0, \alpha, b, \gamma; s_0)$ are piecewise C^1 in (i_0, α, b, γ). Non-differentiabilities arise only at the kink events

$$|T(s_0; \alpha, b)| = \gamma \qquad \text{and} \qquad |z_{\alpha,b,\gamma}(s_0)| = \gamma,$$

where the selection switches regime and the projection term becomes active. Under the maintained Gaussian specification for s_0, these events have probability zero whenever $T(s_0; \alpha, b)$ and $z_{\alpha,b,\gamma}(s_0)$ are nondegenerate (a condition that fails only on knife-edge parameterizations). Thus, for generic parameters we may differentiate under the expectation by dominated convergence, exactly as in the proof of Lemma 5 (the only difference being that here the integrand is piecewise smooth rather than globally smooth). The partial derivatives extend continuously to $\gamma = 0$ when interpreted as right derivatives at the boundary.

Define

$$\tilde{G}_{\alpha,b,\gamma}(s_0) := 2A_y^{ie}\left(A_y^r r_0^n + A_y^u u_0 - \sigma i_0 + A_y^{ie} i_1^e\right) + 2\lambda_{cb} A_\pi^{ie}\left(A_\pi^r r_0^n + A_\pi^u u_0 - \sigma\kappa i_0 + A_\pi^{ie} i_1^e\right).$$

By the chain rule, for $j \in \{\alpha, b, \gamma\}$,

$$\frac{\partial \tilde{\Psi}}{\partial j} = \mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma} \frac{\partial i_1^e}{\partial j}\right], \qquad \frac{\partial \tilde{\Psi}}{\partial i_0} = \mathbb{E}\left[\tilde{G}_{\alpha,b,\gamma} \frac{\partial i_1^e}{\partial i_0}\right] - 2\sigma\, \mathbb{E}\left[y_0 + \lambda_{cb}\kappa\pi_0\right].$$

From the equilibrium inflation equation,

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

we obtain

$$\frac{\partial \pi_0}{\partial i_0} = \frac{-\sigma\kappa + A_\pi^{ie}\gamma \frac{\partial \vartheta_{\alpha,b,\gamma}}{\partial i_0}}{1 - \alpha A_\pi^{ie}}, \qquad \frac{\partial \pi_0}{\partial \alpha} = \frac{A_\pi^{ie}}{1 - \alpha A_\pi^{ie}}\left(\pi_0 + \gamma \frac{\partial \vartheta_{\alpha,b,\gamma}}{\partial \alpha}\right),$$

$$\frac{\partial \pi_0}{\partial b} = \frac{A_\pi^{ie}}{1 - \alpha A_\pi^{ie}}\left(1 + \gamma \frac{\partial \vartheta_{\alpha,b,\gamma}}{\partial b}\right), \qquad \frac{\partial \pi_0}{\partial \gamma} = \frac{A_\pi^{ie}}{1 - \alpha A_\pi^{ie}}\left(\vartheta_{\alpha,b,\gamma} + \gamma \frac{\partial \vartheta_{\alpha,b,\gamma}}{\partial \gamma}\right).$$

Substituting into $i_1^e = \alpha\pi_0 + b + \gamma\vartheta_{\alpha,b,\gamma}$ gives

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

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

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

and

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

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

$$\frac{\partial z_{\alpha,b,\gamma}}{\partial \alpha} = \pi_0 + \alpha\frac{\partial \pi_0}{\partial \alpha},\quad \frac{\partial z_{\alpha,b,\gamma}}{\partial b} = 1 + \alpha\frac{\partial \pi_0}{\partial b},\quad \frac{\partial z_{\alpha,b,\gamma}}{\partial \gamma} = \alpha\frac{\partial \pi_0}{\partial \gamma},\quad \frac{\partial z_{\alpha,b,\gamma}}{\partial i_0} = \alpha\frac{\partial \pi_0}{\partial i_0}.$$

Substituting the partial derivatives above into the chain-rule expressions yields the stated first-order conditions. As stated in the proposition, the condition for γ is the complementary slackness condition associated with $\gamma \geq 0$, and the derivatives extend to $\gamma = 0$ as right derivatives.

C Additional Results

C.1 Second-Order Conditions

This section provides easily checkable conditions that guarantee candidates that satisfy the first-order conditions in Proposition 4 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.

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

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

with Ψ *and all constants as in Propositions 1 and 4, 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*

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

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* $p, q \in \{\alpha, b, \gamma\}$

$$\frac{\partial^2\Psi}{\partial p\partial q} = \xi_i \, \mathbb{E}\left[\frac{\partial I}{\partial p}\frac{\partial I}{\partial q}\right] - \frac{2\delta_{i^e}^2}{1 + \lambda_{cb}\kappa^2} \, \mathbb{E}\left[\frac{\partial I}{\partial p}\right] \mathbb{E}\left[\frac{\partial I}{\partial q}\right] + \mathbb{E}\left[G_{\alpha,b,\gamma}(s_0)\frac{\partial^2 I}{\partial p\partial q}\right], \tag{5}$$

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

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

with

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

and

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

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. Then, for each fixed (α, b) with $\alpha \neq a_{ps}$, the right second derivative with respect to γ exists and equals the limit of the interior second derivatives, i.e.,*

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

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

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

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 7. *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*

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

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

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

Then the following are equivalent:

1. $(\alpha^*, b^*, \gamma^*)$ *is an isolated strict local minimizer of* $\mathcal{J}$,

2. $H \succ 0$,

3. $c_\gamma > 0$ *and* $S \succ 0$.

Then, $c_\gamma > 0$ *and thus, for any parameter* $\theta \in \mathbb{R}$, $S \succ 0$ *implies*

$$\frac{\partial}{\partial\theta}\begin{pmatrix}\alpha^*\\ b^*\\ \gamma^*\end{pmatrix} = -H^{-1}\frac{\partial}{\partial\theta}\nabla_{(\alpha,b,\gamma)}\mathcal{J}\Big|_{(\alpha^*,b^*,\gamma^*)}.$$

II. ***Boundary*** **(**$\gamma^* = 0$**).** *Assume the KKT conditions in Proposition 4 with* $\gamma^* = 0$ *hold. Let*

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

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* $(\alpha^*, b^*, 0)$ *and* $R \succ 0$, *then* $(\alpha^*, b^*, 0)$ *is an isolated strict local minimizer. Moreover, for any parameter* $\theta \in \mathbb{R}$,

$$\frac{\partial}{\partial\theta}\begin{pmatrix}\alpha^*\\ b^*\\ \gamma^*\end{pmatrix} = -\begin{pmatrix}R^{-1}\frac{\partial}{\partial\theta}\nabla_{(\alpha,b)}\mathcal{J}\\ 0\end{pmatrix}\Bigg|_{(\alpha^*,b^*,0)},$$

all evaluated at $(\alpha^*, b^*, 0)$.

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

$$R \succ 0 \qquad \text{and} \qquad c_\gamma > (H^+_{x\gamma})^\top R^{-1} H^+_{x\gamma}.$$

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 $A_\pi^{ie} - \kappa A_y^{ie} = \beta\sigma\kappa$, we have

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

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) and that the FOC map is invertible. The Implicit Function Theorem then yields the expression in the statement.

II. **Boundary.**
Let $d = (u, t)$ be any feasible direction at $(\alpha^*, b^*, 0)$ with $u \in \mathbb{R}^2$ and $t \in \mathbb{R}_+$, and define $\psi(s) := \mathcal{J}(\alpha^* + su_1, b^* + su_2, st)$. By Lemma 7, the right second derivative exists and equals

$$\psi''(0^+) = u^\top R u + 2tu^\top H_{x\gamma}^+ + t^2 c_\gamma.$$

(a) Under strict complementarity, $t = 0$ (γ is fixed at 0 locally). Thus, $\psi''(0^+) = u^\top R u > 0$ for all $u \neq 0$ if and only if$R \succ 0$, which yields a strict local minimum on the boundary. The reduced IFT follows by fixing $\gamma \equiv 0$ and applying it to $\nabla_{(\alpha,b)} \mathcal{J}(\alpha, b, 0) = 0$.

