Built and Spent by Words: Forward Guidance under Inflation Targeting*

Lukasz A. Drozd

Federal Reserve Bank of Philadelphia

lukasz.drozd@phil.frb.org

August 2026

Abstract

Modern central banks communicate through a standing inflation target and discretionary forward guidance. This rule-plus-shift architecture can provide an informationally efficient representation of Ramsey policy. Here we take this institutional arrangement as given and study its positive consequences when operational promises are extended under discretion and imperfectly enforceable (delegated institutional design). The key tension is between the target and forward guidance; the target anchors expectations but makes large departure promises less credible. In a Gaussian–logistic benchmark developed herein, implemented guidance rises, reaches a finite frontier, and then declines—tracing an endogenous Laffer curve of communication. Words can both build and spend central bank credibility capital, and traction depends on the state and the precision with which the public infers from the central bank's words and actions from that state. The leading application of our theory shows looking through large supply shocks can stabilize expectations or backfire when communication is interpreted as policy preference rather than information.

Keywords: central-bank communication, forward guidance, inflation targeting, credibility, asymmetric information, optimal monetary policy.
JEL codes: E31, E43, E52, E58, D82, D83.

*I thank Jonas Arias, João Gomes, Urban Jermann, Makoto Nakajima, and Alessandro Dovis for insightful comments at various stages of this project. Alessandro's comments helped clarify the institutional interpretation developed here and in a companion paper. The views expressed herein are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Correspondence: Lukasz A. Drozd, Federal Reserve Bank of Philadelphia, Research Department, 10 N. Independence Mall West, Philadelphia, PA 19106.

1 Introduction

Modern central banks rely on two legs of communication: a standing commitment to price stability around a concrete numerical target—reinforced by institutionalized accountability for deviations from that target—and operational forward guidance about future policy that defines intentional departures from the policy path implied by the standing regime. Operational guidance can be adjusted quickly; changes to the standing institution move slowly, or not at all when imposed legislatively. The key point of the current paper is that the standing targeting regime shapes how the public interprets operational statements. There is no true silence: absent a discretionary announcement, the central bank reaffirms the targeting rule. The literature typically treats these two legs jointly. Here we separate them by taking the standing targeting regime as given and studying the limits it imposes on operational communication. We remain agnostic about the regime's precise form: it is stationary, anchors long-run inflation at a fixed numerical target, and may condition on the exogenous state, but not on the central bank's current private information or on the endogenous history of its own policy and communication.

This tension is apparent in two core uses of forward guidance. The first is steering the economy out of an effective-lower-bound recession, a strategy broadly employed after both the Global Financial Crisis and the COVID-19 pandemic. Not being able to lower interest rate further, a central bank may want to announce that it will allow the economy to "run hot" after liftoff becomes possible, effectively promising inflation above target in order to stimulate activity today. The conflict arises when the public anticipates that delivering such guidance may become too costly institutionally, making it self-defeating.

The second use is a policy of "looking through" a large supply shock, vividly displayed during the post-pandemic inflation outbreak across advanced economies, including the United States—as recently characterized by Nakamura et al. (2025), English et al., eds (2024) and Forbes et al. (2025). Here the central bank accommodates the cost-push component of the shock in order to support the recovery while communicating that the disturbance is temporary and reassuring the public of a swift return to target—provided its internal analysis (private information) supports such an assessment. If believed, this communication can restrain inflation expectations and reduce the sacrifice ratio of subsequent disinflation, as these authors describe was the case. But it can also lock the central bank into a policy that becomes difficult to deliver if the shock proves more persistent than expected, which again was on full display after the pandemic. Honoring the promise may then require a severe contraction, whereas abandoning it conflicts with the institution's commitment to the target, inflicting institutional penalties implied by the delegated targeting regime, and

damaging credibility.

Both conflicts arise when the numerical target alone is insufficient to characterize desirable policy. Under divine coincidence, stabilizing inflation also stabilizes the output gap, leaving little independent role for operational guidance. Effective-lower-bound constraints and cost-push shocks break this coincidence. They create a stabilization trade-off and make communication about future policy valuable, while bringing that communication into conflict with the standing regime that set out to characterize. In richer models, it should be noted that cost-push disturbances arise from sectoral shocks or asymmetries. Even when a version of divine coincidence survives, it may apply to a specialized price index rather than the index targeted in practice (Rubbo, 2024a), in effect forcing accommodation relative to the actual target.

The proposed theory has two key elements. First, inflation targeting is not just a number; it is an institutional constraint on operational departures from the numerical target and promises made as part of forward guidance. The inherited rule can condition on publicly observed and verifiable states, but not on the central bank's private information about stabilization needs or future shocks. Private information gives rise to communication need (Bassetto, 2019).

Second, private information may concern the severity or persistence of future shocks, or it may represent a temporary change in the bank's perceived stabilization needs—its preference for inflation versus employment. Issuance is discretionary: after observing its private signal, the bank chooses whether and how far to shift the inherited rule. Delivery is imperfectly committed: the announcement affects expectations because future policy is tied to it, but the bank may subsequently waive the promise (commitment to carry out forward guidance ex post) at an institutional cost. This flexibility protects the institution from carrying out a disastrous promise after circumstances change, but the possibility of waiver limits the promise's traction when it is issued.

The rule-plus-guidance architecture is not an arbitrary modeling choice here, and in addition to positive motivation it finds support in normative theory. In a companion paper, Drozd (2026) shows that, in an infinite-horizon linear–quadratic New Keynesian benchmark, the Ramsey plan can be represented as a public state-contingent rule plus a scalar history-dependent shift,

$$\Theta^R_{t+1}(S_{t+1}) = \bar{\Theta}^R(S_{t+1}) + g^R_t, \tag{1}$$

where $\bar{\Theta}^R(S_{t+1})$ is the inflation policy center based on the observable state and g^R_t is the forward guidance shift of that rule. This scalar is a normalized recursive co-state and

is the only changing statistic that current private decisions need from the continuation plan. Under a mutual-information processing cost, devoting attention to other schedule details is wasteful. Institutional design of inflation target as a simple rule and institutional penalties can be seen as informationally efficient approach of implementing something close to Ramsey plan. That being said, here, our object is the positive problem created by that gap between the ideal representation and the institution that implements it.

The delegated monetary authority operates in one of two persistent and hidden modes that differ in their resistance to an ex-post waiver. The ordinary mode faces a higher institutional cost of abandoning an outstanding promise; the override mode faces a lower cost. The modes share the same numerical target, default rule, social objective, information structure, and policy instruments; they differ only in commitment technology. The public learns about the active mode first from the announcement and later from realized policy relative to the promise. Credibility is a form of institutional capital that evolves through both words and actions. This friction leads to additional interaction of communication traction and credibility, adding a dynamic dimension to the model that can be relevant in some applications.

Our theoretical analysis proceeds in two steps. We first derive a local linear–quadratic, second-order representation of the optimized value of releasing an inherited promise around the target equilibrium. The calculation compares the fully reoptimized honor and waiver programs and identifies the local temptation to waive. We then combine this Bellman equation with Gaussian message likelihoods and logistic waiver cost structure. In that benchmark, we show that the implemented effect of guidance follows a unique Laffer curve of communication. A larger announcement initially moves expectations by more. As its distance from the inherited rule grows, however, two forces erode its effectiveness. First, the announcement becomes increasingly characteristic of the override mode and is discounted on delivery by the public (shifts expectations less). Second, future implementation becomes more costly, increasing the probability that the announcement will be waived. The implemented expectation shift reaches a finite frontier and then declines: beyond that point, saying more achieves less while causing greater institutional damage.

Words can both build and spend credibility over time. A moderate announcement may improve the public's assessment of the institution because it is more likely to have originated under the ordinary mode. Yet making that same message more extreme shifts beliefs toward the override mode, draining credibility. As the departure grows, the credibility gain diminishes and eventually becomes a credibility loss.

The key element that substantive analysis focuses on is that traction is state dependent

because the central bank is benevolent and the public understands its incentives in an noisy way. The public evaluates an announcement jointly with the inherited credibility state and the observable macroeconomic environment, both of which affect the bank's anticipated incentive to waive the announcement when it comes due. The same words can therefore have different effects across episodes.

This margin implies that communication traction—measured by its impact on inflation expectations—can shrink precisely when it is most valuable, in adverse economic states or at low levels of inherited credibility. Under such circumstances, the public anticipates a stronger incentive to waive and discounts the announcement more heavily. Our quantitative application quantifies this mechanism.

Our leading quantitative application focuses on the theory's implications for the post-pandemic inflation surge and the use of the so-called look-through response to supply disturbances that initially had a large cost-push component (Rubbo, 2024b; Drozd and Tavares, 2026). When the central bank privately expects a shock to be less severe or less persistent than the public believes, it may wish to communicate a rapid return to target while accommodating the current shock. This logic featured prominently in the response to post-COVID inflation across advanced economies with highly credible monetary authorities—whether by design or in subsequent post-mortem interpretations. Many highly credible central banks, including the Federal Reserve, initially described inflation as transitory, later tightened rapidly, and experienced an unusual combination of persistently anchored expectations and disinflation at a relatively low sacrifice ratio, as emphasized by Nakamura et al. (2025) and Forbes et al. (2025).[1]

In our model, not surprisingly, the inherited flexible-targeting rule prescribes partial accommodation of the cost-push shock. The central bank uses its private assessment of persistence to announce a wedge promising a faster return toward target. Its public forecast is part of communication policy, whether or not it is formally described as such, the early claims of transitory inflation were policy through the lens of the model. When inherited credibility is high, the announcement can lower expected inflation, dampen the initial outbreak, and subsequently reduce the real cost of disinflation.

The strategy creates a fork down the road. If the shock dissipates as the central bank expects, the promise is easy to honor;. If the shock persists, however, the bank is caught between two costly alternatives. Honoring the promised center requires a major

[1] Ihrig and Waller (2024), writing from inside the FOMC, provide a detailed account of pandemic-era guidance and the dilemmas it created in the United States. They argue that the Federal Reserve's 2020 threshold guidance for liftoff and asset purchases eventually became too constraining as inflation rose. Their account connects the two uses of communication studied here: aggressive lower-bound guidance and the subsequent attempt to look through a supply shock.

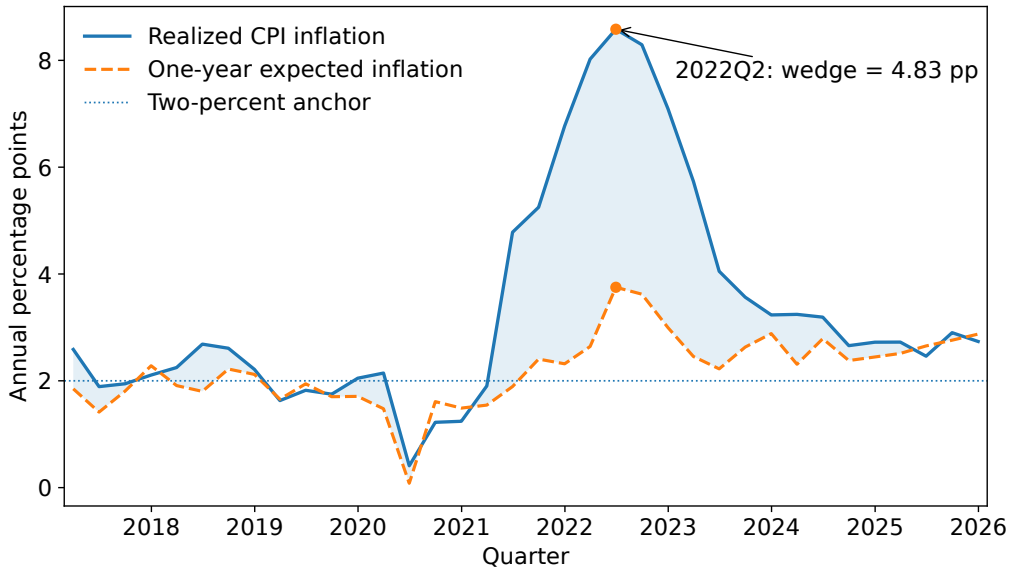

Figure 1: Inflation, one-year expectations, and the implied target weight

Notes: Quarterly CPI inflation and the Federal Reserve Bank of Cleveland's model-based measure of one-year expected inflation. Realized inflation peaked near 9 percent, whereas one-year expectations rose only modestly above 4 percent. The figure also reports the illustrative weight on the 2 percent target implied by the forecasting relation quantified in Section 5. The weight is omitted when realized inflation is too close to target for the ratio to be informative.

contraction, whereas waiving it incurs an institutional cost and damages future credibility. Even then, the earlier restraint of expectations may retain a measurable benefit by lowering the cumulative output cost of disinflation.

The outcome then depends critically on what the public believes the announcement reveals. The failed attempts to look through the supply shocks of the 1970s and the more successful post-pandemic experience illustrate two different possibilities generated by the proposed model. In the rational-expectations benchmark, the public understands the bank's policy rule, and the bank reoptimizes in response to the traction it anticipates. But this benchmark is optimistic because in practice reading what the public infers from data about the central bank's policy is difficult. We thus consider a counterfactual that holds the announced policy path fixed and exposes it to alternative public readings. These fixed-path exercises price the risk embedded in the original look-through plan, in case the bank mis-perceives public's reading of commitment strength. Nonetheless, when the public views the announcement primarily as information about shock persistence, expectations remain comparatively anchored even if the policy choice proves mistaken and over-accommodates (relative to rational expectations baseline). But when the same path is interpreted instead as evidence of unstable policy preferences, inflation expectations rise sharply and a 1970s-style outbreak emerges.

The success of a look-through strategy according to the model thus depends not only on inherited credibility, but also on how the public interprets the information contained in economic data and in the bank's past actions. The target supplies a durable anchor, but operational statements are filtered through the central bank's perceived incentives and institutional history. The post-COVID outcome is consistent with a public that placed substantial weight on the central bank's assessment of the shock, even while discounting its promise, and continued to regard its underlying policy preferences as stable. This distinction helps explain how the transitory call could fail without producing an unanchoring comparable to that of the 1970s.

The final part of our paper shows that the U.S. data was broadly consistent with the above account. One-year inflation expectations remained far below realized inflation, indicating that the standing target retained substantial influence, almost beyond rational given the duration (Figure 1). Inflation proved more persistent than initially communicated, exposing the Fed e to criticism and credibility loss, while disinflation occurred with a relatively modest increase in unemployment.[2] Consistent with the model, before the pandemic scheduled FOMC announcement days carried substantial news for short-horizon inflation compensation relative to ordinary trading days, but that traction disappeared during the 2021–23 inflation episode. The decoupling is consistent with the key prediction of the model: when communication is aggressively used to shift expectations, both marginal traction and the variance of communication wanes.

Beyond this applications, our framework has two additional implications worth pointing out. Adverse institutional evidence accumulates with the horizon and attenuates the conventional forward-guidance puzzle from the policy side. In an auxiliary local implementation benchmark, the standing target also enlarges the determinacy region: if target-based communication anchors a fraction $\bar{\tau}$ of model-consistent inflation expectations, the feedback rule need satisfy only $\phi_\pi > 1 - \bar{\tau}$, rather than the standard Taylor-principle condition $\phi_\pi > 1$.

The remainder of the paper is organized as follows. Section 2 relates the framework to the literature. Section 3 presents the inherited targeting rule, private information, institutional modes, waivers, and Bayesian learning. Section 4 derives the local Bellman value of a waiver and develops the Gaussian–logistic communication benchmark. Section 5 presents the post-pandemic cost-push application. Section 6 concludes. Proofs, numerical details, and data construction appear in the appendix.

[2]See calculations in Forbes et al. (2025).

2 Related Literature

The most direct antecedents are King et al. (2008), Lu et al. (2016), and King and Lu (2026), who study evolving reputation for commitment: the public learns from noisy inflation outcomes whether announced plans are carried out, and in King and Lu (2026) optimal policy under forward-looking expectations manages that learning strategically. The leverage coefficient of Lu et al. (2016) is the closest antecedent of the communication traction studied here. Our framework draws on their setup but differs in both the locus of learning and the commitment technology. In those papers the announcement *is* the plan and carries no information in equilibrium—the type without commitment issues the same announcement, so beliefs move only with delivered inflation. Here, guidance is evaluated against an inherited, state-contingent targeting rule, and the two hidden modes generate different message distributions. The message itself moves beliefs at issuance; this content dependence generates the finite communication frontier at the center of our analysis. Moreover, commitment is endogenous to the promise being made.[3]

The closest recent policy-design papers are Dovis and Kirpalani (2021) and Kostadinov and Roldán (2025). Dovis and Kirpalani design the rule itself when society is unsure that the policymaker can follow it; we take the existing rule as a positive feature of the environment and study communication *around* it. An important point of contact is the endogenous decision to sever the commitment bridge. Kostadinov and Roldán design an announced inflation path to discipline a policymaker without commitment, with reputation attached to a behavioral type that follows the plan mechanically. Here both modes optimize, the ordinary mode supplies the commitment, and a waiver relaxes one outstanding announcement without changing either the target or the rule. Relative to both papers, the distinctive positive contribution is the inherited rule as a message-stage benchmark and the resulting frontier for deviations from it.[4]

To give communication an independent role, we follow Bassetto (2019) and introduce private information. Under symmetric information, purely Odyssean guidance is redundant, so operational communication matters only because the central bank possesses

[3]Two classic antecedents of institutionalied delegation stressed by our paper and the related literature are Lohmann (1992) and Walsh (1995). Lohmann (1992) permits costly override of delegated policy, and Walsh (1995) considers an explicit inflation-contingent delegation contract.