(b) In a degenerate corner, the feasible directions are $\{(u, t) : t \geq 0\}$. Using $u^\top R u \geq \lambda_{\min}(R) \|u\|^2$ and $2tu^\top H_{x\gamma}^+ \geq -2tL^+ \|u\|$, the minimum of ψ'' over $\|u\| \geq 0$ for fixed t equals $t^2(c_\gamma - (L^+)^2/m)$. Hence $\psi''(0^+) > 0$ for all feasible directions if and only if $\lambda_{\min}(R) > 0$ and $c_\gamma \lambda_{\min}(R) > (L^+)^2$. This condition is therefore necessary and sufficient for an isolated strict local minimizer of the constrained problem at the degenerate corner.

□

C.2 The Quality of Signals and Optimal Announcements

How do optimal announcements change with shifts in signal quality? The next proposition addresses this question[31] by characterizing the optimal response of each announcement dimension keeping the other fixed. The purpose is to clarify the workings of the model and organize the intuition for the numerical comparative statics in the endogenous signals model (Section 5). Figure 10 below graphically illustrates the content of the first two elements of the proposition.

[31] We focus on the $\overline{r}^n = 0$ case (Proposition 3), as this allows us to anchor our understanding on the intuitions just built in the previous discussion.

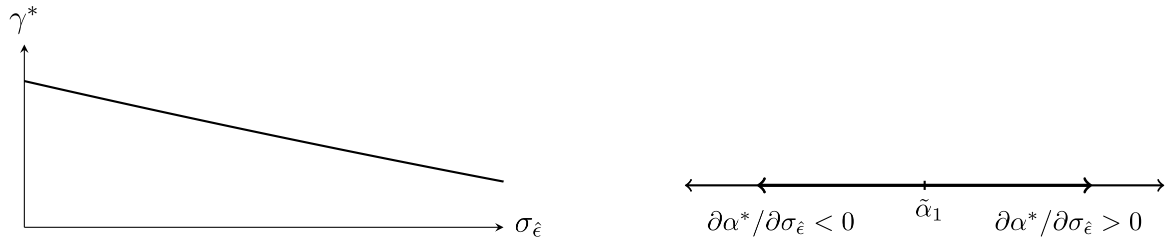

(a) Vagueness (fixing α) (b) Data dependence (fixing γ)

Figure 10: Optimal announcements as the non-conditionable noise increases.

Proposition 8. *Suppose $\overline{r}^n = 0$. Then, the following holds.*

(i) Fixing α, optimal vagueness γ^* is decreasing in the idiosyncratic noise of the non-conditionable signal $\sigma_{\hat{\epsilon}}$.

(ii) Fixing γ, optimal data dependence α^* is decreasing in the idiosyncratic noise of the non-conditionable signal $\sigma_{\hat{\epsilon}}$ if $\alpha^* < \tilde{\alpha}_1$ and increasing if $\alpha^* > \tilde{\alpha}_1$. If $\alpha^* = \tilde{\alpha}_1$, α^* is independent of $\sigma_{\hat{\epsilon}}$.

(iii) Fixing α, optimal vagueness γ^* is increasing in the idiosyncratic noise of the conditionable signal σ_{ϵ} if $\alpha^* > \overline{\alpha}$, where

$$\overline{\alpha} = \max\left\{a_{ps}, \tilde{\alpha}_1, \frac{1}{\tilde{\omega}_r}, \frac{\Gamma(\lambda_{ps})}{\tilde{\omega}_u}, -\frac{\xi_r}{\xi_i \tilde{\omega}_r \rho_r}, -\frac{\xi_u}{\xi_i \tilde{\omega}_u \rho_u}\right\}.$$

A noisier non-conditionable signal makes flexibility less valuable. The intuition for (i) is fairly straightforward. As the non-conditionable signal becomes less informative, the central bank optimally relies less on it to an extent that the overall volatility of deviations from the center of the band to the optimal policy rate at $t = 1$ decreases. This makes vagueness less valuable at $t = 1$. Put differently, the value of flexibility decreases as the information available to make use of that flexibility gets worse. Naturally, as the non-conditionable signal is not on the private sector's radar, and vagueness affects $t = 0$ only through these expectations, changes in $\sigma_{\hat{\epsilon}}$ do not change the $t = 0$ cost of vagueness.

A noisier non-conditionable signal pushes data dependence away from its $t = 1$ target value. The intuition behind (ii) is slightly subtler. First, note that the target $\tilde{\alpha}_1$ is independent of idiosyncratic noise of the non-conditionable signal, as it is simply the projection of the optimal policy rate on the conditionable signal. What does change is the relative weight on the $t = 1$ target. In particular, as $\sigma_{\hat{\epsilon}}$ increases, the probability of being constrained ex-post decreases for a given degree of vagueness, thus making $t = 1$ considerations less important for optimal data dependence. Consequently, as the marginal effect of data dependence on $t = 0$ is unchanged, the optimal α^* moves away from $\tilde{\alpha}_1$ as $\sigma_{\hat{\epsilon}}$ increases: if $\alpha^* < \tilde{\alpha}_1$, it is because $t = 0$ incentives are originally pulling α down relative to what is optimal purely from a $t = 1$ perspective, so the increase in $\sigma_{\hat{\epsilon}}$, by increasing the relative weight of $t = 0$ incentives, pushes α^* further down. The opposite happens if $\alpha^* > \tilde{\alpha}_1$.

The effect of a noisier conditionable signal on optimal vagueness: many moving parts. The mechanisms behind (iii) are more involved, as the noise on the conditionable signal shapes both $t = 1$ and $t = 0$ incentives, the latter through multiple channels. Indeed, as $\sigma_{\tilde{\epsilon}}$ increases, the volatility of deviations from the center of the band to the optimal policy rate at $t = 1$ changes for a given degree of vagueness, thus affecting the weight of $t = 1$ incentives on the determination of the optimal γ. If $\alpha > \tilde{\alpha}_1$, the volatility of these deviations increases with $\sigma_{\tilde{\epsilon}}$, thus making vagueness more valuable at $t = 1$ and pushing γ^* up. If $\alpha < \tilde{\alpha}_1$, on the other hand, the effect on the volatility of deviations is ambiguous. This is because two forces work in opposite directions. First, higher $\sigma_{\tilde{\epsilon}}$ raises $\text{Var}(\tilde{s}_1)$, which tends to increase the volatility of deviations. Second, as $\tilde{\alpha}_1$ decreases with $\sigma_{\tilde{\epsilon}}$, the gap $\alpha - \tilde{\alpha}_1$ (negative in this case) shrinks in magnitude, which tends to decrease the volatility. Usually the first force dominates, but for large negative gaps the second can more than offset it, making the overall effect negative. In this case, the weight of $t = 1$ incentives on the determination of γ^* decreases with $\sigma_{\tilde{\epsilon}}$, pushing γ^* down.

The effect on $t = 0$ incentives (i.e., on the marginal effect of vagueness on $t = 0$ output and inflation volatility) is threefold. This is because the effect of shocks on (y_0, π_0), and thus on the marginal effect of vaguer announcements, has three components: 1) a direct effect (through the contemporaneous realization of the shocks and the private sector's associated change in expectations of their $t = 1$ realizations); 2) an indirect effect through expectations of the mean-band for the policy rate; and 3) an indirect effect through the expectations of the worst-case selection for the policy rate.

The covariance between the direct effect of shocks and the marginal effect of a vaguer announcement has the opposite sign of $\sigma_{r\theta}$. To understand why, consider, for example, a positive natural rate shock $r_0^n > 0$ and a positive covariance $\sigma_{r\theta}$. While the positive rate shock has a positive direct effect on y_0 and π_0, it also increases the expected worst-case selection θ and thus the negative impact of marginally higher expected policy rates on inflation and output. The opposite happens for a negative covariance $\sigma_{r\theta}$. On the other hand, the covariance between the mean-band indirect effect and the marginal effect of a vaguer announcement has the same sign as the corresponding covariance ($\sigma_{r\theta}$ or $\sigma_{u\theta}$). To understand why, consider again a positive natural rate shock $r_0^n > 0$ and a positive covariance $\sigma_{r\theta}$. The natural rate shock increases the expected mean band for the policy rate, hence pushing contemporaneous inflation and output down, while increasing the negative marginal effect of a vaguer announcement by increasing θ. Whenever $\alpha > -\xi_r/(\xi_i \tilde{\omega}_r \rho_r)$, this indirect effect carries more weight than the direct one and thus the overall effect of the covariance between the marginal effect of a vaguer announcement and r_0^n is increasing in $\sigma_{r\theta}$ ($\omega_r \rho_r$ appear because these are the weights on r_0^n of the conditional mean of the signal, $\mu_{\tilde{s}}$, and thus scale the relevance of the indirect channel). An analogous reasoning applies to the cost-push shock u_0. Furthermore, as $\sigma_{\tilde{\epsilon}}$ increases, both $\sigma_{r\theta}$ and $\sigma_{u\theta}$ decrease if and only if $\alpha > 1/\tilde{\omega}_r$ and $\alpha > \Gamma(\lambda_{ps})/\tilde{\omega}_u$, respectively. Intuitively, large α make $t = 1$ policy so responsive to $\tilde{s}_1$ that natural rate shocks increase the private sector's beliefs that rates will be too high and thus that the worst-case rate is the upper bound of the interval, hence decreasing the marginal effect of $\sigma_{\tilde{\epsilon}}$ on the average worst-case selection (the higher θ is, the lower the effect of $\sigma_{\tilde{\epsilon}}$ on it). Thus, if α is large enough, the covariance terms decrease with $\sigma_{\tilde{\epsilon}}$, and this decreases the marginal effect of vaguer announcements on $t = 0$ losses, pushing γ^* up.

Finally, the second indirect effect is associated with σ_θ^2, the volatility of the worst-case selection. This volatility also depends on $\sigma_{\tilde{\epsilon}}$. The direct effect of an increase in $\sigma_{\tilde{\epsilon}}$ is naturally to decrease this volatility. This is because, as $\sigma_{\tilde{\epsilon}}$ increases, the private sector is less able to predict the worst-case rate based on the observed shocks s_0. In turn, this decreases the volatility of this selection from the central bank's perspective. However, in a similar vein as with $t = 1$ incentives, there is an indirect effect through the change in the private sector's target, a_{ps} that in principle can offset the direct effect. In particular, as $\sigma_{\tilde{\epsilon}}$ increases, a_{ps} decreases as the private sector's optimal weight on the signal decreases. If $\alpha > a_{ps}$ this increases the gap $\alpha - a_{ps}$, thus decreasing the private sector's ability to predict the worst-case selector and hence reinforcing the direct effect. If $\alpha < a_{ps}$, however, the gap decreases, thus increasing the private sector's ability to predict the worst-case selector and hence offsetting the direct effect. The overall effect of $\sigma_{\tilde{\epsilon}}$ on the volatility of the worst-case selector is thus ambiguous if $\alpha < a_{ps}$ and negative if $\alpha > a_{ps}$. Lower volatility of the worst-case selector, in turn, naturally decreases the marginal cost of vaguer announcements at $t = 0$, pushing γ^* up.

Thus, overall, if α is large enough all the components of $t = 0$ incentives and $t = 1$ incentives push γ^* up as $\sigma_{\tilde{\epsilon}}$ increases. When this is not the case, some effects go in different directions and it's not possible to sign the overall effect in general. Further, while it is not guaranteed that the inequality $\alpha > \alpha^\star$ will hold at the optimal announcement for all reasonable parameter configurations, the previous discussion allows to dissect and understand the forces at play when the conditionable signal becomes less informative.

On the effect of a noisier conditionable signal on optimal data dependence. What about the effect of $\sigma_{\tilde{\epsilon}}$ on α^*? The discussion above already highlights that the effects of $\sigma_{\tilde{\epsilon}}$ on optimal announcements are not straightforward. In the case of optimal data dependence, offsetting forces prevent a clean prediction in general. While the natural intuition that a noisier signal pushes down the central bank reliance on it holds at $t = 1$, $t = 0$ effects may go in the opposite direction. To see this, note that the $t = 1$ target is given by

$$\tilde{\alpha}_1 = \frac{\text{Cov}(i_1^{\text{FI}}(\lambda_{cb}), \tilde{s}_1)}{\text{Var}(\tilde{s}_1)}.$$

As $\sigma_{\tilde{\epsilon}}$ increases, $\text{Var}(\tilde{s}_1)$ rises while the covariance term in the numerator is unaffected. Thus, in line with the intuition above, the $t = 1$ target decreases. This by itself pushes α^* downward. Note, further, that an analogous reasoning shows that a_{cb} decreases, so ex-post the central bank would always like to put less weight on the conditionable signal as it becomes less informative as per the standard intuition.

However, the idiosyncratic noise on the conditionable signal also shapes $t = 0$ data dependence incentives in a similar way as in the case of optimal vagueness discussed above. In particular, the marginal effect of data dependence on $t = 0$ losses depends on the covariances between shocks and the worst-case selection $\sigma_{r\theta}, \sigma_{u\theta}$ and the covariance between the marginal effect of α on the worst-case selector and the marginal effect of rates on losses as reflected in the h term in Proposition 3. A small enough α makes the effect of $\sigma_{\tilde{\epsilon}}$ on $\sigma_{r\theta}, \sigma_{u\theta}$ positive (the central bank's baseline data dependence is so low that a more volatile signal amounts to a higher probability that the worst-case rates are too low according to the private sector, hence increasing the effect of $\sigma_{\tilde{\epsilon}}$ on this selection), thus pushing α^* down too in line with $t = 1$ incentives. However, the effect in h is much less straightforward to sign, as it pertains to the covariance between shocks and the worst-case selection with the effect of $\sigma_{\tilde{\epsilon}}$ on the marginal effect of α on this selection. This effect can go in either direction depending on the parameters and the level of α, hence potentially pushing α^* up as $\sigma_{\tilde{\epsilon}}$ increases, offsetting the other forces discussed. Nevertheless, this latter effect is usually weaker than the other two, so in general, at least for low levels of α, one would expect data dependence to decrease as $\sigma_{\tilde{\epsilon}}$ increases.

C.2.1 Proof of Proposition 8

We prove (i)–(iii) as partial (direct, single-FOC) effects, i.e., holding the other component of the optimal announcement fixed.

(i) *Behavior of γ^* as $\sigma_{\hat{\epsilon}}$ increases.* Recall that if $\overline{r}^n = 0$ we have $b^* = 0$ and the γ–FOC is

$$\underbrace{\xi_i \sigma_\theta^2 \gamma + 4W_1 \left[\gamma \left(1 - \Phi(\gamma/\sigma_z)\right) - \sigma_z \phi(\gamma/\sigma_z)\right] + \text{terms independent of } \sigma_{\hat{\epsilon}} \text{ and } \gamma}_{=:F(\gamma,\sigma_z)} = 0,$$

where ϕ, Φ are the standard normal pdf and cdf, and $\sigma_z^2 = \operatorname{Var}(z_{\alpha,0})$. Clearly, F depends on $\sigma_{\hat{\epsilon}}$ only through σ_z. We have that

$$\frac{\partial F}{\partial \gamma} = \xi_i \sigma_\theta^2 + 4W_1 \frac{\partial}{\partial \gamma} \left[\gamma(1 - \Phi(t)) - \sigma_z \phi(t)\right] = \xi_i \sigma_\theta^2 + 4W_1(1 - \Phi(t)) > 0,$$

and

$$\frac{\partial F}{\partial \sigma_z} = \frac{\partial}{\partial \sigma_z} \left[\gamma(1 - \Phi(t)) - \sigma_z \phi(t)\right] = -4W_1 \phi(t) < 0.$$

Thus, by the Implicit Function Theorem,

$$\frac{d\gamma^*}{d\sigma_z} = -\frac{\frac{\partial F}{\partial \sigma_z}}{\frac{\partial F}{\partial \gamma}} > 0.$$