[4]Related models of loose or quasi-commitment include Schaumburg and Tambalotti (2007). This literature has evolved into parallel analyses of de-anchoring. Christoffel and Farkas (2025) and Gobbi et al. (2019) model de-anchoring as a switch to a permanently depressed state. In a different vein, Choi and Foerster (2021) study optimal policy across switching growth regimes in which the inflation target itself optimally adjusts. In our setting the target is fixed by the regime; what varies is the traction of communication around it. These two views nevertheless have a point of contact: the rule is slowly shaped by ongoing communication and interaction, as it is never stated as a firm state-contingent function.

information not yet available to the public. The action-learning mechanism and our treatment of credibility draw on the seminal work of Cukierman and Meltzer (1986) and subsequent contributions such as Erceg and Levin (2003) and Amador and Phelan (2026). Because the baseline model contains no Barro–Gordon inflation bias, the inherited target does not primarily discipline an inflation-prone authority; it instead provides the public benchmark against which operational messages are evaluated.

3 Model

Time is discrete and infinite, $t = 0, 1, \ldots$. The economy comprises a private sector and a regulated central bank. The central bank maximizes social welfare but operates within an inherited inflation-targeting regime. The regime specifies a default inflation center as a function of the publicly observable state. It is not a Taylor rule or another instrument rule: $\bar{\Theta}$ maps the public state directly into the inflation center, while the nominal interest rate that implements the resulting allocation is recovered from the IS equation below. The rule may be flexible and richly state contingent; what is fixed at the operational horizon is the mapping itself.

3.1 Economy

Let s_t denote the publicly observed state of the economy. In the baseline, this state is exogenous. It contains the current cost-push and natural-rate disturbances, together with any other public variables relevant for monetary policy. The central bank additionally observes a scalar private state v_t. All agents know the exogenous state probabilistic transition laws:

$$v_t \sim Q_v(\cdot \mid s_t, v_{t-1}), \qquad s_{t+1} \sim Q_s(\cdot \mid s_t, v_t). \tag{2}$$

Let the private sector's exogenous state be summarized by $\widetilde{S}_t \equiv (s_t, v_{t-1})$and the central bank state be $S_t \equiv (s_t, v_t)$. The realization v_t becomes public before the next communication cycle. While this assumption keeps the state space tractable, the exact one-period delay in revelation is not essential for our main results.

Current cost-push and natural-rate shocks are functions of s_t: $u_t = u(s_t)$ and $r_t^n = r^n(s_t)$. The private state v_t shifts the conditional distribution of future public states. Our preferred interpretation is that it summarizes the central bank's private assessment of the persistence or future propagation of the current stabilization problem. Our analysis, however, requires only that v_t be relevant for future policy and not yet publicly observed. With a minor variation in the setup, v_t could be interpreted as a privately observed shock

to the policymaker's stabilization weight.

For the analytical and quantitative applications, we assume

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v, \qquad |\rho_v| < 1, \qquad \varepsilon_t^v \sim \mathcal{N}(0, \sigma_v^2). \tag{3}$$

Equation (3) is a specialization of Q_v in (2); the transition Q_s remains unrestricted in the general theory. The separated-episode economy is nested by $\rho_v = 0$.

Inflation π_t and the output gap x_t satisfy the standard log-linear New Keynesian equations:

$$\pi_t = \beta \widetilde{\mathbb{E}}_t \pi_{t+1} + \kappa x_t + u(s_t) + \varepsilon_t^\pi, \tag{4}$$

$$x_t = \widetilde{\mathbb{E}}_t x_{t+1} - \frac{1}{\sigma}\left[i_t - \widetilde{\mathbb{E}}_t \pi_{t+1} - r^n(s_t)\right], \tag{5}$$

where $\beta \in (0,1)$, $\kappa, \sigma > 0$ are parameters and ε_t^π is i.i.d. white noise. Per our convention, the operator $\widetilde{\mathbb{E}}_t$ denotes the private sector's expectation after observing the central bank's date-t communication.

The central bank chooses a conditional inflation-center policy θ_t. Realized inflation is subject to mean-zero control noise,

$$\pi_t = \theta_t + \varepsilon_t^\pi, \qquad \varepsilon_t^\pi \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\pi^2). \tag{6}$$

The same ε_t^π appears in the Phillips curve and in this policy representation; it is control noise, not an additional shock. Hence $\mathbb{E}\left[\pi_t \mid \theta_t, \widetilde{S}_t\right] = \theta_t$. Combining the above equation with (4) gives the conditional-mean Phillips curve

$$\theta_t = \beta \widetilde{\mathbb{E}}_t \pi_{t+1} + \kappa x_t + u(s_t). \tag{7}$$

Equation (7) is the key policy-implied constraint in the model. Conditional on the current shock, any two of the current inflation center, expected future inflation, and the output gap determine the third. If an outstanding promise fixes θ_t, an adverse cost-push shock must be absorbed either through expected future inflation or a contractionary output gap. This tension is crucial: because the central bank resolves this trade-off differently under commitment, its behavior may provide an informative signal that anchors the public's Bayesian inference of institutional mode—as we explain in the next section.

The central bank's period payoff is described by the usual LQ objective:

$$U(\theta_t, x_t) = -\frac{1}{2}\left(\alpha \theta_t^2 + \vartheta x_t^2\right), \qquad \alpha, \vartheta > 0. \tag{8}$$

Because the control noise in (6) is independent and has constant variance, $\mathbb{E}_t \pi_t^2 = \theta_t^2 + \sigma_\pi^2$. The omitted term $-\alpha\sigma_\pi^2/2$ is independent of policy, so (8) is equivalent to expected quadratic welfare in realized inflation.

When the nominal rate is unrestricted, the optimal allocation is determined jointly by (7), the welfare objective, and the communication equilibrium. The supporting nominal rate is then recovered from the IS equation:

$$i_t = r^n(s_t) + \widetilde{\mathbb{E}}_t \pi_{t+1} + \sigma \left(\widetilde{\mathbb{E}}_t x_{t+1} - x_t \right). \tag{9}$$

The IS equation governs implementation rather than the choice of the unconstrained allocation. At the effective lower bound, however, (9) becomes an implementability constraint. Assuming no sunspot equilibria, this serves as the sole implementation constraint on the central bank's problem.

3.2 Inflation-Targeting Rule

Inflation targeting is a durable communication institution. We represent it by three objects:

(i) a numerical long-run inflation target, normalized to zero;

(ii) a public, stationary, state-contingent rule $\bar{\Theta}(s_t)$ for the inflation center—the accommodation the regime itself permits—which conditions on the exogenous state s_t alone; the state may carry the history of the shocks, but the rule excludes endogenous history—the central bank's own action, policy, and communication record—and its current private information; and

(iii) an institutionalized cost of departing from the rule: an ex post cost of not delivering a shifted center, formalized below. Issuing guidance carries no separate resource cost; the ex ante discipline on the size of a shift is endogenous, operating through the credibility the central bank spends by speaking away from the rule.

The restriction on the inherited rule can be written directly as

$$\bar{\Theta}(s_t, v_t, g_{t-1}) \equiv \bar{\Theta}(s_t) \qquad \text{for every } (v_t, g_{t-1}). \tag{10}$$

The rule can therefore respond flexibly to the exogenous state—the shocks and their history—but it cannot depend on endogenous history, the bank's own action, policy, and communication record (including its prior announcements g_{t-1}), nor on its current private information. Those components must instead travel through the operational shift g_t.

Importantly, $\bar{\Theta}$ is a targeting rule for the inflation center, not a Taylor rule or another instrument rule.[5] It maps the public state directly into the center of inflation, up to the implementation noise in (6); the supporting nominal interest rate is recovered separately from (9). The institution is part reaction function, part accountability structure. Society delegates monetary policy to a benevolent central bank but constrains it: the authority retains the flexibility to respond to observable, explainable shocks, while responses to information the public cannot verify are excluded from the rule itself and must instead travel through priced guidance. The conditioning restriction in (ii) does the work. Because $\bar{\Theta}$ is measurable with respect to the current public state s, the public can compute the default center itself, so any gap between what the bank announces and $\bar{\Theta}(s_t)$ is public—this is what turns operational communication into an accountable object. Away from the lower bound, demand and productivity shocks pose no inflation–activity trade-off: by the standard divine coincidence, they are fully offset by the policy rate at zero inflation (Galí, 2015). The default rule therefore accommodates only disturbances that break that coincidence—cost-push shocks in our applications and, in richer regimes, verifiable natural-rate shifts as well. In our model the residual contingencies—a demand shortfall the constrained rate cannot absorb, and information the public cannot yet verify—are precisely what the operational layer of communication introduced in the next subsection exists to carry. The rule itself is fixed outside the model and inherited by the current authority as a time-invariant, s-measurable schedule with $\bar{\Theta}(\bar{s}) = 0$ in the normal state: whenever there is nothing verifiable to accommodate, the rule reaffirms the numerical target, so accommodation is episodic and cannot serve as a vehicle for average inflation away from it. The recursive model takes $\bar{\Theta}$ as an inherited institutional primitive. The companion paper Drozd (2026) derives a normative rule-plus-shift benchmark and clarifies the information-design problem that precedes the positive institution studied here.

3.3 Forward guidance

The flexibility built into the rule is distinct from forward guidance. Although $\bar{\Theta}$ may respond richly to the publicly observable state, it is fixed at the operational horizon and

[5]Following Svensson (2020), $\bar{\Theta}$ is a targeting rule for the inflation center, not an instrument rule for the policy rate—a schedule of the exogenous state s_t, strict when $\bar{\Theta} \equiv 0$ and flexibly accommodating verifiable shocks otherwise. The exogenous state may itself carry the history of the shocks; the operative restriction is that the rule does not depend on endogenous history—the central bank's own action, policy, and communication record—or on its current private information. This rules out forecast targeting, which adapts to the policymaker's evolving judgment, and makeup strategies such as price-level, temporary price-level, and average-inflation targeting. The standing rule is discretionary; commitment enters only through the operational shift g_t. What the rule keeps out is a crucial distinction: we posit this form here, while Drozd (2026) analyzes its deeper normative justification.

does not condition on v_t or g_{t-1}. The central bank's communication instrument is instead a scalar shift of the inherited schedule; by saying nothing, the central bank affirms exact adherence to the rule, for which it is similarly accountable ex post. That is the sense in which the target in part (i) is a durable form of communication.

The outstanding forward guidance shift issued in the preceding communication cycle is denoted by g_{t-1}, later referred to as the forward guidance shift, or simply the announcement. Let λ_t denote the public probability that the institution is operating in the override mode immediately before the date-t message. The public state observed by the private sector is therefore

$$\widetilde{\Omega}_t \equiv \left(\widetilde{S}_t, \lambda_t, g_{t-1}\right) = \left(s_t, v_{t-1}, \lambda_t, g_{t-1}\right); \tag{11}$$

and the corresponding state observed by the central bank is

$$\Omega_t \equiv \left(S_t, \lambda_t, g_{t-1}\right) = \left(s_t, v_t, \lambda_t, g_{t-1}\right). \tag{12}$$

Throughout the paper, a tilde denotes an object measurable with respect to public information. The central bank additionally observes its institutional mode and its current fixed-cost realization, which are represented separately below.

The scalar-shift restriction is motivated by the companion Ramsey benchmark in Drozd (2026). In its stationary linear–quadratic specialization, the complete continuation plan can be represented as a public state-contingent rule plus one scalar recursive commitment state. The companion paper also shows that this scalar is the decision-relevant operational statistic under a mutual-information processing cost. The present paper does not assume that observed targeting institutions implement Ramsey policy exactly. The rule is inherited and deliberately coarse, the operational shift responds to private information, and delivery is imperfectly enforceable. We study the credibility and inference consequences of those institutional restrictions.

3.4 Credibility capital

The central bank has a persistent but hidden mode of operation: $j_t \in \{\mathsf{G}, \mathsf{R}\}$. The ordinary mode G has greater internal resistance to waiving an outstanding promise ex post. The override mode R has lower resistance. Both modes share the same inflation target, default rule, welfare objective, discount factor, information structure, and policy instruments.

Conditional on mode j, the one-period commitment can be waived at a random cost:

$$F_{j,t} = \bar{F}_j(\Omega_t) + \sigma_F \nu_t, \qquad \nu_t \sim G_F, \tag{13}$$

where G_F is a continuous distribution function. We impose

$$\bar{F}_{\mathsf{G}}(\Omega) > \bar{F}_{\mathsf{R}}(\Omega) \tag{14}$$

at every relevant state. Because the modes share the same scale and normalized variation, (14) implies that ordinary-mode waiver costs first-order stochastically dominate override-mode costs. Either mode may waive; the ordinary mode simply perceives a higher cost of doing so. The cost is noisy and the public does not see its realization.

The key element that gives rise to time-varying credibility of commitment is that the mode is persistent. The timing convention is that the mode which generates a forward guidance shift remains responsible for the honor decision when that announcement comes due. Only after the promise has been resolved does the institutional mode transition. Across completed promise cycles, the mode follows a two-state Markov transition matrix

$$\mathbf{P} = (P_{jk})_{j,k\in\{\mathsf{G},\mathsf{R}\}}.$$

The applications use the parameterization

$$P_{\mathsf{GR}} = (1-\rho_\lambda)\bar{\lambda}, \qquad P_{\mathsf{RR}} = (1-\rho_\lambda)\bar{\lambda} + \rho_\lambda. \tag{15}$$

The ergodic probability of the override mode is $\bar{\lambda}$, while $1-\rho_\lambda$ is the rate at which institutional identity is renewed. The parameter ρ_λ determines how strongly current credibility affects the value of future communication.

3.5 The central bank's policy problem

At date t, the central bank issues a forward guidance shift g_t. In the next period, it chooses hidden action: honor (H) guidance or waive (W). The message is equivalent to announcing the shifted schedule as the future policy center ($t+1$-policy center)

$$s' \longmapsto \bar{\Theta}(s') + g_t. \tag{16}$$

Issuing the announcement carries no direct resource cost; the discipline on its magnitude is endogenous, operating through the public's inference and the cost of a subsequent waiver rather than through a size cost. If the newly issued announcement is honored when it becomes due, the future inflation center is

$$\Theta^H_{t+1}(s_{t+1}|g_t) = \bar{\Theta}(s_{t+1}) + g_t. \tag{17}$$

The announcement inherited at date t is g_{t-1}, and it is part of the state space. Its honored current center is $\Theta_t^H(s_t|g_{t-1}) = \bar{\Theta}(s_t) + g_{t-1}$. Issuing a waiver (W) in the next period releases this commitment, which means that the future inflation center is given by a maximization under discretion. Subsequent periods are unaffected by the waiver.

The sequence of events is as follows:

(i) The private sector observes $\widetilde{\Omega}_t$; the central bank observes Ω_t, its mode j_t, and the fixed-cost realization $F_{j_t,t}$ of issuing the waiver.

(ii) The central bank chooses whether to honor or waive the outstanding announcement and issues the new announcement g_t.

(iii) The private sector observes g_t, updates its joint belief about (j_t, v_t), and forms $\widetilde{\mathbb{E}}_t \pi_{t+1}$ and $\widetilde{\mathbb{E}}_t x_{t+1}$.

(iv) Under the honor action (H), the current inflation center is fixed at Θ_t^H; under the waiver action (W), the central bank chooses a reset center after expectations have been formed. The joint waiver program below is the reduced-form representation of this timing: expectations are predetermined when the reset center is selected.

(v) Inflation is realized according to $\pi_t = \theta_t + \varepsilon_t^\pi$; the public updates its belief about the institutional mode using the ex post inflation outcome.

(vi) The private state v_t becomes public and g_t becomes the outstanding announcement; the mode that generated g_t remains in force through its honor decision. After that decision, the mode transitions according to (15), so it persists with the corresponding transition probability.

Let $\mathcal{L}_j^g\left(g \mid \widetilde{\Omega}_t\right)$ denote the predictive density of the forward guidance shift induced by the equilibrium strategy in mode j. This density integrates over the private state, the fixed-cost draw, and the endogenous honor-or-waive decision.

After observing g_t, the public posterior probability of mode R is

$$q_t^g = \frac{\lambda_t \mathcal{L}_{\mathsf{R}}^g\left(g_t \mid \widetilde{\Omega}_t\right)}{(1-\lambda_t)\mathcal{L}_{\mathsf{G}}^g\left(g_t \mid \widetilde{\Omega}_t\right) + \lambda_t \mathcal{L}_{\mathsf{R}}^g\left(g_t \mid \widetilde{\Omega}_t\right)}. \tag{18}$$

The same observation changes the public distribution of v_t. Let

$$\mu_t(j, \mathrm{d}v) \equiv \mathbb{P}\left(j_t = j,\, v_t \in \mathrm{d}v \mid \widetilde{\Omega}_t, g_t\right) \tag{19}$$

denote the resulting joint posterior.