Hence, $\sigma_{\hat{\epsilon}}$ affects the optimal γ only through σ_z and it does so positively. Thus, it only remains to verify how does σ_z change with $\sigma_{\hat{\epsilon}}$. Let $\pi_1 := \operatorname{Cov}(a_{cb}^0, \tilde{s}_1)/\operatorname{Var}(\tilde{s}_1)$ so that $a_{cb}^0 = \pi_1 \tilde{s}_1 + (a_{cb}^0 - \mathbb{E}[a_{cb}^0 \mid \tilde{s}_1])$ with the last term orthogonal to $\tilde{s}_1$. Then, we can write

$$z_{\alpha,0} = (\alpha - a_{cb} - \pi_1)\tilde{s}_1 - \left(a_{cb}^0 - \mathbb{E}[a_{cb}^0 \mid \tilde{s}_1]\right),$$

and the two summands are orthogonal. Hence, as because of the Gaussian structure the conditional variance of a_{cb}^0 does not depend on the realization of $\tilde{s}_1$, we have

$$\begin{aligned} \sigma_z^2 &= (\alpha - a_{cb} - \pi_1)^2 \operatorname{Var}(\tilde{s}_1) + \operatorname{Var}(a_{cb}^0 \mid \tilde{s}_1) \\ &= (\alpha - \tilde{\alpha}_1)^2 \operatorname{Var}(\tilde{s}_1) + \operatorname{Var}(a_{cb}^0 \mid \tilde{s}_1) \end{aligned}$$

where the second equality follows from the definition of π_1 and $\tilde{\alpha}_1$. As proved above, $\tilde{\alpha}_1$ does not move with changes in $\sigma_{\hat{\epsilon}}$ and naturally neither does $\operatorname{Var}(\tilde{s}_1)$. Thus, we only need to check how $\operatorname{Var}(a_{cb}^0 \mid \tilde{s}_1)$ changes with $\sigma_{\hat{\epsilon}}$. To do so, first recall there exists $(w_0, \hat{w})$ such that

$$a_{cb}^0 = w_0' s_0 + \hat{w}\hat{s}_1.$$

Projecting $\hat{s}_1$ onto $(s_0, \tilde{s}_1)$ yields

$$\hat{s}_1 = B_0' s_0 + \tilde{B}\tilde{s}_1 + \eta$$

where $\eta = \hat{s}_1 - \mathbb{E}[\hat{s}_1 \mid s_0, \tilde{s}_1] \perp (s_0, \tilde{s}_1)$ and clearly $(B_0, \tilde{B})$ does not depend on $\sigma_{\hat{\epsilon}}$. Substituting,

$$a^0_{cb} = (w_0' + \hat{w} B_0')s_0 + \hat{w}\tilde{B}\tilde{s}_1 + \hat{w}\eta.$$

Hence,

$$\operatorname{Var}(a^0_{cb} \mid \tilde{s}_1) = \operatorname{Var}\left((w_0 + \hat{w}B_0)s_0 \mid \tilde{s}_1\right) + \hat{w}^2 \operatorname{Var}(\eta).$$

By Frisch-Waugh-Lovell,

$$\hat{w} = \frac{\operatorname{Cov}\left(i_1^{\text{FI}}(\lambda_{cb}) - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1], \hat{s}_1 - \mathbb{E}[\hat{s}_1 \mid s_0, \tilde{s}_1]\right)}{\operatorname{Var}(\eta)}$$

and thus

$$\hat{w}^2 \operatorname{Var}(\eta) = \frac{\operatorname{Cov}\left(i_1^{\text{FI}}(\lambda_{cb}) - \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1], \hat{s}_1 - \mathbb{E}[\hat{s}_1 \mid s_0, \tilde{s}_1]\right)^2}{\operatorname{Var}(\eta)}.$$

The numerator clearly does not depend on $\sigma_{\hat{\epsilon}}$ (the first argument of the covariance does not depend on $\hat{\epsilon}$). On the other hand,

$$\operatorname{Var}(\eta) = \operatorname{Var}(\hat{s}_1) + \operatorname{Var}(\mathbb{E}[\hat{s}_1 \mid s_0, \tilde{s}_1]) - 2\operatorname{Cov}(\hat{s}_1, \mathbb{E}[\hat{s}_1 \mid s_0, \tilde{s}_1]).$$

The last two terms do not depend on $\sigma_{\hat{\epsilon}}$, while $\operatorname{Var}(\hat{s}_1)$ equals $\sigma^2_{\hat{\epsilon}}$ plus terms independent on $\sigma^2_{\hat{\epsilon}}$. Thus, $\operatorname{Var}(\eta)$ is increasing in $\sigma^2_{\hat{\epsilon}}$ and therefore $\hat{w}^2 \operatorname{Var}(\eta)$ is decreasing in $\sigma^2_{\hat{\epsilon}}$. Finally, note that $w_0' + \hat{w}B_0'$ is precisely the coefficient on s_0 of the projection of $i_1^{\text{FI}}(\lambda_{cb})$ onto $(s_0, \tilde{s}_1)$, which is independent of $\sigma^2_{\hat{\epsilon}}$. Hence, $\operatorname{Var}\left((w_0' + \hat{w}B_0')s_0 \mid \tilde{s}_1\right)$ does not depend on $\sigma^2_{\hat{\epsilon}}$.

Putting all the above together, we conclude that σ^2_z is decreasing in $\sigma^2_{\hat{\epsilon}}$.

Consequently, fixing α, γ^* is decreasing in $\sigma^2_{\hat{\epsilon}}$ as stated.

(ii) *Behavior of α^* as $\sigma_{\hat{\epsilon}}$ increases.* Recall that if $\bar{r}^n = 0$ we have $b^* = 0$ and the α–FOC is

$$\underbrace{4W_1(\alpha - \tilde{\alpha}_1)\operatorname{Var}(\tilde{s}_1)\left(1 - \Phi(\gamma/\sigma_z)\right) + \text{terms independent of } \sigma_{\hat{\epsilon}}}_{=:G(\alpha, \sigma_{\hat{\epsilon}})} = 0,$$

where $\tilde{\alpha}_1$ is defined in equation (4) above. We will prove that the target $\tilde{\alpha}_1$ does not depend on $\sigma_{\hat{\epsilon}}$. To do so, first let $m(s_0, \tilde{s}_1, \hat{s}_1) := \mathbb{E}[i_1^{\text{FI}}(\lambda_{cb}) \mid s_0, \tilde{s}_1, \hat{s}_1]$. Then, we can write

$$\tilde{\alpha}_1 = \frac{\operatorname{Cov}\left(m(s_0, \tilde{s}_1, \hat{s}_1), \tilde{s}_1\right)}{\operatorname{Var}(\tilde{s}_1)}.$$

Naturally, the residual $i_1^{\mathrm{FI}}(\lambda_{cb}) - m(s_0, \tilde{s}_1, \hat{s}_1)$ is orthogonal to any measurable function of $(s_0, \tilde{s}_1, \hat{s}_1)$, hence in particular to $\tilde{s}_1$. Thus,

$$\mathrm{Cov}\left(m(s_0, \tilde{s}_1, \hat{s}_1), \tilde{s}_1\right) = \mathrm{Cov}(i_1^{\mathrm{FI}}(\lambda_{cb}), \tilde{s}_1)$$

and then

$$\tilde{\alpha}_1 = \frac{\mathrm{Cov}(i_1^{\mathrm{FI}}(\lambda_{cb}), \tilde{s}_1)}{\mathrm{Var}(\tilde{s}_1)}.$$

Crucially, this depends only on the joint law of $(i_1^{\mathrm{FI}}(\lambda_{cb}), \tilde{s}_1)$. Increasing the noise in $\hat{s}_1$ leaves the joint law of $(i_1^{\mathrm{FI}}(\lambda_{cb}), \tilde{s}_1)$ unchanged, so $\tilde{\alpha}_1$ does not depend on $\sigma_{\hat{\epsilon}}$ and thus neither does $\mathrm{Cov}(\tilde{s}_1, z_{\alpha,0})$. Thus, the only way $\sigma_{\hat{\epsilon}}$ affects the α–FOC is through the $\Phi(\gamma/\sigma_z)$ term. In particular,

$$\frac{\partial G}{\partial \sigma_{\hat{\epsilon}}} = \frac{4W_1 \gamma \phi(\gamma/\sigma_z)\,\mathrm{Var}(\tilde{s}_1)}{\sigma_z^2}(\alpha - \tilde{\alpha}_1)\frac{\partial \sigma_z}{\partial \sigma_{\hat{\epsilon}}}.$$

As shown in (i) above, σ_z is decreasing in $\sigma_{\hat{\epsilon}}$. Consequently,

$$\frac{\partial G}{\partial \sigma_{\hat{\epsilon}}} > 0$$

if and only if $\alpha < \tilde{\alpha}_1$. Further, by the Implicit Function Theorem,

$$\frac{d\alpha^*}{d\sigma_z} = -\frac{\frac{\partial G}{\partial \sigma_{\hat{\epsilon}}}}{\frac{\partial G}{\partial \alpha}}$$

and the denominator is positive by the local convexity of the objective around the minimizer characterized by the FOC. Thus, fixing γ, if $\alpha^* < \tilde{\alpha}_1$, α^* is decreasing in $\sigma_{\hat{\epsilon}}$. If instead $\alpha^* > \tilde{\alpha}_1$, α^* is increasing in $\sigma_{\hat{\epsilon}}$. Finally, if $\alpha^* = \tilde{\alpha}_1$, α^* is independent of $\sigma_{\hat{\epsilon}}$.

(iii) *Behavior of γ as $\sigma_{\hat{\epsilon}}$ increases.* Recall that if $\overline{r}^n = 0$ we have $b^* = 0$ and the γ–FOC is

$$\underbrace{(\xi_r + \alpha\xi_i\tilde{\omega}_r\rho_r)\sigma_{r\theta} + (\xi_u + \alpha\xi_i\tilde{\omega}_u\rho_u)\sigma_{u\theta} + \xi_i\sigma_\theta^2\gamma + 4W_1\left[\gamma\left(1 - \Phi(\gamma/\sigma_z)\right) - \sigma_z\phi(\gamma/\sigma_z)\right]}_{=:F(\gamma,\sigma_{\hat{\epsilon}})} = 0,$$

First, note that as proved in (i) above,

$$\frac{\partial F}{\partial \sigma_z} < 0,$$

and, clearly,

$$\frac{\partial F}{\partial \sigma_\theta^2} = \xi_i\gamma \geq 0.$$

Now, by assumption

$$\alpha > -\frac{\xi_r}{\xi_i \tilde{\omega}_r \rho_r}, -\frac{\xi_u}{\xi_i \tilde{\omega}_u \rho_u}$$

and thus

$$\frac{\partial F}{\partial \sigma_{r\theta}} = \xi_r + \alpha \xi_i \tilde{\omega}_r \rho_r > 0, \qquad \frac{\partial F}{\partial \sigma_{u\theta}} = \xi_u + \alpha \xi_i \tilde{\omega}_u \rho_u > 0.$$

Furthermore,

$$\frac{\partial F}{\partial \sigma_{\tilde{\epsilon}}} = \frac{\partial F}{\partial \sigma_z} \frac{d\sigma_z}{d\sigma_{\tilde{\epsilon}}} + \frac{\partial F}{\partial \sigma_\theta^2} \frac{d\sigma_\theta^2}{d\sigma_{\tilde{\epsilon}}} + \frac{\partial F}{\partial \sigma_{r\theta}} \frac{d\sigma_{r\theta}}{d\sigma_{\tilde{\epsilon}}} + \frac{\partial F}{\partial \sigma_{u\theta}} \frac{d\sigma_{u\theta}}{d\sigma_{\tilde{\epsilon}}}.$$

Thus, to sign the total effect we need to sign $d\sigma_z/d\sigma_{\tilde{\epsilon}}$, $d\sigma_\theta^2/d\sigma_{\tilde{\epsilon}}$, $d\sigma_{r\theta}/d\sigma_{\tilde{\epsilon}}$ and $d\sigma_{u\theta}/d\sigma_{\tilde{\epsilon}}$. Let's work through then in that order. As above, let $a_{cb}^0(s_0, \hat{s}_1) = w_0' s_0 + \hat{w}\hat{s}_1$ and

$$\pi := \frac{\mathrm{Cov}(\tilde{s}_1, a_{cb}^0)}{\mathrm{Var}(\tilde{s}_1)}.$$

Thus, $\tilde{\alpha}_1 = a_{cb} + \pi$. Define $\Delta = \alpha - \tilde{\alpha}_1$. We have

$$\mathbb{E}\left[z_{\alpha,0} \mid \tilde{s}_1\right] = \left(\alpha - a_{cb} - \pi\right) \tilde{s}_1 = \Delta \tilde{s}_1,$$

hence, by the law of total variance,

$$\mathrm{Var}(z_{\alpha,0}) = \Delta^2 \,\mathrm{Var}(\tilde{s}_1) + \mathrm{Var}\left(a_{cb}^0 \mid \tilde{s}_1\right). \tag{6}$$

Let $V := \mathrm{Var}(\tilde{s}_1) = V_0 + \sigma_{\tilde{\epsilon}}^2$, where V_0 is the part independent of $\sigma_{\tilde{\epsilon}}^2$. Let $C := \mathrm{Cov}(\tilde{s}_1, a_{cb}^0)$ (which does not depend on $\sigma_{\tilde{\epsilon}}^2$ because $\tilde{\epsilon}$ is independent of $(s_0, \hat{s}_1)$). Then, $\pi = C/V$ and using the law of total variance again,

$$\mathrm{Var}\left(a_{cb}^0 \mid \tilde{s}_1\right) = \mathrm{Var}(a_{cb}^0) - \frac{C^2}{V}.$$

Thus,

$$\frac{d}{d\sigma_{\tilde{\epsilon}}^2} \mathrm{Var}\left(a_{cb}^0 \mid \tilde{s}_1\right) = \frac{C^2}{V^2} > 0.$$

Let $\tilde{s}_1^\perp := \tilde{s}_1 - \mathbb{E}[\tilde{s}_1 \mid s_0, \hat{s}_1]$ and $Y^\perp := Y - \mathbb{E}[Y \mid s_0, \hat{s}_1]$, i.e., the residuals of $\tilde{s}_1$ and $Y = i_1^{\mathrm{FI}}(\lambda_{cb})$ when projecting on $(s_0, \hat{s}_1)$. Then, by Frisch–Waugh–Lovell,

$$a_{cb} = \frac{\mathrm{Cov}(Y^\perp, \tilde{s}_1^\perp)}{\mathrm{Var}(\tilde{s}_1^\perp)}.$$

Increasing $\sigma_{\tilde{\epsilon}}$ raises $\mathrm{Var}(\tilde{s}_1^\perp)$ one-for-one but leaves $\mathrm{Cov}(Y^\perp, \tilde{s}_1^\perp)$ unchanged. Hence,

$$\frac{da_{cb}}{d\sigma_{\tilde{\epsilon}}^2} = -\frac{\mathrm{Cov}(Y^\perp, \tilde{s}_1^\perp)}{\mathrm{Var}(\tilde{s}_1^\perp)^2} = -\frac{a_{cb}}{\mathrm{Var}(\tilde{s}_1^\perp)} < 0.$$

From (6) we have

$$\sigma_z^2 = \Delta^2 V + \mathrm{Var}(a_{cb}^0) - \frac{C^2}{V}.$$

Recall $\Delta = \alpha - a_{cb} - \frac{C}{V}$. Differentiating with respect to V, we get

$$\frac{d\sigma_z^2}{dV} = 2\Delta\frac{\partial\Delta}{\partial V}V + \Delta^2 + \frac{C^2}{V^2}, \qquad \frac{\partial\Delta}{\partial V} = -\frac{da_{cb}}{dV} + \frac{C}{V^2}.$$