The probability relevant for current expectations is the probability that the newly issued announcement will be honored when it comes due. Let $\chi_j^+\left(g_t, s', v, \widetilde{\Omega}_t\right)$ denote the mode-j probability of honoring g_t in the future public state s', after integrating over the future fixed-cost draw. Because the mode that generated g_t governs honoring, prospective plan survival probability is

$$p_t\left(g_t, \widetilde{\Omega}_t\right) = \sum_{j\in\{\mathsf{G},\mathsf{R}\}} \int \int \chi_j^+\left(g_t, s', v, \widetilde{\Omega}_t\right) Q_s\left(\mathrm{d}s' \mid s_t, v\right) \mu_t(j, \mathrm{d}v). \tag{20}$$

The Markov transition matrix does not enter (20), because institutional turnover occurs after the outstanding promise has been resolved.

Let

$$h_t\left(g_t, \widetilde{\Omega}_t\right) \equiv \mathbb{E}\left[\bar{\Theta}(s_{t+1}) + g_t \mid b_{t+1} = 0, \widetilde{\Omega}_t, g_t\right] \tag{21}$$

be the expected center if the shifted rule is honored, where $b_{t+1} = 0$ denotes honor. Let

$$d_t\left(g_t, \widetilde{\Omega}_t; e\right) \equiv \mathbb{E}\left[\Theta^W_{j_t,t+1} \mid b_{t+1} = 1, \widetilde{\Omega}_t, g_t, \widetilde{\mathbb{E}}_t\pi_{t+1} = e\right] \tag{22}$$

be the expected center under waiver. The waiver center may respond to predetermined expected inflation, which is why e appears as an argument. Because implementation noise has zero mean, rationally expected inflation equals the expected future policy center:

$$e_t \equiv \widetilde{\mathbb{E}}_t\pi_{t+1} = p_t h_t + (1-p_t) d_t\left(g_t, \widetilde{\Omega}_t; e_t\right). \tag{23}$$

Define $\Phi_t(e) \equiv p_t h_t + (1-p_t) d_t\left(g_t, \widetilde{\Omega}_t; e\right)$. We impose the following contraction condition on the sensitivity of the waiver policy to predetermined expectations:

$$(1-p_t)\sup_e \left|\frac{\partial d_t(g_t, \widetilde{\Omega}_t; e)}{\partial e}\right| < 1. \tag{24}$$

Then Φ_t is a contraction and (23) has a unique solution. In the local linear representation, $d_t(g_t, \widetilde{\Omega}_t; e) = \bar{d}_t + \delta_W(e - \bar{d}_t)$, $\delta_W \in [0,1)$, condition (24) is automatically satisfied—or there is indeterminacy of equilibrium in the mode—a case we eliminate by assumption.

Let $\mathcal{L}_j^\pi\left(\pi_t \mid g_t, \widetilde{\Omega}_t\right)$ denote the equilibrium predictive density of realized inflation under institutional mode $j \in \{\mathsf{R}, \mathsf{G}\}$. This density integrates over the private state, the fixed-cost draw, the honoring decision, the mode-specific reset policy, and control noise.

In the local Gaussian representation,

$$\mathcal{L}_j^{\pi}\left(\pi_t \mid g_t, \widetilde{\Omega}_t\right) = \chi_{j,t}\phi_{\sigma_\pi}\left(\pi_t - \Theta_t^H\right) + (1-\chi_{j,t})\phi_{\sigma_\pi}\left(\pi_t - \Theta_{j,t}^W\right). \tag{25}$$

After observing inflation, the posterior probability of mode R is thus given by

$$q_t^{\pi} = \frac{q_t^g \mathcal{L}_{\mathsf{R}}^{\pi}\left(\pi_t \mid g_t, \widetilde{\Omega}_t\right)}{(1-q_t^g)\mathcal{L}_{\mathsf{G}}^{\pi}\left(\pi_t \mid g_t, \widetilde{\Omega}_t\right) + q_t^g \mathcal{L}_{\mathsf{R}}^{\pi}\left(\pi_t \mid g_t, \widetilde{\Omega}_t\right)}; \tag{26}$$

equivalently,

$$\operatorname{logit} q_t^{\pi} = \operatorname{logit}\lambda_t + \log\frac{\mathcal{L}_{\mathsf{R}}^{g}(g_t \mid \widetilde{\Omega}_t)}{\mathcal{L}_{\mathsf{G}}^{g}(g_t \mid \widetilde{\Omega}_t)} + \log\frac{\mathcal{L}_{\mathsf{R}}^{\pi}(\pi_t \mid g_t, \widetilde{\Omega}_t)}{\mathcal{L}_{\mathsf{G}}^{\pi}(\pi_t \mid g_t, \widetilde{\Omega}_t)}. \tag{27}$$

where $\operatorname{logit}\alpha \equiv \log\left[\alpha/(1-\alpha)\right]$ is the log-odds (the logarithm of the odds). The first likelihood increment is evidence from words; the second is evidence from the ex post inflation outcome. After a completed promise cycle, the posterior is mapped through the mode transition matrix. Under (15), the prior for the next communication cycle is

$$\lambda_{t+1} = (1-\rho_\lambda)\bar{\lambda} + \rho_\lambda q_t^{\pi}. \tag{28}$$

3.6 Operational policy design

Conditional on the central bank's information, the output gap implied by honor is

$$x_t^H\left(g_t; \Omega_t\right) = \frac{\Theta_t^H - \beta e_t(g_t, \widetilde{\Omega}_t) - u(s_t)}{\kappa}. \tag{29}$$

The output gap implied by a waiver policy θ_t is

$$x_t^W\left(\theta_t, g_t; \Omega_t\right) = \frac{\theta_t - \beta e_t(g_t, \widetilde{\Omega}_t) - u(s_t)}{\kappa}. \tag{30}$$

Let $\mathcal{V}_j(\widetilde{\Omega}_t)$ denote the mode-j continuation payoff before the current private state and waiver cost have been observed. The honor branch takes the inherited center Θ_t^H as given and chooses the new forward guidance shift, implying the Bellman equation of the form:

$$\mathcal{V}_j^H(\Omega_t) = \max_{g_t}\left\{U\left(\Theta_t^H(s_t|g_{t-1}), x_t^H(g_t; \Omega_t)\right) + \beta\mathbb{E}_t^j\mathcal{V}_{j_{t+1}}\left(\widetilde{\Omega}_{t+1}\right)\right\}. \tag{31}$$

Because a waiver releases the inherited center by definition, in that case we have

$$\mathcal{V}_j^W(\Omega_t) = \max_{\theta_t, g_t} \left\{ U\left(\theta_t, x_t^W(\theta_t, g_t; \Omega_t)\right) + \beta \mathbb{E}_t^j \mathcal{V}_{j_{t+1}}\left(\widetilde{\Omega}_{t+1}\right) \right\}. \tag{32}$$

The expectation in (31)–(32) integrates over the public-state transition, control noise, Bayesian belief updating, and institutional turnover after the promise has been resolved. The next public state is

$$\widetilde{\Omega}_{t+1} = \left(\widetilde{S}_{t+1}, \lambda_{t+1}, g_t\right) = \left(s_{t+1}, v_t, \lambda_{t+1}, g_t\right), \tag{33}$$

which, note, depends on g_t chosen in the current period by the central bank. Before observing v_t and $F_{j,t}$, the mode-j continuation payoff is

$$\mathcal{V}_j(\widetilde{\Omega}_t) = \mathbb{E}\left[\max\left\{ \mathcal{V}_j^H(\Omega_t), \mathcal{V}_j^W(\Omega_t) - F_{j,t} \right\} \middle| \widetilde{\Omega}_t, j \right]. \tag{34}$$

Define the gross payoff gain from waiving before paying the fixed cost:

$$\Delta\mathcal{V}_j(\Omega_t) \equiv \mathcal{V}_j^W(\Omega_t) - \mathcal{V}_j^H(\Omega_t). \tag{35}$$

The institution honors if and only if

$$F_{j,t} \geq \Delta\mathcal{V}_j(\Omega_t). \tag{36}$$

The mode-j honoring probability before the fixed-cost realization is therefore

$$\chi_j(\Omega_t) = \mathbb{P}\left(F_{j,t} \geq \Delta\mathcal{V}_j(\Omega_t) \mid \Omega_t, j\right), \tag{37}$$

which, under (13), gives

$$\chi_j(\Omega_t) = 1 - G_F\left(\frac{\Delta\mathcal{V}_j(\Omega_t) - \bar{F}_j(\Omega_t)}{\sigma_F}\right). \tag{38}$$

The outstanding announcement g_{t-1} restricts the current center under honor and is a state variable. The central bank does not commit at date zero to an unrestricted sequence of future announcements. Such ex ante commitment could alter its choices because it would be made behind a veil of ignorance. We conclude by defining the equilibrium.

Definition 1 (Equilibrium)**.** *A recursive (Markov) equilibrium consists of mode-specific guidance strategies, honor and waiver policies, private-sector forecasts, message and inflation likelihoods,*

Bayesian beliefs, value functions, and an interest-rate implementation rule such that: i) the branch policies solve (31) and (32); ii) the branch decision satisfies (34), and honoring probabilities satisfy (38); iii) private expectations satisfy (23); iv) message beliefs satisfy (18); v) inflation-outcome beliefs satisfy (26); vi) the next prior satisfies (28); vii) the likelihoods $\mathcal{L}_j^g$ and $\mathcal{L}_j^\pi$ are generated by the equilibrium guidance and implementation strategies; viii) under honor, the current center equals $\bar{\Theta}(s_t) + g_{t-1}$; ix) the IS equation implements the selected allocation and respects the effective lower bound whenever it binds.

4 Equilibrium Properties and Results

The recursive equilibrium in Section 3 does not admit a useful global closed form. The reason is substantive: whether an inherited promise is honored depends on the value of relaxing it, and that value is obtained only after the central bank has reoptimized both its current inflation center and its new announcement, taking into account the belief state and the promise carried into the next period. We therefore proceed in two steps. We first derive a second-order representation of this optimized Bellman value around the target equilibrium. We then combine that local Bellman object with Gaussian message likelihoods and logistic waiver costs to obtain a closed-form communication benchmark. The frontier results below are exact in that approximate economy. Restrictions needed only for the closed form are stated when they are introduced; they are not imposed on the recursive model itself.

To simplify timing, this section suppresses date subscripts. Write

$$\Omega = (S, \lambda, g_{-1}) \qquad \text{and} \qquad \widetilde{\Omega}' = (\widetilde{S}', \lambda', g),$$

where $S = (s, v)$ is the central bank's current macroeconomic and information state, $\widetilde{S}' = (s', v)$ is the public state at the start of the next communication cycle, g_{-1} is the inherited announcement, and g is the announcement chosen now. Thus the same g becomes next period's inherited announcement. This is the undated counterpart of (12) and (33).

4.1 The local value of a waiver

Let $W_j = -\mathcal{V}_j$ denote loss rather than payoff, and let $\mathcal{J}_j(\theta, g; \Omega)$ denote the current loss plus the expected continuation loss generated by current center θ and new announcement g. It includes the effects of both choices on private expectations, realized inflation, the posterior λ', and every component of $\widetilde{\Omega}'$. The two Bellman branches can then be written

as

$$W_j^H(\Omega) = \min_g \mathcal{J}_j\big(\bar{\Theta}(s) + g_{-1}, g; \Omega\big)\,, \tag{39}$$

$$W_j^W(\Omega) = \min_{\theta,g} \mathcal{J}_j(\theta, g; \Omega). \tag{40}$$

Thus both branches reoptimize the new announcement, while the waiver branch also reoptimizes the current center. The gross payoff gain from waiving is equivalently

$$\Delta\mathcal{V}_j(\Omega) = W_j^H(\Omega) - W_j^W(\Omega) \geq 0. \tag{41}$$

Consider a regular target reference state $\bar{\Omega} = (\bar{S}, \bar{\lambda}, 0)$ at which the honored center coincides with the unconstrained center and the optimal new announcement is zero. Assume that $\mathcal{J}_j$ is twice continuously differentiable in a neighborhood of this state, that both branch solutions are locally interior, and that its Hessian with respect to (θ, g) is positive definite. The next proposition is the key local Bellman result.

Proposition 1 (The local value of a waiver)**.** *For inherited announcements g_{-1} near zero, the optimized waiver center and the gross waiver gain satisfy*

$$\theta_j^W(g_{-1}) = \bar{\theta} + \omega_j g_{-1} + o(g_{-1}), \tag{42}$$

$$\Delta\mathcal{V}_j(g_{-1}) = \frac{1}{2}\Psi_j g_{-1}^2 + o(g_{-1}^2), \qquad \Psi_j = \mathcal{A}_j(1-\omega_j)^2 \geq 0. \tag{43}$$

Here $\mathcal{A}_j > 0$ is the curvature of the continuation-adjusted loss in the reset direction after the new announcement has been optimally adjusted, and ω_j is the local pass-through of the inherited promise into the waiver center. The full derivative formula for Ψ_j, which allows the honor and waiver branches to choose different new announcements, is given in Appendix 7.1.

The gap in Proposition 1 attains a local minimum at the target because a waiver is optional. Honoring keeps the inherited center and chooses the best new announcement; waiving solves the same problem with the added right to reset the center, so the extra option cannot raise loss and $\Delta\mathcal{V}_j(g_{-1}) = W_j^H - W_j^W \geq 0$ for every inherited announcement. At $g_{-1} = 0$ there is nothing to waive: the two branches coincide, the gap vanishes, and by the envelope theorem its first derivative vanishes as well. A smooth nonnegative function that is flat at the origin carries no linear term, so the leading behavior is the quadratic $\frac{1}{2}\Psi_j g_{-1}^2$ and (43) holds with an $o(g_{-1}^2)$ remainder. Nonnegativity fixes the sign and removes the linear term, but it does not by itself exclude a cubic. When the target is symmetric, so that the continuation-adjusted loss is even in g_{-1}—as it is under the

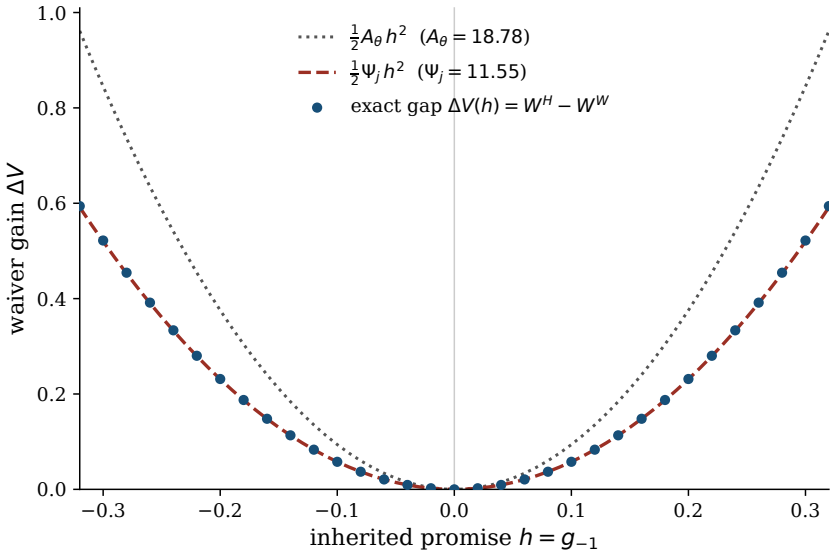

Figure 2: The local waiver gain is quadratic in the inherited promise.

Notes: Local waiver gain $\Delta\mathcal{V}(h) = W_j^H(h) - W_j^W(h)$ at the target reference state, on the post-pandemic calibration of Section 5, computed by nested minimization of the continuation-adjusted objective (the next-period promise priced at the fixed-point curvature): the honor branch locks the current center at $\theta = \bar{\theta} + h$ and reoptimizes the announcement, while the waiver branch reoptimizes both and pays the reputational tether. The exact gap (markers) lies below the period-loss parabola $\frac{1}{2}A_\theta h^2$ ($A_\theta = 18.78$, dotted) and closely tracks the quadratic $\frac{1}{2}\Psi_j h^2$ at the optimized curvature $\Psi_j = 11.55$ (dashed), up to a small higher-order residual.

quadratic Bellman benchmark used below for the closed form—the odd terms vanish and the remainder sharpens to $O(g_{-1}^4)$; on that benchmark the gap is exactly $\frac{1}{2}\Psi_j g_{-1}^2$. This is the linear–quadratic approximation the paper uses throughout. Figure 2 shows that it is accurate on the post-pandemic calibration of Section 5: the exact nested-minimization gap stays close to $\frac{1}{2}\Psi_j g_{-1}^2$ across the relevant range of inherited promises—up to a small higher-order residual—at the optimized curvature Ψ_j rather than the steeper period-loss curvature A_θ.

The result isolates the economically relevant notion of validation cost. A promise is costly to honor only to the extent that a waiver would undo it after all other controls have been reoptimized. If $\omega_j = 1$, the waiver center tracks the inherited promise one for one and the promise creates no local temptation to waive. If it is partly reversed, the gain from waiver is quadratic, with curvature Ψ_j. This curvature is a Bellman object: it contains current stabilization costs, endogenous belief updating, and continuation values. The familiar current-loss coefficient $\alpha + \vartheta/\kappa^2$ is recovered only in the static special case in which the new announcement and continuation value are irrelevant.

The maximum over honor and waiver in (34) also has an exact envelope representation.