Substituting the latter and rearranging,

$$\frac{d\sigma_z^2}{dV} = -2\Delta V\frac{da_{cb}}{dV} + \underbrace{\Delta^2 + \tfrac{2\Delta C}{V} + \tfrac{C^2}{V^2}}_{=(\alpha - a_{cb})^2}.$$

Using this and substituting da_{cb}/dV, we get

$$\frac{d\sigma_z^2}{d\sigma_{\tilde{\epsilon}}^2} = \frac{d\sigma_z^2}{dV} = (\alpha - a_{cb})^2 + 2\,(\alpha - \tilde{\alpha}_1)\,\frac{\mathrm{Var}\,\tilde{s}_1}{\mathrm{Var}(\tilde{s}_1^\perp)}a_{cb}.$$

Hence,

- If $\alpha \geq \tilde{\alpha}_1$ (i.e., the announcement more data-dependent than the $t = 1$ target), the second term is non-negative, so $d\sigma_z^2/d\sigma_{\tilde{\epsilon}}^2 \geq (\alpha - a_{cb})^2 > 0$. In this region σ_z^2 increases *strictly* with $\sigma_{\tilde{\epsilon}}^2$.
- If $\alpha < \tilde{\alpha}_1$, the second term is non-positive and partially offsets $(\alpha - a_{cb})^2$. The net sign is ambiguous.

Now let's turn to $\sigma_\theta^2 = \mathrm{Var}(\theta_{\alpha,0})$. Write

$$\theta_{\alpha,0}(s_0) = 2\Phi\left(\frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) - a_{ps}^0(s_0)}{|\alpha - a_{ps}|\sigma_{\tilde{s}}}\right) - 1, \qquad \sigma_{\tilde{s}}^2 := \mathrm{Var}(\tilde{s}_1 \mid s_0)$$

and define

$$T(s_0) := \frac{(\alpha - a_{ps})\mu_{\tilde{s}}(s_0) - a_{ps}^0(s_0)}{|\alpha - a_{ps}|\sigma_{\tilde{s}}}, \quad \theta_{\alpha,0} = g(T), \quad g(x) := 2\Phi(x) - 1.$$

Its variance is

$$s_T^2 := \mathrm{Var}\left(T_\alpha(s_0)\right) = \frac{V_*}{(\alpha - a_{ps})^2 \sigma_{\tilde{s}}^2},$$

where

$$V_* := \mathrm{Var}\left(\alpha \mu_{\tilde{s}}(s_0) - m_Y(s_0)\right), \qquad m_Y(s_0) := \mathbb{E}\left[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0\right].$$

Note that V_* does not depend on $\sigma_{\tilde{\epsilon}}^2$. On the other hand, as $\sigma_{\tilde{\epsilon}}^2$ increases, $\sigma_{\tilde{s}}^2 = \mathrm{Var}(\tilde{s}_1 \mid s_0)$ increases one-for-one, while the conditional covariance in the numerator of

$$a_{ps} = \frac{\mathrm{Cov}(i_1^{\mathrm{FI}}(\lambda_{ps}), \tilde{s}_1 \mid s_0)}{\mathrm{Var}(\tilde{s}_1 \mid s_0)}$$

is unchanged, so a_{ps} is decreasing in $\sigma_{\tilde{\epsilon}}^2$. Because $\alpha > a_{ps}$, $(\alpha - a_{ps})^2$ increases. Hence the denominator

$$(\alpha - a_{ps})^2 \sigma_{\tilde{s}}^2$$

is increasing in $\sigma_{\tilde{\epsilon}}^2$, implying $ds_T^2/d\sigma_{\tilde{\epsilon}}^2 < 0$. Finally, since g is odd and increasing, $\mathrm{Var}\left(g(T_\alpha)\right)$ is increasing in $\mathrm{Var}(T_\alpha)$. Hence,

$$\frac{d}{d\sigma_{\tilde{\epsilon}}^2} \mathrm{Var}\left(\theta_{\alpha,0}\right) = \frac{d}{ds_T^2} \mathrm{Var}\left(g(T_\alpha)\right) \frac{ds_T^2}{d\sigma_{\tilde{s}}^2} < 0,$$

so the variance of the worst-case selector $\theta_{\alpha,0}$ decreases as $\sigma_{\tilde{\epsilon}}$ increases.

It remains to sign $d\sigma_{r\theta}/d\sigma_{\tilde{\epsilon}}$ and $d\sigma_{u\theta}/d\sigma_{\tilde{\epsilon}}$. To do so, note first that by a standard dominated convergence argument used many times already (see, e.g., the proof of Lemma 5),

$$\frac{d\sigma_{r\theta}}{d\sigma_{\tilde{\epsilon}}} = \mathbb{E}\left[r_0^n \frac{\partial \theta_{\alpha,0}}{\partial \sigma_{\tilde{\epsilon}}}\right]$$

Further,

$$\frac{\partial \theta_{\alpha,0}}{\partial \sigma_{\tilde{\epsilon}}} = 2\phi(T) \frac{\partial T}{\partial \sigma_{\tilde{\epsilon}}}$$

To compute $\partial T/\partial \sigma_{\tilde{\epsilon}}$, note first that by definition, $\mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0, \tilde{s}_1] = a_{ps}\tilde{s}_1 + a_{ps}^0(s_0)$. Taking $\mathbb{E}[\cdot \mid s_0]$ on both sides gives $\mathbb{E}[i_1^{\mathrm{FI}}(\lambda_{ps}) \mid s_0] = a_{ps}\mu_{\tilde{s}}(s_0) + a_{ps}^0(s_0)$. The left-hand side does not depend on $\sigma_{\tilde{\epsilon}}$, hence pointwise in s_0

$$\frac{d}{d\sigma_{\tilde{\epsilon}}}\left(a_{ps}\mu_{\tilde{s}}(s_0) + a_{ps}^0(s_0)\right) = 0.$$

Therefore, letting $\Delta := \alpha - a_{ps}$,

$$\frac{d}{d\sigma_{\tilde{\epsilon}}}\left(\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0)\right) = -\frac{da_{ps}}{d\sigma_{\tilde{\epsilon}}}\mu_{\tilde{s}}(s_0) - \frac{da^0_{ps}}{d\sigma_{\tilde{\epsilon}}}(s_0) = 0,$$

so the numerator of T is invariant w.r.t. $\sigma_{\tilde{\epsilon}}$.

On the other hand, the denominator is given by

$$D(\sigma_{\tilde{\epsilon}}; s_0) = |\Delta|\sigma_{\tilde{s}}.$$

Since the numerator of T does not depend on $\sigma_{\tilde{\epsilon}}$,

$$\frac{\partial T}{\partial \sigma_{\tilde{\epsilon}}} = -\frac{\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0)}{D(\sigma_{\tilde{\epsilon}}; s_0)^2}\frac{dD(\sigma_{\tilde{\epsilon}}; s_0)}{d\sigma_{\tilde{\epsilon}}}.$$

Moreover,

$$\begin{aligned}\frac{dD}{d\sigma_{\tilde{\epsilon}}} &= \frac{1}{2D}\frac{d}{d\sigma_{\tilde{\epsilon}}}\left(\Delta^2\sigma^2_{\tilde{s}}\right) \\ &= \frac{1}{D}\left(\Delta\sigma^2_{\tilde{s}}\frac{d\Delta}{d\sigma_{\tilde{\epsilon}}} + \Delta^2\sigma_{\tilde{\epsilon}}\right).\end{aligned}$$

As argued before, $d\Delta/d\sigma_{\tilde{\epsilon}} = -da_{ps}/d\sigma_{\tilde{\epsilon}} > 0$. Furthermore, by assumption $\alpha > a_{ps}$, so $\Delta > 0$. Hence $dD/d\sigma_{\tilde{\epsilon}} \geq 0$ (strictly if $\Delta > 0$). Therefore,

$$\frac{\partial \theta_{\alpha,0}}{\partial \sigma_{\tilde{\epsilon}}} = 2\phi(T)\frac{\partial T}{\partial \sigma_{\tilde{\epsilon}}} = 2\phi(T)\left[-\frac{\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0)}{D^2}\frac{dD}{d\sigma_{\tilde{\epsilon}}}\right],$$

so its sign is the opposite of the sign of $\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0)$. Hence,

$$\mathbb{E}\left[r^n_0\frac{\partial \theta_{\alpha,0}}{\partial \sigma_{\tilde{\epsilon}}}\right] = -\mathbb{E}\left[r^n_0\phi(T)\frac{dD/d\sigma_{\tilde{\epsilon}}}{D^2}\left(\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0)\right)\right].$$

The weight $\phi(T)\frac{dD/d\sigma_{\tilde{\epsilon}}}{D^2}$ is strictly positive, so the sign equals the negative of $\mathbb{E}\left[r^n_0(\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0))\right]$. Then,

$$\mathbb{E}\left[r^n_0(\Delta\mu_{\tilde{s}}(s_0) - a^0_{ps}(s_0))\right] = \Delta\,\mathrm{Cov}(r^n_0, \mu_{\tilde{s}}) - \mathrm{Cov}(r^n_0, a^0_{ps}) = \sigma^2_r\left(\Delta\tilde{\omega}_r\rho_r - w_r\right),$$

where $w_r > 0$ is the loading of r^n_0 in $a^0_{ps}(s_0)$. By definition, we know

$$\mathbb{E}\left[i^{\mathrm{FI}}_1(\lambda_{ps}) \mid s_0\right] = a_{ps}\mathbb{E}\left[\tilde{s}_1 \mid s_0\right] + a^0_{ps}(s_0),$$

hence

$$a^0_{ps}(s_0) = \underbrace{\mathbb{E}\left[i^{\mathrm{FI}}_1(\lambda_{ps}) \mid s_0\right]}_{=\rho_r r^n_0 + \Gamma_{ps}\rho_u u_0} - a_{ps} \underbrace{\mathbb{E}\left[\tilde{s}_1 \mid s_0\right]}_{=\tilde{\omega}_r \rho_r r^n_0 + \tilde{\omega}_u \rho_u u_0}.$$

Therefore,

$$w_r = \rho_r \left(1 - a_{ps}\tilde{\omega}_r\right).$$

and thus, as $\alpha > 1/\tilde{\omega}_r$,

$$\Delta\tilde{\omega}_r\rho_r - w_r = (\alpha\tilde{\omega}_r - 1)\rho_r > 0$$

Hence,

$$\frac{\partial\sigma_{r\theta}}{\partial\sigma_{\tilde{\epsilon}}} = \mathbb{E}\left[r^n_0 \frac{\partial\theta_{\alpha,0}}{\partial\sigma_{\tilde{\epsilon}}}\right] < 0.$$

A fully analogous argument applies to $E[u_0 \partial\theta_{\alpha,0}/\partial\sigma_{\tilde{\epsilon}}]$. Putting all the above together, we have that provided $\alpha \geq \alpha^\star$,

$$\frac{\partial F}{\partial\sigma_{\tilde{\epsilon}}} = \underbrace{\frac{\partial F}{\partial\sigma_z}}_{<0}\underbrace{\frac{d\sigma_z}{d\sigma_{\tilde{\epsilon}}}}_{>0} + \underbrace{\frac{\partial F}{\partial\sigma^2_\theta}}_{\geq 0}\underbrace{\frac{d\sigma^2_\theta}{d\sigma_{\tilde{\epsilon}}}}_{<0} + \underbrace{\frac{\partial F}{\partial\sigma_{r\theta}}}_{>0}\underbrace{\frac{d\sigma_{r\theta}}{d\sigma_{\tilde{\epsilon}}}}_{<0} + \underbrace{\frac{\partial F}{\partial\sigma_{u\theta}}}_{>0}\underbrace{\frac{d\sigma_{u\theta}}{d\sigma_{\tilde{\epsilon}}}}_{<0} < 0.$$

Then, because

$$\frac{\partial F}{\partial\gamma} > 0,$$

as proven in (i) above, we have that the Implicit Function Theorem yields

$$\frac{\partial\gamma^*}{\partial\sigma_{\tilde{\epsilon}}} = -\frac{\frac{\partial F}{\partial\sigma_{\tilde{\epsilon}}}}{\frac{\partial F}{\partial\gamma}} > 0.$$

and so optimal vagueness γ^* increases with $\sigma_{\tilde{\epsilon}}$ for fixed α.

D Proofs of Additional Technical Lemmas

Proof of Lemma 4. Let $x := \begin{pmatrix} r^n_1 & u_1 \end{pmatrix}^T$. Conditional on (r^n_0, u_0), we have

$$x \mid (r^n_0, u_0) \sim \mathcal{N}(\mu_{|0}, P_{|0}), \quad \mu_{|0} := \begin{pmatrix} (1-\rho_r)\overline{r}^n + \rho_r r^n_0 \\ \rho_u u_0 \end{pmatrix}, \quad P_{|0} := \mathrm{diag}(\sigma^2_{\epsilon_r}, \sigma^2_{\epsilon_u}),$$

Define

$$\tilde{\Omega} := (\tilde{\omega}_r \tilde{\omega}_u), \qquad \hat{\Omega} := (\hat{\omega}_r \hat{\omega}_u),$$

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,

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

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 $\varepsilon \sim \mathcal{N}(0, R)$. Again, by standard results given the Gaussian structure,

$$\mathbb{E}[x \mid y] = \mu_{|0} + K(y - H\mu_{|0}), \qquad K := PH^\top(HPH^\top + R)^{-1}.$$

Let $S := HPH^\top + R$ with entries

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

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

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

Writing $\eta := y - H\mu = \begin{pmatrix}\eta_1 \\ \eta_2\end{pmatrix}$, 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 $\mu + K\eta$ is

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

and the second component is

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

Consequently,

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

where

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

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

Proof of Lemma 3. 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 $m_{ps}(s, s_0) := \mathbb{E}[i_1^{\text{FI}}(\lambda_{ps}) \mid s_0, \tilde{s}_1 = s] = a_{ps}s + a^0_{ps}(s_0)$. 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 $\gamma(s) := \frac{U(s)-L(s)}{2} \geq 0$,

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

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

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

as stated. □

Proof of Lemma 5. 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

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

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

We treat the $t = 1$ and $t = 0$ component of the objective separately for clarity.

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

$$\frac{\partial h(z,\gamma)}{\partial z} = 2(|z| - \gamma)_+ \operatorname{sign} z, \qquad \frac{\partial h(z,\gamma)}{\partial \gamma} = -2(|z| - \gamma)_+,$$

and these formulas hold for all z except possibly on the set $\{|z| = \gamma\}$. Here $z = (\alpha - a_{cb})\tilde{s}_1 + (b - a_{cb}^0)$ is an affine function of jointly Gaussian variables $(\tilde{s}_1, \hat{s}_1, r_0^n, u_0)$, thus z is Gaussian with finite second moment and $\{|z| = \gamma\}$ has measure zero.