At a fixed Ω, let $\widehat{W}_j(\Omega)$ denote loss after the private state is known but before the fixed cost is drawn. Holding the exogenous distribution of F_j locally fixed with respect to the inherited announcement, define

$$\mathcal{D}_j(g_{-1}) \equiv \mathbb{E}_F \min\{\Delta\mathcal{V}_j(g_{-1}), F_j\}.$$

Then

$$\widehat{W}_j = W_j^W + \mathcal{D}_j, \qquad \mathcal{D}_j''(0) = \chi_{j0}\Psi_j, \qquad \chi_{j0} \equiv \mathbb{P}(F_j \geq 0), \tag{44}$$

The logistic waiver-cost draw has full support and can therefore be negative, so χ_{j0} need not equal one. The pre-draw loss in the recursive problem is the conditional expectation of $\widehat{W}_j(\Omega)$ over the private state. We refer to

$$\Xi_j \equiv \chi_{j0}\Psi_j \tag{45}$$

as the expected implementation curvature. It is the local cost anticipated at issuance: the optimized cost of validating a marginal promise, multiplied by the probability that the institution will actually bear it. When the delivery state is random, the Ξ_j entering the issuance problem is the corresponding conditional expectation over that state; in the fixed-state benchmark below it is exactly (45).

4.2 Endogenous caution and message likelihoods

The Bellman curvature affects what each mode is willing to say. Let $\widetilde{v} = v - \mathbb{E}(v \mid \widetilde{S})$ be the private stabilization innovation. A second-order expansion of the mode-j issuance problem can be written as

$$J_j(g, \widetilde{v}) = J_{j0} - \zeta_j g\widetilde{v} + \frac{1}{2}(\gamma_j + \beta\Xi_j)g^2 + o(g^2 + \widetilde{v}^2), \tag{46}$$

where ζ_j is the marginal stabilization value of private information and $\gamma_j > 0$ collects all remaining local curvature—current stabilization and continuation effects common to the two branches—after the incremental honor–waiver implementation term has been separated out. Hence

$$g_j(\widetilde{v}) = K_j\widetilde{v} + o(|\widetilde{v}|), \qquad K_j = \frac{\zeta_j}{\gamma_j + \beta\Xi_j}. \tag{47}$$

Equation (47) clarifies what is, and is not, implied by the model. A mode that expects a larger implementation burden responds less to the same private stabilization signal, other

things equal. An unconditional ordering of message slopes does not follow merely from Gaussianity or from the ordering of fixed waiver costs, because other local derivatives can also differ by mode. For the analytical benchmark below, we impose the transparent local symmetry restriction $\zeta_G = \zeta_R > 0$ and $\gamma_G = \gamma_R > 0$. Under this restriction,

$$\Xi_G > \Xi_R \quad \Longrightarrow \quad 0 < K_G < K_R. \tag{48}$$

A sufficient specialization is common Bellman curvature $\Psi_G = \Psi_R$ together with $\chi_{G,0} > \chi_{R,0}$, as implied by the ordered waiver-cost distributions.

For use below, write the local policy rule with additive communication noise as

$$g_{j,t} = \mu_j(\widetilde{\Omega}_t) + K_j(\widetilde{\Omega}_t)\widetilde{v}_t + \varepsilon^g_{j,t}. \tag{49}$$

Lemma 1 (Policy-generated local likelihoods)**.** *Suppose* $\widetilde{v}_t \mid \widetilde{S}_t \sim \mathcal{N}(0, \sigma_v^2)$ *and* $\varepsilon^g_{j,t} \sim \mathcal{N}(0, \sigma^2_{g,j})$ *independently. Then*

$$g_{j,t} \mid j, \widetilde{\Omega}_t \sim \mathcal{N}\Big(\mu_j(\widetilde{\Omega}_t), \Sigma_j^2(\widetilde{\Omega}_t)\Big), \qquad \Sigma_j^2 = K_j^2\sigma_v^2 + \sigma^2_{g,j}. \tag{50}$$

If the local intercepts and communication-noise variances are common across modes, the caution ordering in (48) *implies* $\Sigma_G < \Sigma_R$.

Under these local symmetry restrictions, the economic mechanism is simple. The ordinary mode expects to be held to its words more often, so an aggressive announcement is a more expensive response to private urgency. It therefore generates a tighter local message distribution. The target then serves as a statistical sieve: a sufficiently unusual departure is more likely to have come from the override mode.

4.3 The communication frontier

We now impose four additional restrictions to obtain a closed form. First, messages are measured relative to a common local mean and the Gaussian likelihoods in Lemma 1 are extended over the analytical range, with $\Sigma_G < \Sigma_R$. Second, the override mode supplies no incremental commitment to an off-rule announcement. Third, conditional on the ordinary mode, the local Bellman approximation $\Delta\mathcal{V}_G(g) = \Psi_G g^2/2$ is extended over that range and the fixed waiver cost is logistic with constant location $\bar{F}_G$. Fourth, the honored incremental center is g and the waiver fallback is zero. These restrictions define a second-order Bellman–Gaussian benchmark; the results in this subsection are exact within it. We state the frontier on the positive branch $g \geq 0$; the centered benchmark is symmetric on

the negative branch.

Bayes' rule and the logistic waiver decision give

$$q(g) = \frac{C_\lambda e^{\eta_g g^2}}{1 + C_\lambda e^{\eta_g g^2}}, \qquad C_\lambda \equiv \frac{\lambda}{1-\lambda}\frac{\Sigma_{\mathsf{G}}}{\Sigma_{\mathsf{R}}}, \qquad \eta_g \equiv \frac{1}{2}\left(\Sigma_{\mathsf{G}}^{-2} - \Sigma_{\mathsf{R}}^{-2}\right) > 0, \tag{51}$$

$$\chi(g) = \frac{1}{1 + D_F e^{\kappa_F g^2}}, \qquad D_F \equiv e^{-\bar{F}_{\mathsf{G}}/\sigma_F}, \qquad \kappa_F \equiv \frac{\Psi_{\mathsf{G}}}{2\sigma_F} > 0. \tag{52}$$

Here $q(g)$ is the posterior probability of the override mode after the message and $\chi(g)$ is the ordinary mode's probability of honoring it when it comes due. A plan survives only if it was issued under the ordinary mode and is subsequently honored:

$$p(g) = [1 - q(g)]\chi(g). \tag{53}$$

The implemented shift in expected inflation is therefore $m(g) = gp(g)$.

Proposition 2 (The communication frontier)**.** *In the second-order Bellman–Gaussian benchmark,*

$$m(g) = \frac{g}{[1 + C_\lambda e^{\eta_g g^2}][1 + D_F e^{\kappa_F g^2}]} \tag{54}$$

and

$$m_g(g) = p(g)\left[1 - \underbrace{2\eta_g g^2 q(g)}_{\textit{inference}} - \underbrace{2\kappa_F g^2[1 - \chi(g)]}_{\textit{temptation}}\right]. \tag{55}$$

There is a unique $\bar{g} > 0$ such that m rises on $(0, \bar{g})$, is maximized at $\bar{g}$, and declines thereafter to zero. The frontier is characterized by

$$2\bar{g}^2\left\{\eta_g q(\bar{g}) + \kappa_F[1 - \chi(\bar{g})]\right\} = 1. \tag{56}$$

Holding the other benchmark objects fixed, $\bar{g}$ decreases with prior doubt C_λ, the informativeness of extremity η_g, baseline waiver risk D_F, and Bellman temptation κ_F.

The direct marginal benefit of a larger literal promise is normalized to one in (55). It is offset by two endogenous wedges. The inference wedge is the loss of ordinary-mode probability caused by the announcement itself. The temptation wedge is the increase in the ordinary mode's future waiver probability, whose curvature is the optimized Bellman coefficient Ψ_{G}. Both wedges rise with the announcement, so they cross the direct benefit exactly once. This is the Laffer curve of communication: beyond $\bar{g}$, saying more shifts expectations by less.

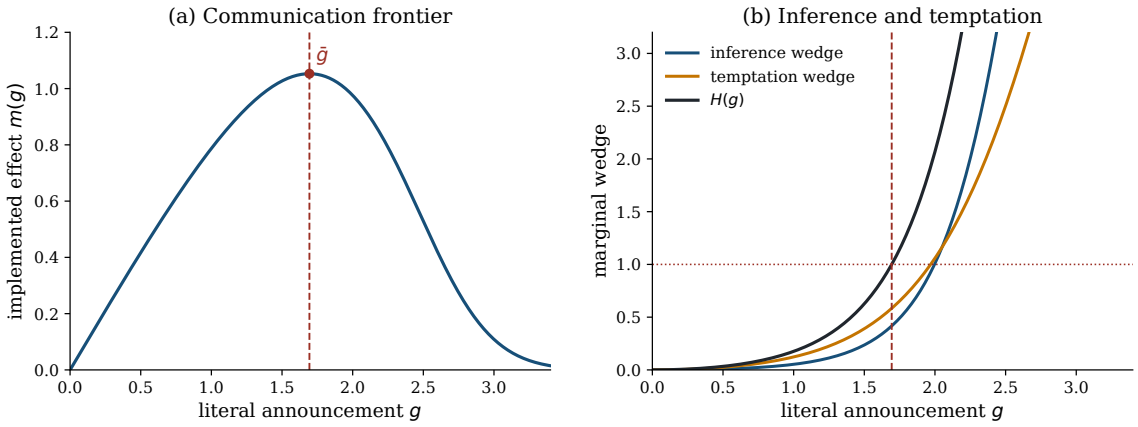

Figure 3: The communication frontier and its two margins

Notes: Panel (a) plots the implemented expectation shift $m(g)$. Panel (b) plots the inference wedge, the temptation wedge, and their sum. The vertical line marks the unique announcement $\bar{g}$ at which the combined marginal wedge equals one.

Corollary 1 (The frontier bounds optimal speech)**.** *Consider an authority that chooses $g \geq 0$ to maximize $B(m(g); s) - A(g)$, where $B_m > 0$ and A is nondecreasing. Every announcement to the right of $\bar{g}$ is weakly dominated by a smaller announcement to its left. Hence an optimum can always be selected in $[0, \bar{g}]$; if A is strictly increasing, every optimum lies in that interval. If, in addition, A is differentiable and $A_g(\bar{g}) > 0$, every positive optimum lies strictly below $\bar{g}$. In the linear-benefit problem $sm(g) - A(g)$, suppose A is continuously differentiable with finite derivative on $[0, \bar{g}]$. Any sequence of interior optima selected on the effective branch $[0, \bar{g}]$ converges to $\bar{g}$ from below as $s \to \infty$.*

A more severe stabilization problem therefore pushes policy toward the frontier but does not move the frontier. The central bank cannot manufacture additional traction by making the promise still larger once the inference and temptation terms have exhausted its marginal effect.

The same logic disciplines distant forward guidance.

Corollary 2 (Distant forward guidance)**.** *Let $\ell^g(g) = \log[\mathcal{L}^g_{\mathsf{R}}(g)/\mathcal{L}^g_{\mathsf{G}}(g)] > 0$ be the adverse log-likelihood increment generated by an extreme repeated message. Along a fixed-mode path on which n conditionally independent message increments add, suppose the raw New Keynesian horizon multiplier is bounded by $\bar{M}e^{\gamma n}$. If $\ell^g(g) > \gamma$, the credibility-adjusted effect of the distant promise converges to zero as $n \to \infty$.*

The result is deliberately conditional on the fixed-mode path. It says that sufficiently extreme guidance cannot be rescued merely by moving it farther into the future: adverse institutional evidence compounds geometrically and can dominate the conventional horizon effect.

4.4 Credibility from words and actions

The Gaussian sieve distinguishes two notions that are easily conflated. A message can improve credibility relative to the prior, yet a marginal increase in that same message can still damage credibility relative to a slightly smaller announcement.

Proposition 3 (Impact credibility and marginal spending)**.** *In the Gaussian benchmark, the announcement lowers the posterior probability of the override mode if and only if $|g| < g_I$, where*

$$g_I^2 = \frac{2\log(\Sigma_{\mathsf{R}}/\Sigma_{\mathsf{G}})}{\Sigma_{\mathsf{G}}^{-2} - \Sigma_{\mathsf{R}}^{-2}}. \tag{57}$$

For every $g > 0$, however,

$$q_g(g) = 2\eta_g g q(g)[1 - q(g)] > 0. \tag{58}$$

Hence a small positive announcement may build credibility in levels while making it larger always shifts beliefs toward the override mode at the margin.

On the effective branch $g < \bar{g}$, the institutional cost per unit of implemented expectation management is

$$\frac{\mathrm{d}q}{\mathrm{d}m} = \frac{2\eta_g g q(g)}{\chi(g)[1 - H(g)]}, \tag{59}$$

$$\frac{\mathrm{d}[1-\chi(g)]}{\mathrm{d}m} = \frac{2\kappa_F g[1-\chi(g)]}{[1-q(g)][1-H(g)]}, \tag{60}$$

where $H(g) = 2g^2\{\eta_g q(g) + \kappa_F[1 - \chi(g)]\}$. Both expressions diverge as $g \uparrow \bar{g}$. Near the frontier, public doubt and waiver risk continue to react to a larger announcement while its marginal effect on expectations vanishes. Diminishing traction therefore makes the credibility spent per unit of achieved stabilization arbitrarily large.

Words are only the first signal. When the promise comes due, noisy inflation outcomes reveal whether the institution acted as a mode with high or low resistance to waiver. Let Θ^H be the honored center and let a waiver instead deliver a common local reset center θ. Conditional on mode j, the outcome density is

$$\mathcal{L}_j^{\pi}(y) = \chi_j \phi_{\sigma_\pi}(y - \Theta^H) + (1 - \chi_j)\phi_{\sigma_\pi}(y - \theta), \qquad \chi_{\mathsf{G}} > \chi_{\mathsf{R}}. \tag{61}$$

Proposition 4 (Belief damage from a delivery failure)**.** *Let q be the override posterior before inflation is observed, $\bar{\chi}_q = (1-q)\chi_{\mathsf{G}} + q\chi_{\mathsf{R}}$, and $\Delta_\theta = \theta - \Theta^H$. Conditional on an actual waiver,*

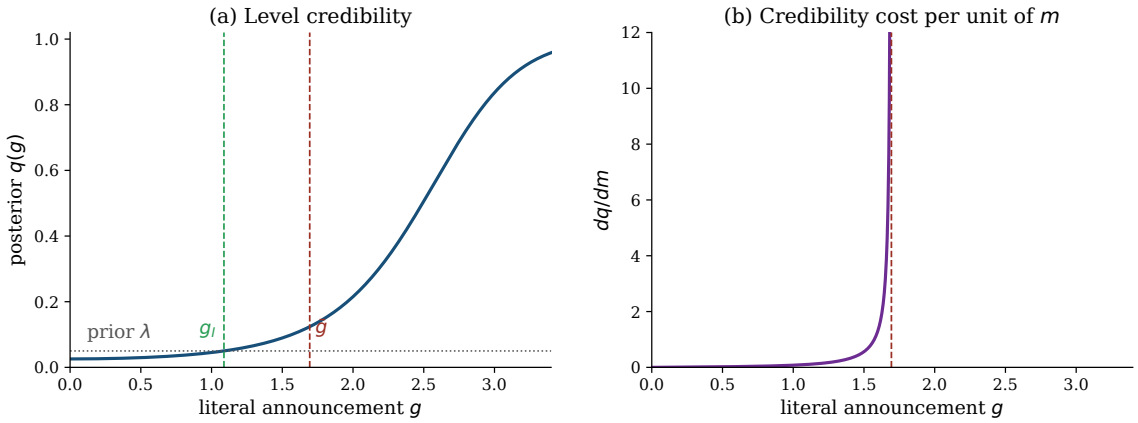

Figure 4: Credibility in levels and the institutional price of traction

Notes: Panel (a) distinguishes the level threshold g_I, below which the message lowers override beliefs relative to the prior, from the communication frontier $\bar{g}$. Panel (b) plots the marginal credibility cost per unit of implemented guidance, which diverges as the frontier is approached.

$y \sim \mathcal{N}(\theta, \sigma_\pi^2)$ *and, as* $\Delta_\theta \to 0$,

$$\mathbb{E}_\theta\left[\log \frac{\mathcal{L}_{\mathsf{R}}^\pi(y)}{\mathcal{L}_{\mathsf{G}}^\pi(y)}\right] = \frac{\chi_{\mathsf{G}}^2 - \chi_{\mathsf{R}}^2}{2\sigma_\pi^2}\Delta_\theta^2 + o(\Delta_\theta^2), \tag{62}$$

$$\mathbb{E}_\theta(q^\pi - q) = \frac{q(1-q)(\chi_{\mathsf{G}} - \chi_{\mathsf{R}})\bar{\chi}_q}{\sigma_\pi^2}\Delta_\theta^2 + o(\Delta_\theta^2). \tag{63}$$

The second expression is positive for every sufficiently small nonzero delivery miss. Consequently, the transition in (28) *carries the same second-order increase in public doubt into the next prior, scaled by* ρ_λ.

The proposition is purely Bayesian: a larger delivery miss is more diagnostic of the mode that waives more often. Turning this belief damage into a policy tether requires an additional value-function restriction. The following conditional result makes that restriction explicit rather than embedding it in the likelihood calculation.

Lemma 2 (Tethered reset under a reputational value restriction)**.** *Let* $W_j^c(\lambda')$ *denote continuation loss after applying the institutional transition, holding the other next-state objects fixed. Suppose* W_j^c *is twice continuously differentiable at the no-miss next prior. Suppose also that* $\rho_\lambda > 0$ *and*

$$W_{j,\lambda}^c > 0, \qquad W_{j,\lambda\lambda}^c \geq 0. \tag{64}$$

Then the expected continuation loss following a small waiver contains a reputational penalty

$$\frac{\varrho_j}{2}(\theta - \Theta^H)^2, \qquad \varrho_j > 0, \tag{65}$$

where the exact local coefficient is given in (136). *Let*

$$A_\theta \equiv \alpha + \frac{\vartheta}{\kappa^2}, \qquad d(e,u) \equiv \frac{\vartheta}{\alpha\kappa^2+\vartheta}(\beta e + u) \tag{66}$$

be the curvature and minimizer of current loss in the inflation center. The unique local waiver policy is

$$\Theta_j^W = \frac{A_\theta d(e,u) + \varrho_j \Theta^H}{A_\theta + \varrho_j}. \tag{67}$$

Proposition 1 describes the local reset without signing ω_j. Under (64), the reputational curvature is positive and the two representations coincide locally: $\omega_j = \varrho_j/(A_\theta + \varrho_j)$. Thus the weighted-average formula below is a signed specialization of the general optimized reset, not a second reset theory.

When credibility has no continuation value, $\varrho_j = 0$ and the central bank selects the static discretionary center. When credibility is valuable, the reset remains tethered to the promise even after the formal commitment is waived. The first condition in (64) says that greater public doubt raises future loss; the second ensures that the additional dispersion of posterior beliefs created by a noisy delivery miss does not reverse that force. These are state-dependent value restrictions. Without them, Proposition 4 remains valid but the policy-tether conclusion need not follow.

4.5 Institutional fragility

The logistic specification smooths the honor decision. As the dispersion of the fixed cost falls, this smooth decision converges to a cliff whose location is pinned down by the optimized Bellman curvature, not by current loss alone.

Proposition 5 (Low-noise waiver cliff)**.** *Hold* $\Psi_G, \bar{F}_G > 0$*, and the Gaussian message coefficients fixed while* $\sigma_F \downarrow 0$*. Define*

$$g^\star \equiv \sqrt{\frac{2\bar{F}_G}{\Psi_G}}. \tag{68}$$

Then $\chi(g) \to \mathbf{1}\{|g| < g^\star\}$ *at every continuity point. Let* $\bar{g}_I$ *be the unique frontier in the inference-only economy,* $2\eta_g \bar{g}_I^2 q(\bar{g}_I) = 1$*. The full communication frontier satisfies*

$$\lim_{\sigma_F \downarrow 0} \bar{g}(\sigma_F) = \min\{\bar{g}_I, g^\star\}. \tag{69}$$

Institutional precision is therefore double edged. Below the threshold, low noise makes honor highly predictable. At the threshold, the same precision turns a small increase in

the promise into an order-one collapse in survival. Inference limits communication first when $\bar{g}_I < g^\star$; otherwise the Bellman waiver threshold does.

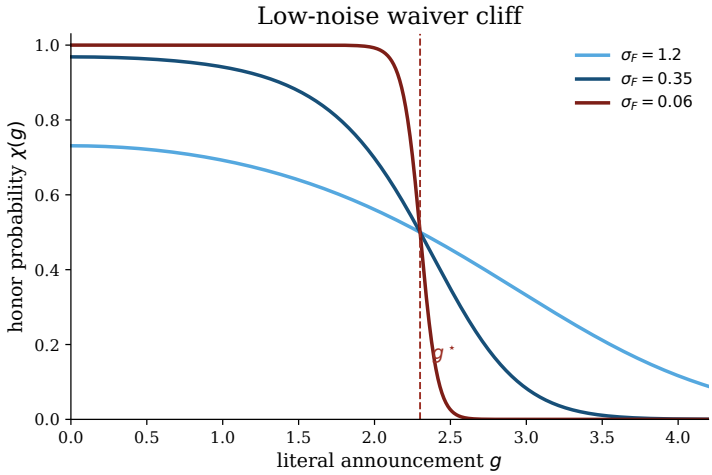

Figure 5: The low-noise waiver cliff

Notes: The ordinary-mode honor probability for several values of the fixed-cost dispersion σ_F. As institutional noise vanishes, the smooth decision converges to a threshold at $g^\star$.

4.6 Speech, interest rates, and implementation

The frontier describes what speech can achieve once it is available. A separate institutional question is whether society would ever commit the central bank not to speak before a future cost-push shock is observed. With aligned objectives and a fixed communication technology, the answer is no: allowing speech weakly expands the feasible set because $g = 0$ remains available. A ban can be valuable only if opening the speech margin changes the delegated authority's incentives or the equilibrium communication technology. Throughout, $g = 0$ should be understood as reaffirming the target, not as a literal absence of communication: the target remains the standing message of the model, and what a ban removes is only the discretionary announcement g around it.

The following auxiliary delegation problem isolates the first possibility. The current interest rate implements x, speech shifts inflation by ξg with $\xi > 0$, and

$$\pi = u + \kappa x + \xi g. \tag{70}$$

After observing u, the delegated central bank minimizes

$$L^A = \frac{1}{2}\left[\alpha\pi^2 + \vartheta(x - \bar{x})^2 + \Gamma_s g^2\right], \tag{71}$$

where $\bar{x} \geq 0$ is its delegated output target and $\Gamma_s > 0$ is the reduced local curvature of the communication margin, including expected implementation and equilibrium effects. It is not a separate primitive guidance-size cost in the baseline model. The ex-ante principal evaluates the same expression with $\bar{x} = 0$ in the output term. Before u is observed, the principal chooses whether to allow g or impose $g = 0$. This is an auxiliary local delegation result, not a claim about the unrestricted signaling game. The companion paper Drozd (2026) studies the broader information-design problem behind access to operational communication.

Proposition 6 (When is ex-ante silence desirable?)**.** *Let*

$$D = 1 + \frac{\alpha\kappa^2}{\vartheta}, \qquad h_s = \frac{\alpha\xi^2}{\Gamma_s},$$

where D is the static stabilization denominator and h_s is the local speech curvature. The principal's state-by-state loss difference between allowing speech and imposing silence is

$$L^P_{\text{allow}}(u) - L^P_{\text{silent}}(u) = -\frac{\alpha h_s}{2D(D + h_s)} \left(u^2 - \kappa^2\bar{x}^2\right). \tag{72}$$

Hence the principal strictly prefers to allow speech when

$$\mathbb{E}(u^2) > \kappa^2\bar{x}^2, \tag{73}$$

strictly prefers silence when the inequality is reversed, and is indifferent at equality. With aligned objectives ($\bar{x} = 0$), a ban is never strictly optimal for a nondegenerate cost-push shock.

The criterion has a direct interpretation. Speech is worth preserving when the prospective stabilization problem is large. A ban can dominate only when the delegated central bank's incentive to use communication for excessive accommodation is large relative to the shocks for which communication would be socially useful. This is the relevant trade-off for committing not to deploy a look-through strategy before the supply shock is known.

4.7 Communication and the interest-rate margin

Conditional on allowing speech, a cost-push disturbance generally calls for both the current rate and guidance. Let local traction be $\bar{\tau} > 0$ and consider

$$\pi = u + \kappa x + \beta\bar{\tau}g. \tag{74}$$

Proposition 7 (Communication and the interest-rate margin)**.** *For the problem*

$$\min_{x,g} \frac{1}{2}\left(\alpha\pi^2 + \vartheta x^2\right) \quad \textit{subject to (74)}, \tag{75}$$

the first-order conditions are

$$\vartheta x = -\alpha\kappa\pi, \qquad \alpha\beta\bar{\tau}\pi = 0. \tag{76}$$

Because guidance carries no resource cost, the announcement alone offsets a cost-push shock: $\pi = 0$ *and* $x = 0$ *at the optimum, with* $g = -u/(\beta\bar{\tau})$*. The current rate is reserved for natural-rate disturbances, which enter the IS equation rather than the Phillips curve and cannot be addressed by guidance.*

This result concerns states in which the current center is free. A contemporaneous rate rule and a promise act on different margins: the rate moves the output gap, whereas the announcement moves the expectations term. If a nonzero inherited announcement is honored, the center is fixed at $\bar{\Theta}(s) + g_{-1}$ and the output gap must absorb the residual in (7). If the promise is waived, the conditional reset in Lemma 2 applies when its value restriction holds.

4.8 Decentralized implementation

The IS curve is not an additional constraint on the unconstrained institutional allocation; it determines the interest rate that implements it. To state a determinacy result, one must specify how the target anchors model-consistent forecasts. Consider the auxiliary local implementation system

$$\pi_t = \beta\delta\mathbb{E}_t\pi_{t+1} + \kappa x_t, \tag{77}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left[i_t - \delta\mathbb{E}_t\pi_{t+1}\right], \tag{78}$$

where $\delta \equiv 1 - \bar{\tau} \in (0, 1]$ is the unanchored share of the private inflation forecast. The same share enters the Phillips curve and the real-rate term.

Proposition 8 (Credibility and local determinacy)**.** *Under the feedback rule* $i_t = \phi_\pi\pi_t$ *with* $\phi_\pi \geq 0$*, the target equilibrium of (77)–(78) is locally determinate if and only if* $\phi_\pi > 1 - \bar{\tau}$*. At equality one eigenvalue is unity. The condition extends by continuity as* $\bar{\tau} \to 1$*.*

The result is an implementation boundary, not a restriction imposed on the planner's allocation. Strong target traction substitutes for part of the contemporaneous rate response

needed to exclude self-fulfilling fluctuations. As traction vanishes, the standard Taylor principle is recovered.

The economics of the relaxed boundary is best seen through the real rate. The Taylor principle requires that the real rate rise when inflation drifts; by Fisher equation, and perfect foresight, the nominal rate must respond more than one for one to inflation to tighten the real rate. This achieves stability with certainty. Here, target-reaffirming communication holds a fraction $\bar{\tau}$ of the priced forecast at the target, so the same drift raises the forecast by only $1 - \bar{\tau}$ per unit: the standing message supplies the missing real-rate response, and the current-rate rule need only cover the residual. The institution as a whole is therefore not stabilizing less—the durable layer of communication performs part of the job the rate coefficient performs in the textbook model. What is weaker is the required rate response, not policy. But there is a caveat. Coefficient-for-coefficient within the same rule class, $i_t = \phi_\pi \pi_t$, for which the textbook threshold is exactly one. And the proposition holds $\bar{\tau}$ fixed: in the full model the anchor is an equilibrium object, and a sustained off-target drift could erode it through the learning block, migrating the boundary back toward one. What this means is that this relaxation is only as durable as the credibility that sustains it. This term has a parallel in the literature. Mathematically, the δ-scaling of the forecast is the same channel by which cognitive discounting relaxes the Taylor principle in behavioral New Keynesian models (Gabaix, 2020) and by which less-than-fully-forward-looking expectations relax it under learning (Bullard and Mitra, 2002); here the discounting is not a behavioral primitive but an equilibrium service supplied by the imposed targeting regime on the delegated institution.

5 Application: Policy Response to Pandemic Inflation

Our leading application studies the post-pandemic inflation surge in the U.S. The exercise is organized around the narrative laid out in Section 1: a credible central bank—here, the Federal Reserve—enters the pandemic with a strong record, deploys guidance at the lower bound, and then faces a large supply shock that it deems transient. It then *pursues* aggressive but *announces* low future inflation—based on its truthful but optimistic private assessment of the shock's persistence. We ask three questions. What does the announced accommodation buy relative to silent adherence to the default target? And how far do outcomes deteriorate as the public's reading of the institution hardens—from second-guessing the promise to reading the central bank's words as preferences rather than information?

Two facts support this narrative. First, the anchor held. Figure 1 plots quarterly CPI

inflation against the Federal Reserve Bank of Cleveland's model-based one-year expected-inflation measure. CPI inflation peaked at 8.58 percent in 2022Q2, while the expectation measure barely breached 3.75 percent, a gap sustained for well over a year—which from the perspective of time looks almost irrational. To quantify the pull of the target, it's instructive to write expected inflation as a weighted average of the long-run target ($\pi^* = 2$) and contemporaneous realized inflation,

$$\pi_t^e = w_t \pi^* + (1 - w_t)\pi_t, \tag{79}$$

so that the implied weight on the target is

$$\widehat{w}_t = 1 - \frac{\pi_t^e - \pi^*}{\pi_t - \pi^*}. \tag{80}$$

Using simple rolling regression analysis shows that, at the peak of the surge, this diagnostic weight remained at 0.734, similar to the whole period 2017–2024—which yields an implied target weight of $1 - \widehat{\tau} = 0.756$.[6] Second, anecdotal evidence and FOMC dot plots show that, in fact, in 2021 and even in 2022 the public was assured that deviations from the target is short lived (relative to what happened ex post). Finally, third, as documented below, the central bank's *marginal* communication lost its market footprint as the model predicts: scheduled FOMC days ceased to carry excess news for short-horizon inflation compensation even as they continued to move nominal yields.

Setup. We trace the episode year by year from 2020 to 2027 through the model of Section 3, solved along the path rather than over the stationary distribution.[7] The macroeconomic block uses a flat, production-calibrated Phillips curve—slope $\kappa = 0.15$ and activity weight $\vartheta = 0.40$—from the capacity-walls production network of Drozd and Tavares (2026), in which supply forces drive the post-COVID surge; Rubbo (2024a,b) likewise find supply forces central to the 2021 inflation. The inherited regime specifies a strict *default target* conditional on the disclosed supply outlook, $\bar{\theta}_t = \phi_{IT}\, \widehat{u}_{t+1|t}$ with loading $\phi_{IT} = 0.35$—the default target stays close to two percent even in the surge—and the central bank announces a center $a_t = \bar{\theta}_t + g_t$, where g_t is the forward guidance shift. The announcement is a *conditional plan on the stance*: honoring delivers the activity gap the announced center

[6] The regression $\Delta_4 \pi_t^e = \alpha + \tau \Delta_4 \pi_t + \varepsilon_t$ yields $\widehat{\tau} = 0.244$ (HAC standard error 0.031). The Cleveland Fed measure structurally conditions on realized inflation, so $\widehat{\tau}$ is partly an artifact of that model; the level fact in Figure 1 is unambiguous.

[7] Solving the model year by year preserves the message posterior, honor–waiver margin, conditional tether, and expectation mapping while making the episode's timing explicit. The message likelihoods remain frozen local approximations, as stated at the start of this section. Computational details are in Section 8.

implied on the central bank's own outlook, and all surprises—supply and expectational—pass through prices, so a miss caused by a shock the disclosed forecast did not contain is a forecast error, not a breach. What the public prices prospectively, and what a skeptical public fears, is the *letter* of the announced center evaluated at the state it believes will arrive.

The path is the episode's. In 2020 a lower-bound recession hits: the natural rate collapses to $r^* = -4.1$, the bound binds, and expected inflation is the only stabilization margin. The 2020 announcement is the one policy object we *impose* rather than optimize: +0.50, an announced center of 2.52 percent, the "moderate overshoot" of the Federal Reserve's 2020 framework. The bound then binds *through 2021*, as in the data: liftoff came only in March 2022, and what held the central bank at the bound in 2021 was the outstanding threshold guidance—honoring it means not lifting off, and since the bound is one-sided, waiving it would be an early liftoff onto the reset path. From 2021 onward every announcement is chosen optimally on the central bank's own information. Each year the central bank privately observes a signal implying that only a fraction 0.35 of the current pressure will carry into the following year, discloses the implied forecast—as in a published projection—and issues its announcement accordingly; the realized process instead decays at 0.70, so each year's adverse surprise arrives after that year's announcement.

Why must the *words* carry the accommodation? On a flat Phillips curve, delivered inflation in a cost-push year is essentially expected inflation plus the shock; demand restraint barely dents it, so matching the episode—high inflation, no recession—requires the policy stance to accommodate nearly the whole shock. That accommodation can be granted only once, by one of two instruments. If the default target itself were generous, the standing regime would already promise it: the optimizer then finds nothing for the announcement to buy and chooses $g \approx 0$, and communication is quantitatively irrelevant. With the strict default target the regime actually inherited, the accommodation must be *announced*: a positive announcement above the default target, a promise that must be priced, can be doubted, and can be spent. Guidance matters exactly when the inherited target regime is stricter than the policy the central bank needs to deliver—the announcement is valuable because it is a departure from the standing message, which is this paper's thesis recovered as a quantitative comparative static. Table 1 lays out who knows what within a year; Table 2 collects the calibrated objects.

We discipline the shock by the episode's real-side facts, not by chasing the inflation path year by year. The recession was in 2020; the surge years were years of recovery and a tight labor market. We therefore match 2020 on *both* margins—the supply disturbance $u_{2020} = 0.15$ and the natural-rate collapse $r^*_{2020} = -4.1$ are solved jointly so that the model

Table 1: Information and beliefs within year t

Object	The central bank	The public
Current pressure u_t	Observed.	Observed.
Next year's pressure u_{t+1}	Private signal implying the point forecast $\rho^{CB}u_t$; disclosed in full, as in a published projection.	Prior (before the disclosure): its own read $\rho^{\text{pub}}u_t$. Posterior (after): $\varpi\,\rho^{CB}u_t + (1-\varpi)\,\rho^{\text{pub}}u_t$, where ϖ is the forecast credence given to the disclosed forecast.
Policy mode	Knows its own mode.	Prior: override probability λ_t. Posterior after hearing the announcement g_t: q_t from the message densities $(\sigma_{\mathsf{G}}, \sigma_{\mathsf{R}})$; updated once more at year-end from words and outcomes.
The announced center $a_t = \bar{\theta}_t + g_t$	Forward guidance shift g_t chosen optimally on the central bank's own disclosed forecast (2020: imposed at +0.50, the FAIT overshoot); the center is fixed at issuance.	Prices honoring the *letter* of a_t at its own believed state u_{t+1}: prospective survival p_t from the honor–waiver margin at threshold dispersion σ_F. A conditional miss—the announced center plus the pass-through of a supply surprise—is not a breach.
Future realizations	Unknown; each year's adverse surprise arrives after that year's announcement.	Unknown.