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

$$\left|(|z| - \gamma)_+ \tilde{s}_1 \operatorname{sign} z\right| \le |z||\tilde{s}_1|,$$

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

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

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

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

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,

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

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

for some constant C that does not depend on the particular (α, b, γ) picked (it depends on d_0 through $\mathcal{N}$ alone). Because all random variables are Gaussian, the right hand sides are integrable. Therefore, by the Dominated Convergence Theorem and the derivatives of h computed above,

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

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

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

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

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

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

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

Further, it's straightforward to verify that

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

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

$$\frac{\partial I}{\partial \alpha} = \mu_{\tilde{s}} + \gamma\frac{\partial \theta_{\alpha,b}}{\partial \alpha}, \quad \frac{\partial I}{\partial b} = 1 + \gamma\frac{\partial \theta_{\alpha,b}}{\partial b}, \quad \frac{\partial I}{\partial \gamma} = \theta_{\alpha,b} \tag{7}$$

with

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

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 $x \in \mathbb{R}$,

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

$$\left|\frac{\partial \theta_{\alpha,b}}{\partial \alpha}\right| \leq \frac{2}{\sqrt{2\pi}} \left(\frac{1}{\varepsilon^2 \sigma_{\tilde{s}}} |b - a^0_{ps}(s_0)|\right) \leq c_\alpha \Big(1 + |r^n_0| + |u_0|\Big),$$

for constants c_b, c_α depending only on $\mathcal{N}$ and parameters (the rightmost inequality uses that a^0_{ps} 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

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

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

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

and clearly grows at most linearly on $\mathcal{N}$. Therefore 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

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

and continuity of the above derivative follows from the same argument (recall the expression inside the expectation in Φ 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 $(|z| - \gamma)_+^2 \to |z|^2$ monotonically. Thus, by the Monotone Convergence Theorem, $\mathbb{E}[(|z| - \gamma)_+^2] \to \mathbb{E}[|z|^2]$. For $\partial/\partial\alpha$ and $\partial/\partial b$, note that $(|z| - \gamma)_+ \operatorname{sign} z \to |z| \operatorname{sign} z = z$. Further, the bounds found above do not depend on γ near 0, so can use those and apply the Dominated Convergence Theorem to get that as $\gamma \downarrow 0$,

$$\frac{\partial W_1 \,\mathbb{E}[(|z_{\alpha,b}| - \gamma)_+^2]}{\partial \alpha} \to 2W_1 \,\mathbb{E}[z_{\alpha,b}\tilde{s}_1],$$

$$\frac{\partial W_1 \,\mathbb{E}[(|z_{\alpha,b}| - \gamma)_+^2]}{\partial b} \to 2W_1 \,\mathbb{E}[z_{\alpha,b}],$$

proving the continuity at $\gamma = 0$ of the derivatives. Moreover, for $\partial/\partial\gamma$, $(|z| - \gamma)_+ \to |z|$ 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 $\gamma \downarrow 0$,

$$I(\alpha, b, \gamma; s_0) \to I_0(s_0) = \alpha\mu_{\tilde{s}}(s_0) + b, \qquad m_I(\alpha, b, \gamma) \to m_{I,0} := \mathbb{E}[I_0],$$

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 $\mu_{\tilde{s}}$ and a_{ps}^0 are affine in (r_0^n, u_0) and $|\theta_{\alpha,b}| \leq 1$, for any fixed $\gamma_1 > 0$,

$$|I(\alpha, b, \gamma; s_0)| \leq |I_0(s_0)| + \gamma_1 \leq C(1 + |r_0^n| + |u_0|), \qquad \gamma \in [0, \gamma_1],$$

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. Therefore $QI^2 + 2HI$ 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,

$$\mathbb{E}[QI^2 + 2HI] \to \mathbb{E}[QI_0^2 + 2HI_0],$$

and the other term in Ψ converges since $m_I(\alpha, b, \gamma) \to m_{I,0}$. Thus $\Psi(\alpha, b, \gamma) \to \Psi(\alpha, b, 0)$.

Now, for $\gamma > 0$, we proved above that

$$\frac{\partial \Psi}{\partial \alpha} = \mathbb{E}\left[(2QI + 2H)\left(\mu_{\tilde{s}} + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha}\right)\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_I\right)\mathbb{E}\left[\mu_{\tilde{s}} + \gamma \frac{\partial \theta_{\alpha,b}}{\partial \alpha}\right],$$

$$\frac{\partial \Psi}{\partial b} = \mathbb{E}\left[(2QI + 2H)\left(1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b}\right)\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_I\right)\mathbb{E}\left[1 + \gamma \frac{\partial \theta_{\alpha,b}}{\partial b}\right].$$

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 $\gamma \in [0, \gamma_1]$,

$$\left|\frac{\partial \theta_{\alpha,b}}{\partial b}\right| \leq c_b, \qquad \left|\frac{\partial \theta_{\alpha,b}}{\partial \alpha}\right| \leq c_\alpha(1 + |r_0^n| + |u_0|),$$

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 $I \to I_0$, $m_I \to m_{I,0}$ and $\gamma \frac{\partial \theta_{\alpha,b}}{\partial p} \to 0$ for $p \in \{\alpha, b\}$, the Dominated Convergence Theorem implies

$$\frac{\partial \Psi}{\partial \alpha} \to \mathbb{E}\left[(2QI_0 + 2H)\mu_{\tilde{s}}\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_{I,0}\right)\mathbb{E}[\mu_{\tilde{s}}],$$

and

$$\frac{\partial \Psi}{\partial b} \to \mathbb{E}\left[(2QI_0 + 2H)\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_{I,0}\right).$$

Finally, using the results above, for $\gamma > 0$,

$$\frac{\partial \Psi}{\partial \gamma} = \mathbb{E}\left[(2QI + 2H)\theta_{\alpha,b}\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_I\right)\mathbb{E}[\theta_{\alpha,b}].$$

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 $I \to I_0$ and $m_I \to m_{I,0}$ we can apply again the Dominated Convergence Theorem to obtain the limit of the right derivative,

$$\frac{\partial \Psi}{\partial \gamma} = \mathbb{E}\left[(2QI_0 + 2H)\theta_{\alpha,b}\right] - \frac{2\delta_{i^e}}{1+\lambda_{cb}\kappa^2}\left(\delta_r \overline{r}^n + \delta_{i^e} m_{I,0}\right)\mathbb{E}[\theta_{\alpha,b}],$$

thus completing the proof of the continuous extension of the derivatives of Ψ, and thus of $\mathcal{J}$ to $\gamma = 0$. □

Proof of Lemma 6. 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

$$2h_z = 2(|z| - \gamma)_+ \mathrm{sgn} z, \quad h_{zz} = 2\mathbf{1}_{\{|z|>\gamma\}}, \quad h_{z\gamma} = -2\mathbf{1}_{\{|z|>\gamma\}} \mathrm{sgn} z, \quad h_{\gamma\gamma} = 2\mathbf{1}_{\{|z|>\gamma\}}.$$

Because z is an affine function of jointly Gaussian variables, $\Pr(|z| = \gamma) = 0$ 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 Proof of 5 for details of an analogous argument). 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 $\Psi = \mathbb{E}[QI^2 + 2HI] - \frac{(\delta_r \bar{r}^n + \delta_{i^e} m_I)^2}{1+\lambda_{cb}\kappa^2}$ 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 (5). Each term is continuous by dominated convergence: I, I_p, I_{pq} grow at most linearly in (r_0^n, u_0) on the neighborhood, while Q, H, G are affine in I so at most linear in (r_0^n, u_0). This yields $\Psi \in C^2$ 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 together using that $\Pr(|z| = \gamma) = 0$ and the ε–separation from $\alpha = a_{ps}$. □

Proof of Lemma 7. Using the expressions for the derivatives in Lemma 6, one gets that for $\gamma > 0$,

$$\begin{aligned}
\frac{\partial^2 \mathcal{J}}{\partial\gamma^2} &= \xi_i\, \mathbb{E}[\theta_{\alpha,b}^2] - \frac{2\delta_{i^e}^2}{1+\lambda_{cb}\kappa^2}\, \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma) \\
&= \xi_i \sigma_\theta^2 + \left(\xi_i - \frac{2\delta_{i^e}^2}{1+\lambda_{cb}\kappa^2}\right) \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma) \\
&= \xi_i \sigma_\theta^2 + \frac{2\lambda_{cb}}{1+\lambda_{cb}\kappa^2}(\beta\sigma\kappa)^2\, \mathbb{E}[\theta_{\alpha,b}]^2 + 2W_1 \mathbb{P}(|z_{\alpha,b}| > \gamma),
\end{aligned}$$

where the last equality follows from the definitions of ξ_i and δ_{i^e} and some algebra.

Note that the first two terms are independent of γ and as $\gamma \downarrow 0$, $\Pr(|z| > \gamma) \to 1$ 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 $\mathbf{1}_{\{|z|>\gamma\}} \to 1$ almost surely. This yields the stated expressions. $\square$