Notes: In the baseline reading the public accepts the disclosed forecast ($\varpi = 1$, $\rho^{\text{pub}} = \rho^{CB}$) and holds the diffuse institutional prior ($\sigma_F = 4.5$, $\lambda_{2020} = 0.12$). The central bank knows neither future realizations nor—in the misreading counterfactuals—the public's actual reading: those counterfactuals lower ϖ, raise ρ^{pub}, sharpen σ_F, and (in Misreading II) raise λ_0 and flatten the message densities, holding the central bank's policy path fixed. In 2020 the disclosed forecast is of essentially no pandemic supply pressure ($\rho^{CB}u_{2020} \approx 0.05$): the 2021 outbreak is a surprise to everyone, and the 2020 announcement works through the mode posterior and the priced promise. In the bound years (2020–21), honoring the outstanding promise means remaining at the bound; waiving it means an early liftoff onto the reset path.

reproduces 2020 core inflation of 1.6 percent and a recession gap of −4.0 percent exactly. Because an annual model whose year-t decisions and shocks produce year-t inflation maps to the price change *during* calendar year t, we target core CPI December-over-December (Dec/Dec): the outbreak's peak $U_{2021} = 2.63$ is estimated so that the model's 2021 inflation equals the Dec/Dec 2021 print exactly, and 2022–25 are left untargeted (root mean squared error 0.82 percentage points). Measured Dec/Dec the surge arrives *in* 2021, where the model puts it; the annual-average convention is a moving-average filter that shifts the apparent peak about half a year later. The single visible residual, 2022, has a name: CPI shelter lags market rents by about a year, so late-2022 core still records much of 2021's inflation, with the remainder the supply-side propagation of Drozd and Tavares (2026)—both outside this model by design, which is also why core, not headline, is the calibration target. The decay rate is set a priori at $\rho_u = 0.70$—a supply disturbance with a roughly two-year half-life, consistent with the normalization of supply-chain pressure by 2023—

rather than estimated: freeing it drives the fitted shock toward permanence in order to buy the 2023–24 inflation tail, at the cost of serial waivers and restraint that never fades. The tail is instead left untargeted, and the model undershoots it (a 2023–24 average of 3.27 against 3.55 percent): with expectations anchored, a cost-push waning at ρ_u cannot hold inflation near five percent alongside the roughly zero output gap observed in those years. That persistent tail belongs to the supply-side propagation studied in Drozd and Tavares (2026), which we deliberately keep out of the estimated shock; core (rather than headline) CPI is the target throughout for the same reason. The discrepancy between the disclosed persistence 0.35 and the realized 0.70 is the model's measure of how optimistic the "transitory" call was. The Cleveland expectation series is nowhere targeted: how the model tracks it is the exercise's out-of-sample test.

Table 2: Post-pandemic application: objects held fixed

Object	Role	Value
κ, ϑ, β_{NK}, α	NK block (flat, production-calibrated)	0.15, 0.40, 0.96, 1
ϕ_{IT}	Default-target loading on disclosed forecast	0.35
U_{2021}	Outbreak peak (matches 2021 core CPI, Dec/Dec)	2.63
ρ_u	Shock decay (set: ∼2-year half-life)	0.70
u_{2020}, r^*_{2020}	Match 2020 core CPI and gap of −4.0	0.15, −4.1
r^*_{2021}	Natural rate recovered; bound still binds	0.0
g_{2020}	Run-hot announcement (imposed; FAIT overshoot)	+0.50
ρ^{CB}	Signal-implied persistence (disclosed)	0.35
$[\underline{g}, \overline{g}]$	Authorization range for announcements	[−1.5, 2.0]
$\bar{F}_{\text{R}}$, $\bar{F}_{\text{G}}$	Waiver-cost locations	0.5, 4.5
σ_F	Waiver-cost dispersion, robust / sharp	4.5 / 0.25
$\bar{\lambda}$, ρ_λ, λ_0	Override prior: mean, persistence, 2020	0.05, 0.87, 0.12
σ_{G}, σ_{R}	Message dispersions	1.2, 3.0
$\bar{\varrho}$	Reset tether, $\varrho = \bar{\varrho}(1 - \lambda_t)$	0.15
ξ_λ	Credibility value	3.0
σ_π	Inflation-control noise	1.10

Notes: Annual periods. The outbreak peaks in 2021 and decays geometrically; U_{2021} is estimated so the model's 2021 core CPI inflation (Dec/Dec) matches the data exactly, with the policy path re-optimized at every trial, and 2022–25 are untargeted. (u_{2020}, r^*_{2020}) are solved jointly so 2020 matches both core inflation and the recession gap. The bound binds in 2020–21 (liftoff is 2022); from 2021 onward the central bank re-optimizes its announcement each year on its own information within the authorization range; counterfactuals below hold the resulting policy path fixed and vary only the public's reading; these rows price an unchanged path and are not equilibrium counterfactuals in which the central bank re-optimizes after observing the darker interpretation. The reset parameters $\bar{\varrho} = 0.15$ and $\xi_\lambda = 3.0$ impose the reputational-value restriction in (64). Model variables are deviations from target; reported levels add the two-percent target.

The waiver schedule uses the local optimized gap, not the difference in current period

losses. A finite-difference Hessian of the continuation-adjusted branch objective gives the reduced reset curvature $\widehat{\mathcal{A}} = 11.69$ and the optimized gap coefficient $\widehat{\Psi} = 11.55$, compared with the period-loss curvature $A_\theta = 18.78$. Thus $\widehat{\Psi}/A_\theta = 0.62$: reoptimization lowers the relevant temptation curvature materially. At the 2021 state, the unconstrained implemented-effect map peaks at an announcement of 1.49, inside the authorization upper bound 2.00. The chosen 2021 announcement is only 0.40, so neither the authorization box nor the communication frontier binds along the baseline path.

5.0.1 Data path versus the model

Table 3 reports the baseline transition and Figure 6 plots it against the data. The path reads as the episode's timeline, with the real side where the data put it: the recession in 2020, a hot economy through the surge, restraint from liftoff that fades as the shock wanes. In 2020 the economy is in the matched recession (a gap of −4.0 percent) and the central bank issues its threshold guidance with the moderate overshoot attached: an announced center of 2.52 percent, and, above all, the commitment not to lift off. In 2021 the outbreak arrives at full force ($u_{2021} = 2.63$) with that guidance outstanding, and the central bank honors it by staying at the bound: policy never tightens against the shock, the economy runs hot (a gap of +0.78 percent), and inflation prints 5.50 percent—high inflation in a recovered economy, the episode's defining feature. The guidance issued in the recession thus comes due in an especially unfavorable state—the mechanism Ihrig and Waller (2024) describe from inside the FOMC when they call the 2020 threshold guidance "too constraining."

Alongside, the central bank discloses its optimistic forecast and issues its largest announcement: a forward guidance shift of +0.40 *above* the strict default target, an announced center of 2.72 percent against a default target of 2.32. This is the paper's mechanism at work: the accommodation the episode requires is not in the inherited target regime—it is spoken. The public prices the promise at survival 0.72, and one-year expected inflation settles at 2.78 percent—the model's counterpart of the Cleveland measure's hump—while realized inflation runs near six: the episode's signature gap between realized and expected inflation, generated by a priced, honorable promise rather than by a breach. In 2022 the central bank lifts off on schedule. The shock has persisted far beyond the disclosed forecast, but the announcement was a conditional plan, and delivered inflation of 4.12 percent is the announced stance delivered into the realized state—a forecast error, not a breach. Restraint arrives with liftoff at the scale of a hiking cycle (−1.7 percent, fading below one by 2025), the announced centers glide down (2.53, 2.36, 2.26, . . .), every promise on the path is honored, and the override prior declines from $\lambda_0 = 0.12$ to 0.01: credibility is rebuilt through delivery, and the "transitory" miss is absorbed as the forecast error it was.

What does the announced accommodation buy? Not lower inflation—on a flat Phillips curve the surge is the shock's doing, and silent adherence to the default target delivers nearly the same peak (5.24 against 5.50 percent). What it buys is the *output the strict default target would have burned*. Left to speak for itself, the strict target's stance grinds against the surge: silence delivers restraint of −4.4 and −3.1 percent in 2022–23—sustained tightening the episode never saw—and the public prices the target's letter at only 0.62. The announced glide path relieves that restraint (gaps of −1.7 and −1.1 percent, the hiking cycle as observed), is priced at 0.72, and cuts the episode's cumulative loss from 69 to 43 in equal-weight loss units—roughly forty percent—at a cost of two-tenths of a point of peak inflation and expectations (2.78 against 2.55). This is the makeup framework's bargain, priced: the accommodation was bought with words, openly, and paid for in a slightly higher and fully anticipated inflation path. And the strategy carries an option on vindication: had the shock faded as the signal implied, the 2021 peak would have been identical, the glide path honored into a rapid return to target, and the episode's cumulative loss cut by roughly a quarter (33 against 43 in equal-weight loss units; a cumulative output cost of 6.7 against 8.9 percentage points)—the model's version of the immaculate disinflation. The fork between that outcome and the realized one is nature's draw, not a difference in policy.

Table 3: Baseline transition: the episode as the model reads it

	2020	2021	2022	2023	2024	2025	2026	2027
Cost-push u_t	0.1	2.6	1.8	1.3	0.9	0.6	0.4	0.3
New wedge g_t	0.50	0.40	0.30	0.20	0.15	0.10	0.05	0.05
Announced center	2.52	2.72	2.53	2.36	2.26	2.18	2.10	2.09
Expected inflation	2.06	2.78	2.56	2.38	2.28	2.19	2.12	2.09
Inflation	1.60	5.50	4.12	3.50	3.05	2.74	2.50	2.34
Output gap	-4.00	0.78	-1.73	-1.08	-0.82	-0.51	-0.39	-0.41
Plan survival p_t	0.72	0.72	0.72	0.73	0.73	0.73	0.73	0.73
Promise honored	–	✓	✓	✓	✓	✓	✓	✓
Override prior λ_t	0.12	0.05	0.04	0.02	0.02	0.01	0.01	0.01

Notes: Annual periods; inflation, expectations, and announced centers are levels in percent. The announced center applies to the following year; "Plan survival" is the public's prospective probability that the newly issued announced center will be honored; "Promise honored" reports the settlement of the *outstanding* promise (none is outstanding in 2020). A conditional miss—delivered inflation equal to the announced center plus the pass-through of the supply surprise—is honoring, not a breach; every promise on the baseline path is honored. λ_t is the start-of-year override prior; with no breach on the path it declines from $\lambda_0 = 0.12$ toward its long-run mean. The bound binds in 2020–21; liftoff is 2022.

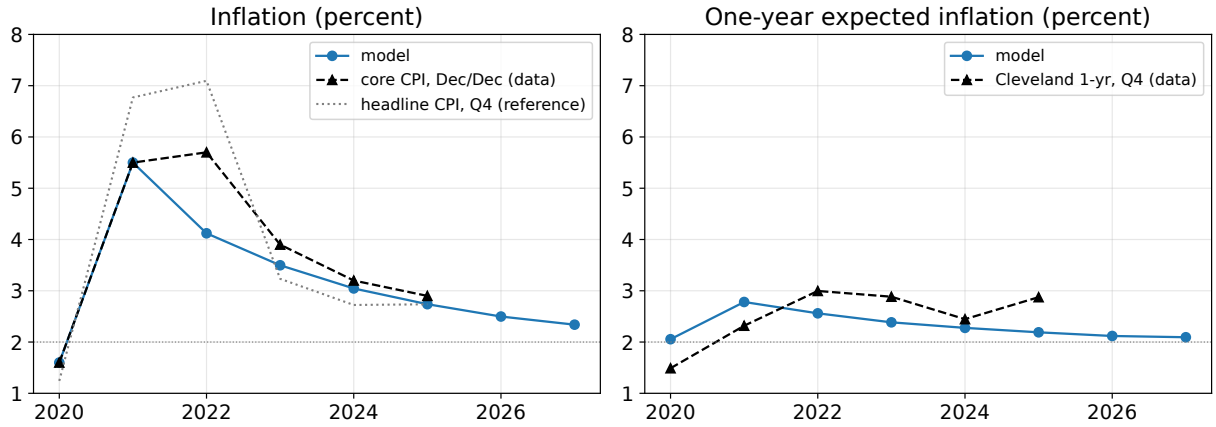

Figure 6: Baseline against the data: inflation and one-year expectations

Notes: Model levels add the two-percent target; both panels share a common scale. Left panel: model inflation against core CPI measured December-over-December (data), with headline CPI at Q4 as a reference. The calibration matches 2020 and the 2021 Dec/Dec print exactly; 2022–25 are untargeted, and the 2022 residual reflects the CPI shelter lag (late-2022 core still recording 2021's inflation) plus the supply-side propagation studied in Drozd and Tavares (2026). Right panel: model one-year expectations against Q4 readings of the Cleveland Fed one-year (headline) measure; this series is nowhere targeted (root mean squared error 0.48 percentage points over 2021–25), and the model's 2020 value reflects the announced run-hot center, a channel absent from the survey-based measure during the pandemic collapse.

5.0.2 Policy risks

Consider now counterfactuals that hold the entire baseline policy path fixed—every announced center, every disclosed forecast—and vary only how the public reads the institution from 2021 onward. Table 4 summarizes; Figure 7 plots the paths. Two cases as illuminating of the policy risks, and how this same scenario helps explain the difference of essentially similar policy approach in the 1970s that turned out very differently.

In the first case the public second-guesses the promise: it prices the letter of each announced center at a persistence of 0.85 rather than the disclosed 0.35, and the institutional threshold is sharp ($\sigma_F = 0.25$)—but the central bank does not know it and issues the baseline announcements. Priced at the state the public fears, the letter of the glide path is brutal to deliver on a flat curve, so prospective survival collapses to 0.031 and, with the anchor weakened, the bound years run hotter: 2021 inflation reaches 6.46 percent with expectations at 3.65. In this equilibrium the misses *are* read as breaches: the override posterior climbs to 0.81 before delivery slowly rebuilds it—credibility spent even though the central bank's conduct never changed. Yet the collapse is bounded, because the announcement still carries information about the shock (the public retains a weight of 0.5 on the disclosed forecast): expected inflation peaks one point above the baseline, not seven, and the cumulative loss rises to 56 against the baseline's 43. A credible central

bank that fails still communicates. And retreating to silence makes matters *worse*, not better: the retreat row—the same sharp reading with no announcements beyond the default target—reaches the same peaks but a cumulative loss of 87, half again the misread central bank's, because silence surrenders the announced accommodation while the sharp reading persists regardless. Under a sharp reading it is the institution being priced, not the message; the words, even when disbelieved as promises, still deliver the stance and the information.

The second case is the Great-Inflation reading, and it is qualitatively different: the public believes the announcements reflect the policymaker's *preferences* rather than its information—the words carry no signal about the shock and next to no signal about the mode—and it believes those preferences are inflationary and variable (a dovish bias of 0.1 priced into the discretionary fallback, a prior of 0.45 on the manipulative mode, and no reputational tether left). Communication then neither commits nor informs. Expected inflation explodes to 9.71 percent and the bound cannot hold: the central bank abandons the guidance and lifts off a year early into discretion at the public's own assessment of the shock augmented by the priced-in bias, delivering inflation of 11.65 percent with expectations at 9.71—the Great-Inflation configuration, from an identical policy path, at more than four times the baseline's cumulative loss (217 against 43). The distance between the two cases, roughly six percentage points of peak expected inflation, is the measured value of the second anchor: being believed to be *wrong-but-informative* rather than *manipulative*. It is also the model's answer to why the post-COVID episode never approached the Great Inflation, and why, ex ante, a central bank's most valuable reputation is that its preferences are steady.

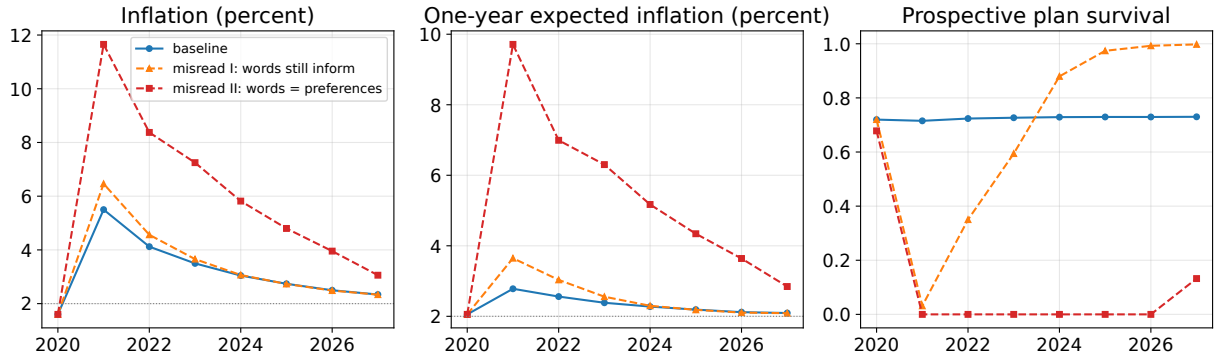

Figure 7: Baseline and the two misreadings

Notes: All three paths share the same policy (announcements and disclosed forecasts); they differ only in the public's reading of the institution from 2021 onward. Misreading I: sharp waiver threshold, skeptical persistence read, announcement retains information weight 0.5. Misreading II: announcements read as variable policy preferences (no information content, dovish bias, no tether).

Table 4: One policy path under alternative readings of the institution

	Peak π	Peak $\mathbb{E}\pi$	Min. survival	Waived	Cum. loss	Peak λ
Silent default target (no guidance)	5.24	2.55	0.62	–	69	0.12
Vindicated look-through	5.50	2.78	0.72	–	33	0.12
Baseline (shock persists)	5.50	2.78	0.72	–	43	0.12
Retreat to default target under sharp reading	6.49	3.67	0.00	–	87	0.82
Misreading I (words still inform)	6.46	3.65	0.03	–	56	0.81
Misreading II (words = preferences)	11.65	9.71	0.00	2021	217	0.88

Notes: The two misreading rows hold the baseline policy path fixed from 2021 onward and vary the public's reading. "Retreat to default target under sharp reading" keeps the 2020 guidance and the disclosed forecasts but issues no announcements beyond the default target from 2021 onward, under the same reading as Misreading I. "Waived" lists the years in which the outstanding promise was abandoned; in a bound year a waiver is an early liftoff, which is how Misreading II leaves the bound in 2021. "Cum. loss" is the cumulative quadratic loss on absolute deviations—inflation from the two-percent target and the output gap, at equal weights—summed over 2020–27. The baseline's loss includes the restraint that honoring its glide path requires after liftoff; regimes that breach fall to the accommodative reset and escape that cost, which is why the "Waived" column must be read alongside the loss.

5.0.3 Test of model's core mechanism

We return to the second empirical fact that tests the model. The standing message can continue to hold down the *level* of expectations even when scheduled announcements carry little incremental news. In the application, marginal traction falls modestly: it is $M_0 = 0.729$ near the rule and $M_E = 0.711$ at the 2021 look-through announcement, so $M_E/M_0 = 0.976$. But once the glide path is announced, the optimal policy center moves smoothly (2.72, 2.53, 2.36, . . .), sharply reducing the variance of unexpected announcement innovations.

The market counterpart to this is an event-study statistic in the tradition of Gürkaynak et al. (2005). For a series x and scheduled FOMC decision days $\mathcal{D}$, define

$$R^x_{\mathcal{D}} = \left[\frac{N_{\mathcal{D}}^{-1} \sum_{t \in \mathcal{D}} (\Delta x_t)^2}{N_{\neg\mathcal{D}}^{-1} \sum_{t \notin \mathcal{D}} (\Delta x_t)^2} \right]^{1/2}, \tag{81}$$

the excess news injected by meeting days. Figure 8 reports this ratio for two-year inflation compensation and the two-year nominal yield. Before the pandemic, FOMC days were major news events for both series (1.66 for inflation compensation, 1.65 for yields). During the 2021–23 surge the series decoupled: the yield ratio stayed elevated at 1.32—meetings moved rates—while the inflation-compensation ratio collapsed to 0.78. Meeting days ceased to carry excess *inflation* news precisely when inflation was discussed the most. As the surge receded, the ratio recovered (1.01): a temporary quieting of the message, not a paradigm shift. To map the statistic to the model, let V_E denote the variance of unexpected

guidance in era E. Up to the ordinary-day normalization,

$$R_E^2 - 1 = \varkappa_R \left(\frac{M_E}{M_0}\right)^2 \left(\frac{V_E}{V_0}\right), \tag{82}$$

where $\varkappa_R = R_0^2 - 1$ is pinned down by the anchored era. Holding announcement variance fixed, the application's traction decline would predict a ratio of 1.63, far above the observed 0.78. In the quiet-message limit, $V_E/V_0 = 0.000$ and the model ratio is one. The empirical estimate of 0.78 lies below one because ordinary-day volatility and finite-sample variation also enter the denominator. Thus the exhibit supports an unusually predictable standing message; it does not establish that marginal traction literally vanished or that the central bank sat at the communication frontier.

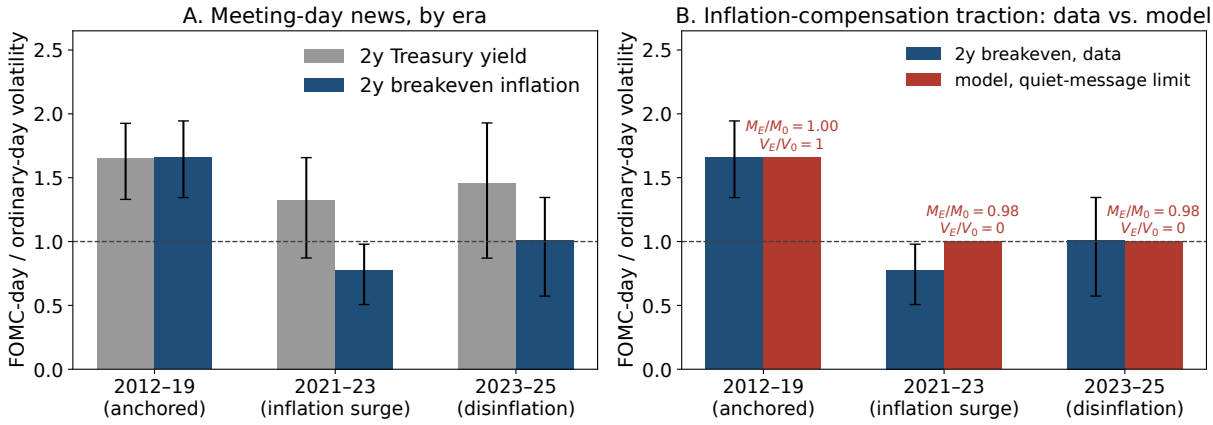

Figure 8: FOMC-day news before, during, and after the inflation surge: data and model

Notes: Ratio (81) on scheduled FOMC days relative to ordinary days; error bars are bootstrap 90-percent intervals. Panel A reports the statistic by era for two-year inflation compensation and the two-year nominal Treasury yield. Panel B overlays, for inflation compensation, the model's quiet-message limit. The annotation reports both the solved marginal-traction ratio and the maintained announcement-innovation variance ratio; the surge collapse is not represented as zero marginal traction. The pre period is 2012–19, the surge 2021–23, the post 2023–25.

6 Conclusion

Inflation targeting is more than a numerical objective. It is a durable institutional regime and a standing form of communication that both supports and constrains the central bank's discretionary use of operational forward guidance. Taking that regime as given, this paper has developed a tractable theory of what operational communication around it can accomplish and what limits the resulting departures.

The analysis establishes several properties. The standing target separates average from marginal traction: it may continue to anchor the level of expectations even when additional statements have little power to move them. Words are a policy margin. Depending on their distance from the inherited rule, they can also build institutional credibility in levels while spending it at the margin.

Most importantly, however, traction is state dependent. The public interprets an announcement jointly with the observable economic state, the central bank's institutional history, and the incentive to waive the promise when it comes due. As the communication frontier is approached, the institutional cost of obtaining an additional unit of expectation management rises sharply. Ignoring this property can lead to policy myopia in which forward guidance is overused, draining credibility that may be useful to respond to future shocks.

Our most important result identifies the central risk of "looking through" large shocks. Under rational expectations, the central bank understands how the public will interpret its announcement and adjusts policy to the traction it anticipates. In practice, however, this mapping is difficult to assess in real time. The danger does not arise simply because communication is imperfectly credible. It depends on what the public believes the communication reveals, even credible communication. If an announcement is read primarily as a data-dependent signal about the persistence of the shock, the consequences of a moderate forecasting error may remain contained. If instead the same announcement is interpreted as evidence of unstable or unusually accommodative policy preferences, credibility and traction can deteriorate much more severely. The critical distinction is thus whether the public reads the central bank's words as information about the economy or as information about a fickle policymaker.

Our framework offers a unified interpretation of several surprising features of the post-pandemic inflation episode. The persistent gap between high realized inflation and comparatively anchored one-year expectations, shown in Figure 1, indicates that the standing target retained substantial influence even after the transitory assessment proved incorrect. The same expectations channel is consistent with the relatively low sacrifice ratio of the subsequent partial disinflation, as documented by Forbes et al. (2025). At the same time, scheduled FOMC announcement days ceased during the inflation episode to carry the excess short-horizon inflation-compensation news that they had carried before COVID-19, even though they continued to move nominal yields. This decoupling is consistent with a diminished marginal market footprint of new inflation communication. In the model it stems from a combination of limited incremental traction and increasingly predictable policy communication.

The broader institutional lesson is that there is no true silence under inflation targeting: choosing no discretionary guidance wedge reaffirms the inherited rule. The relevant policy problem is therefore not whether the central bank should communicate, but when it should depart from its standing message, how far it should go, what it should worry about, and how tightly it should hold on to that departure after the fact. Our framework tractably captures these margins.

References

Amador, Manuel and Christopher Phelan, "Central Bank Reputation with Noise," *Journal of Monetary Economics*, 2026, *158*, 103906.

Bassetto, Marco, "Forward Guidance: Communication, Commitment, or Both?," *Journal of Monetary Economics*, 2019, *108*, 69–86.

Bullard, James and Kaushik Mitra, "Learning about Monetary Policy Rules," *Journal of Monetary Economics*, 2002, *49* (6), 1105–1129.

Choi, Jason and Andrew Foerster, "Optimal Monetary Policy Regime Switches," *Review of Economic Dynamics*, 2021, *42*, 333–346.

Christoffel, Kai and Mátyás Farkas, "Managing the Risks of Inflation Expectation De-Anchoring," Working Paper Series 3082, European Central Bank 2025.

Cukierman, Alex and Allan H. Meltzer, "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 1986, *54* (5), 1099–1128.

Dovis, Alessandro and Rishabh Kirpalani, "Rules without Commitment: Reputation and Incentives," *The Review of Economic Studies*, 2021, *88* (6), 2833–2856.

Drozd, Lukasz A., "Information Architecture of Central Bank Communication," 2026. Working paper.

_ **and Marina M. Tavares**, "Capacity Walls in Production Networks: Inflation after COVID," Working Paper, Federal Reserve Bank of Philadelphia July 2026. Forthcoming.

English, Bill, Kristin J. Forbes, and Angel Ubide, eds, *Monetary Policy Responses to the Post-Pandemic Inflation*, London: CEPR Press, 2024.

Erceg, Christopher J. and Andrew T. Levin, "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 2003, *50* (4), 915–944.

Forbes, Kristin, Jongrim Ha, and M. Ayhan Kose, "Tradeoffs over Rate Cycles: Activity, Inflation and the Price Level," NBER Working Paper 33825, National Bureau of Economic Research 2025.

Gabaix, Xavier, "A Behavioral New Keynesian Model," *American Economic Review*, 2020, *110* (8), 2271–2327.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, 2 ed., Princeton University Press, 2015.

Gobbi, Lucio, Ronny Mazzocchi, and Roberto Tamborini, "Monetary Policy, De-Anchoring of Inflation Expectations, and the "New Normal"," *Journal of Macroeconomics*, 2019, *61*, 103070.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 2005, *1* (1), 55–93.

Ihrig, Jane and Christopher J. Waller, "The Federal Reserve's Responses to the Post-COVID Period of High Inflation," in Bill English, Kristin J. Forbes, and Angel Ubide, eds., *Monetary Policy Responses to the Post-Pandemic Inflation*, London: CEPR Press, 2024, chapter 7, pp. 131–142.

King, Robert G. and Yang K. Lu, "Evolving Reputation for Commitment: Understanding Inflation and Inflation Expectations," July 2026. Working paper, Boston University and Hong Kong University of Science and Technology.

_ , _ , and Ernesto S. Pastén, "Managing Expectations," *Journal of Money, Credit and Banking*, 2008, *40* (8), 1625–1666.

Kostadinov, Rumen and Francisco Roldán, "Reputation and the Credibility of Inflation Plans," *Journal of Economic Theory*, 2025, *229*, 106062.

Lohmann, Susanne, "Optimal Commitment in Monetary Policy: Credibility versus Flexibility," *American Economic Review*, March 1992, *82* (1), 273–286.

Lu, Yang K., Robert G. King, and Ernesto S. Pastén, "Optimal Reputation Building in the New Keynesian Model," *Journal of Monetary Economics*, 2016, *84*, 233–249.

Nakamura, Emi, Venance Riblier, and Jón Steinsson, "Beyond the Taylor Rule," NBER Working Paper 34200, National Bureau of Economic Research September 2025.

Rubbo, Elisa, "Networks, Phillips Curves, and Monetary Policy," *Econometrica*, 2024. Forthcoming.

_ , "What Drives Inflation? Lessons from Disaggregated Price Data," Working Paper 32194, National Bureau of Economic Research 2024.

Schaumburg, Ernst and Andrea Tambalotti, "An Investigation of the Gains from Commitment in Monetary Policy," *Journal of Monetary Economics*, 2007, *54*, 302–324.

Svensson, Lars E. O., "Monetary Policy Strategies for the Federal Reserve," *International Journal of Central Banking*, 2020, *16* (1), 133–193.

Walsh, Carl E., "Optimal Contracts for Central Bankers," *American Economic Review*, March 1995, *85* (1), 150–167.

Online Appendix

Not for Publication

7 Proofs and Derivations

7.1 The optimized Bellman gap

This subsection derives Proposition 1 while retaining the distinct control choices on the honor and waiver branches. Let

$$\mathcal{J}_j(\theta, g; \Omega) = -U\big(\theta, x(\theta, g; \Omega)\big) + \beta \mathbb{E}^j W_{j'}\Big(\widetilde{\Omega}'(\theta, g, \varepsilon^\pi; \Omega)\Big) \tag{83}$$

be the continuation-adjusted loss introduced in the text. The expectation is written using fixed underlying innovations. Thus differentiation of the continuation term includes both derivatives of the value function and derivatives of the transition map. For any scalar arguments $y, z \in \{\theta, g, g_{-1}\}$,

$$\begin{aligned}\partial_{yz} \mathbb{E}^j W_{j'}(\widetilde{\Omega}') = \mathbb{E}^j \Big[&(D_y \widetilde{\Omega}')^\top W_{j',\widetilde{\Omega}\widetilde{\Omega}}(\widetilde{\Omega}') D_z \widetilde{\Omega}' \\ &+ W_{j',\widetilde{\Omega}}(\widetilde{\Omega}') D_{yz} \widetilde{\Omega}' \Big].\end{aligned} \tag{84}$$

In particular, the derivatives below include the effects of the current controls on λ', on the promise g inherited next period, and on all other components of $\widetilde{\Omega}'$.

Let $h \equiv g_{-1}$ and hold (S, λ) at the reference value $(\bar{S}, \bar{\lambda})$. Write the waiver deviation from the honored center as

$$\delta \equiv \theta - (\bar{\theta} + h), \qquad \bar{\theta} \equiv \bar{\Theta}(\bar{s}), \tag{85}$$

where $\bar{s}$ is the public component of $\bar{S}$, and define

$$Q_j(\delta, g, h) \equiv \mathcal{J}_j(\bar{\theta} + h + \delta, g; \bar{S}, \bar{\lambda}, h). \tag{86}$$

Honor imposes $\delta = 0$ and minimizes over g; waiver minimizes over both (δ, g). At the reference point $(0, 0, 0)$, the derivatives with respect to (δ, g) vanish. The second-order expansion is

$$\begin{aligned}Q_j(\delta, g, h) = Q_{j0} + q_{j,h} h \\ + \frac{1}{2}\left[a_j \delta^2 + 2 b_j \delta g + c_j g^2 + 2 r_j \delta h + 2 s_j g h + t_j h^2\right] + o(\delta^2 + g^2 + h^2),\end{aligned} \tag{87}$$

where all derivatives are evaluated at the reference point. Assume the control Hessian

$$H_j \equiv \begin{pmatrix} a_j & b_j \\ b_j & c_j \end{pmatrix} \tag{88}$$

is positive definite. Hence $c_j > 0$ and

$$\mathcal{A}_j \equiv a_j - \frac{b_j^2}{c_j} > 0. \tag{89}$$

Also define

$$\mathcal{B}_j \equiv r_j - \frac{b_j s_j}{c_j}. \tag{90}$$

Proof of Proposition 1. Under honor, $\delta = 0$. The first-order condition for the new announcement in the quadratic approximation is

$$c_j g + s_j h = 0,$$

so the exact policy has the local expansion

$$g_j^H(h) = -\frac{s_j}{c_j} h + o(h). \tag{91}$$

Substituting into (87) gives

$$W_j^H(h) = Q_{j0} + q_{j,h} h + \frac{1}{2}\left(t_j - \frac{s_j^2}{c_j}\right) h^2 + o(h^2). \tag{92}$$

Under waiver, the first-order conditions are

$$a_j \delta + b_j g + r_j h = 0, \tag{93}$$

$$b_j \delta + c_j g + s_j h = 0. \tag{94}$$

Solving (94) for g and substituting into (93) yields

$$\left(a_j - \frac{b_j^2}{c_j}\right)\delta + \left(r_j - \frac{b_j s_j}{c_j}\right) h = 0. \tag{95}$$

Therefore

$$\delta_j^W(h) = -\frac{\mathcal{B}_j}{\mathcal{A}_j}h + o(h), \tag{96}$$

$$g_j^W(h) = -\frac{s_j}{c_j}h + \frac{b_j\mathcal{B}_j}{c_j\mathcal{A}_j}h + o(h). \tag{97}$$

Equation (97) makes explicit that the two branches generally choose different new announcements:

$$g_j^W(h) - g_j^H(h) = \frac{b_j\mathcal{B}_j}{c_j\mathcal{A}_j}h + o(h). \tag{98}$$

For the optimized waiver value, let $z = (\delta, g)^\top$ and $k_j = (r_j, s_j)^\top$. The control-dependent quadratic is

$$\frac{1}{2}\left(z^\top H_j z + 2hk_j^\top z + t_j h^2\right).$$

Its minimizer is $z = -H_j^{-1}k_j h$, and its minimized coefficient on h^2 is $t_j - k_j^\top H_j^{-1}k_j$. Direct calculation gives

$$k_j^\top H_j^{-1}k_j = \frac{s_j^2}{c_j} + \frac{\mathcal{B}_j^2}{\mathcal{A}_j}. \tag{99}$$

Hence

$$W_j^W(h) = Q_{j0} + q_{j,h}h + \frac{1}{2}\left[t_j - \frac{s_j^2}{c_j} - \frac{\mathcal{B}_j^2}{\mathcal{A}_j}\right]h^2 + o(h^2). \tag{100}$$

Subtracting (100) from (92) gives

$$\Delta\mathcal{V}_j(h) = W_j^H(h) - W_j^W(h) = \frac{1}{2}\frac{\mathcal{B}_j^2}{\mathcal{A}_j}h^2 + o(h^2). \tag{101}$$

The waiver center is

$$\theta_j^W(h) = \bar{\theta} + h + \delta_j^W(h) = \bar{\theta} + \left(1 - \frac{\mathcal{B}_j}{\mathcal{A}_j}\right)h + o(h).$$

Thus

$$\omega_j = 1 - \frac{\mathcal{B}_j}{\mathcal{A}_j}, \qquad \Psi_j = \frac{\mathcal{B}_j^2}{\mathcal{A}_j} = \mathcal{A}_j(1-\omega_j)^2, \tag{102}$$

which proves the proposition. Since $\mathcal{A}_j > 0$, $\Psi_j \geq 0$, with strict inequality if and only if $\mathcal{B}_j \neq 0$. □

The restriction $\omega_j \in [0, 1]$ does not follow from convexity alone; it requires $0 \leq \mathcal{B}_j \leq \mathcal{A}_j$.

The static current-loss formula is nested as a special case. If there is no continuation term and the new announcement is irrelevant at delivery, then $b_j = s_j = 0$ and $a_j = r_j = \alpha + \vartheta/\kappa^2$. It follows that $\mathcal{A}_j = \mathcal{B}_j = \Psi_j = \alpha + \vartheta/\kappa^2$.

7.2 The pre-draw maximum and expected implementation curvature

Fix Ω and integrate only over the continuous fixed cost. Define $\widehat{W}_j(\Omega)$ as loss after the private state is known but before F_j is drawn. In loss notation,

$$\begin{aligned}\widehat{W}_j(\Omega) &= \mathbb{E}_F \min\{W_j^H(\Omega), W_j^W(\Omega) + F_j\} \\ &= W_j^W(\Omega) + \mathbb{E}_F \min\{W_j^H(\Omega) - W_j^W(\Omega), F_j\} \\ &= W_j^W(\Omega) + \mathbb{E}_F \min\{\Delta\mathcal{V}_j(\Omega), F_j\}. \end{aligned} \tag{103}$$

Taking the conditional expectation over the private state recovers $W_j(\widetilde{\Omega})$ in (34). This proves the exact identity used in the text.

Let $C_j(z) = \mathbb{E}\min\{z, F_j\}$. If F_j has a continuous distribution that is locally independent of h, differentiation under the integral gives

$$C_j'(z) = \mathbb{P}(F_j \geq z). \tag{104}$$

Therefore, with $z = \Delta\mathcal{V}_j(h)$,

$$\mathcal{D}_j'(h) = \chi_j(h)\Delta\mathcal{V}_j'(h), \qquad \chi_j(h) = \mathbb{P}(F_j \geq \Delta\mathcal{V}_j(h)). \tag{105}$$

At the reference state, $\Delta\mathcal{V}_j'(0) = 0$ and $\Delta\mathcal{V}_j''(0) = \Psi_j$, so

$$\mathcal{D}_j''(0) = \chi_{j0}\Psi_j = \Xi_j, \tag{106}$$

which establishes (44).

7.3 Local issuance and Gaussian likelihoods

The first-order condition associated with (46) is

$$-\zeta_j\widetilde{v} + (\gamma_j + \beta\Xi_j)g + o(|g| + |\widetilde{v}|) = 0. \tag{107}$$

The implicit-function theorem therefore gives

$$g_j(\widetilde{v}) = \frac{\zeta_j}{\gamma_j + \beta\Xi_j}\widetilde{v} + o(|\widetilde{v}|), \tag{108}$$

which is (47). Under the common-baseline restriction $\zeta_{\mathsf{G}} = \zeta_{\mathsf{R}} > 0$ and $\gamma_{\mathsf{G}} = \gamma_{\mathsf{R}} > 0$, the slope is strictly decreasing in Ξ_j, proving (48).

Proof of Lemma 1. Conditional on $(j, \widetilde{\Omega}_t)$, (49) is the sum of a constant, the Gaussian variable $K_j\widetilde{v}_t$, and independent Gaussian noise $\varepsilon^g_{j,t}$. Its mean is $\mu_j(\widetilde{\Omega}_t)$ and its variance is $K_j^2\sigma_v^2 + \sigma^2_{g,j}$, proving (50). With common noise variance and positive slopes, $K_{\mathsf{G}} < K_{\mathsf{R}}$ implies $\Sigma_{\mathsf{G}} < \Sigma_{\mathsf{R}}$. □

7.4 The Bellman–Gaussian communication frontier

Proof of Proposition 2. For centered Gaussian densities,

$$\frac{\mathcal{L}^g_{\mathsf{R}}(g)}{\mathcal{L}^g_{\mathsf{G}}(g)} = \frac{\Sigma_{\mathsf{G}}}{\Sigma_{\mathsf{R}}}\exp(\eta_g g^2). \tag{109}$$

Bayes' rule gives posterior odds $q(g)/[1-q(g)] = C_\lambda e^{\eta_g g^2}$, and hence (51). Under the quadratic Bellman approximation and a logistic fixed cost,

$$\begin{aligned}\chi(g) &= 1 - G_F\left(\frac{\Psi_{\mathsf{G}} g^2/2 - \bar{F}_{\mathsf{G}}}{\sigma_F}\right)\\ &= \frac{1}{1 + D_F e^{\kappa_F g^2}},\end{aligned} \tag{110}$$

which is (52). The polar commitment restriction implies $p(g) = [1-q(g)]\chi(g)$, so multiplication by the honored announcement g yields (54).

Logarithmic differentiation gives

$$\partial_g \log[1-q(g)] = -2\eta_g g q(g), \tag{111}$$

$$\partial_g \log \chi(g) = -2\kappa_F g[1-\chi(g)]. \tag{112}$$

Therefore

$$\frac{p_g(g)}{p(g)} = -2\eta_g g q(g) - 2\kappa_F g[1-\chi(g)], \tag{113}$$

and $m_g = p + g p_g$ gives (55).

Define

$$H(g) \equiv 2g^2\{\eta_g q(g) + \kappa_F[1 - \chi(g)]\}. \tag{114}$$

Then $m_g = p(1 - H)$. For $g > 0$,

$$\begin{aligned} H'(g) = {} & 4g\{\eta_g q + \kappa_F(1 - \chi)\} \\ & + 2g^2\{\eta_g q_g - \kappa_F \chi_g\} > 0, \end{aligned} \tag{115}$$

because $q_g > 0$ and $\chi_g < 0$. Moreover $H(0) = 0$ and $H(g) \to \infty$ as $g \to \infty$. Thus there is a unique $\bar{g} > 0$ satisfying $H(\bar{g}) = 1$. The sign of m_g is positive before that root and negative after it. Finally,

$$0 < m(g) \leq \frac{g}{C_\lambda D_F \exp[(\eta_g + \kappa_F)g^2]} \longrightarrow 0, \tag{116}$$

which proves the tail result.

For comparative statics, $H(g)$ rises pointwise with each of C_λ, η_g, D_F, and κ_F. Since H is strictly increasing in g, the unique solution to $H(\bar{g}) = 1$ falls with each parameter. Notice that this does not imply a global monotonicity in σ_F, which changes D_F and κ_F in opposite directions. □

Proof of Corollary 1. Because m is strictly increasing on $[0, \bar{g}]$, strictly decreasing on $[\bar{g}, \infty)$, and converges to zero, every $g > \bar{g}$ has a counterpart $\widehat{g} < \bar{g}$ such that $m(\widehat{g}) = m(g)$. Since A is nondecreasing, replacing g with $\widehat{g}$ weakly raises the objective. Thus every announcement to the right of the frontier is weakly dominated and an optimum can be selected in $[0, \bar{g}]$. If A is strictly increasing, the dominance is strict, so every optimum lies in that interval. If $A_g(\bar{g}) > 0$, the derivative of the objective at the frontier is $-A_g(\bar{g}) < 0$, so a positive optimum cannot equal $\bar{g}$ and must lie strictly below it.

For the linear-benefit problem, let g_s be an interior maximizer selected on the effective branch. Its first-order condition is

$$s m_g(g_s) = A_g(g_s). \tag{117}$$

Every accumulation point lies in $[0, \bar{g}]$. If an accumulation point were strictly below $\bar{g}$, then m_g would be bounded away from zero along a subsequence, whereas A_g is bounded on $[0, \bar{g}]$. The left side of (117) would therefore diverge, a contradiction. Hence $g_s \to \bar{g}$ from below. □

Proof of Corollary 2. Let $O_n^G = (1 - q_n)/q_n$ be the posterior odds in favor of the ordinary

mode. Additive log odds imply

$$O_n^G = O_0^G e^{-n\ell^g(g)}. \tag{118}$$

Since $1 - q_n = O_n^G/(1 + O_n^G) \le O_n^G$, the credibility-adjusted effect is bounded by

$$M_n(1 - q_n) \le \bar{M} O_0^G e^{-n[\ell^g(g)-\gamma]}. \tag{119}$$

This converges to zero when $\ell^g(g) > \gamma$. □

7.5 Credibility accounting and delivery evidence

Proof of Proposition 3. The message lowers override beliefs relative to the prior if and only if its likelihood ratio in (109) is below one. The unique positive equality point satisfies

$$\log \frac{\Sigma_{\mathsf{G}}}{\Sigma_{\mathsf{R}}} + \eta_g g_I^2 = 0, \tag{120}$$

which gives (57). Differentiating the posterior odds yields $q_g = 2\eta_g g q(1-q) > 0$ for $g > 0$.

Since $m_g = p(1 - H)$ on the effective branch,

$$\frac{\mathrm{d}q}{\mathrm{d}m} = \frac{q_g}{m_g} = \frac{2\eta_g g q}{\chi(1-H)}, \tag{121}$$

$$\frac{\mathrm{d}(1-\chi)}{\mathrm{d}m} = \frac{-\chi_g}{m_g} = \frac{2\kappa_F g(1-\chi)}{(1-q)(1-H)}. \tag{122}$$

At $\bar{g}$, all numerator terms and the remaining denominator terms are finite and strictly positive, while $1 - H \downarrow 0$. This proves the divergence asserted after the proposition. □

Proof of Proposition 4. Let

$$z \equiv \frac{y-\theta}{\sigma_\pi}, \qquad t \equiv \frac{\Delta_\theta}{\sigma_\pi}. \tag{123}$$

Conditional on an actual waiver, $z \sim \mathcal{N}(0,1)$. Relative to the waiver density, the honor-density ratio is

$$R(z,t) \equiv \frac{\phi_{\sigma_\pi}(y - \Theta^H)}{\phi_{\sigma_\pi}(y-\theta)} = \exp(-zt - t^2/2). \tag{124}$$

Set $u_R \equiv R - 1$. Because R is a density ratio,

$$\mathbb{E}_\theta u_R = 0, \qquad \mathbb{E}_\theta u_R^2 = t^2 + o(t^2). \tag{125}$$

Factoring out the waiver density from (61) gives

$$\frac{\mathcal{L}_j^\pi(y)}{\phi_{\sigma_\pi}(y-\theta)} = 1 + \chi_j u_R. \tag{126}$$

Hence

$$\log \frac{\mathcal{L}_{\mathsf{R}}^\pi}{\mathcal{L}_{\mathsf{G}}^\pi} = (\chi_{\mathsf{R}} - \chi_{\mathsf{G}})u_R + \frac{1}{2}(\chi_{\mathsf{G}}^2 - \chi_{\mathsf{R}}^2)u_R^2 + o(t^2). \tag{127}$$

Taking expectations and applying (125) proves (62).

Bayes' rule also gives the exact expression

$$q^\pi = \frac{q(1+\chi_{\mathsf{R}}u_R)}{1+\bar{\chi}_q u_R}. \tag{128}$$

Expanding the reciprocal to second order,

$$q^\pi = q\left[1 + (\chi_{\mathsf{R}} - \bar{\chi}_q)u_R + (\bar{\chi}_q^2 - \chi_{\mathsf{R}}\bar{\chi}_q)u_R^2\right] + o(t^2). \tag{129}$$

Using

$$\chi_{\mathsf{R}} - \bar{\chi}_q = -(1-q)(\chi_{\mathsf{G}} - \chi_{\mathsf{R}}), \tag{130}$$

$$\bar{\chi}_q^2 - \chi_{\mathsf{R}}\bar{\chi}_q = \bar{\chi}_q(1-q)(\chi_{\mathsf{G}} - \chi_{\mathsf{R}}) \tag{131}$$

and (125) yields (63). Applying the affine transition (28) multiplies the result by ρ_λ. □

Proof of Lemma 2. Let

$$\Delta\chi \equiv \chi_{\mathsf{G}} - \chi_{\mathsf{R}} > 0, \qquad \bar{\chi}_q = (1-q)\chi_{\mathsf{G}} + q\chi_{\mathsf{R}}, \qquad t = \frac{\Delta_\theta}{\sigma_\pi}.$$

The expansion in (128) can be written as

$$q^\pi - q = a_q u_R + b_q u_R^2 + o(t^2), \qquad a_q = -q(1-q)\Delta\chi, \qquad b_q = q(1-q)\bar{\chi}_q\Delta\chi. \tag{132}$$

Using (125),

$$\mathbb{E}_\theta(q^\pi - q) = b_q t^2 + o(t^2), \tag{133}$$

$$\mathbb{E}_\theta[(q^\pi - q)^2] = a_q^2 t^2 + o(t^2). \tag{134}$$

Let

$$\lambda_0' = (1-\rho_\lambda)\bar{\lambda} + \rho_\lambda q$$

be the next prior when the delivery miss is zero. Since $\lambda' - \lambda_0' = \rho_\lambda(q^\pi - q)$, a second-order expansion of continuation loss around λ_0' gives

$$\begin{aligned}\beta \mathbb{E}_\theta W_j^c(\lambda') &= \text{constant} + \beta W_{j,\lambda}^c \rho_\lambda b_q t^2 + \frac{\beta}{2} W_{j,\lambda\lambda}^c \rho_\lambda^2 a_q^2 t^2 + o(t^2) \\ &= \text{constant} + \frac{\varrho_j}{2}\Delta_\theta^2 + o(\Delta_\theta^2), \end{aligned} \tag{135}$$

where all value derivatives are evaluated at λ_0' and

$$\varrho_j = \frac{2\beta}{\sigma_\pi^2}\left[\rho_\lambda W_{j,\lambda}^c q(1-q)\bar{\chi}_q \Delta\chi + \frac{\rho_\lambda^2}{2} W_{j,\lambda\lambda}^c q^2(1-q)^2(\Delta\chi)^2\right]. \tag{136}$$

The conditions in (64) imply $\varrho_j > 0$. The first term in (136) prices the mean increase in public doubt; the second prices the additional posterior dispersion generated by noisy delivery.

To derive current loss, substitute $x = [\theta - (\beta e + u)]/\kappa$ into the quadratic objective:

$$\frac{1}{2}\left[\alpha\theta^2 + \frac{\vartheta}{\kappa^2}\{\theta - (\beta e + u)\}^2\right] = \text{constant} + \frac{A_\theta}{2}[\theta - d(e,u)]^2, \tag{137}$$

where A_θ and $d(e,u)$ are given in (66). The local waiver problem is consequently

$$\min_\theta \left\{\frac{A_\theta}{2}[\theta - d(e,u)]^2 + \frac{\varrho_j}{2}(\theta - \Theta^H)^2\right\}. \tag{138}$$

Its first-order condition is

$$A_\theta[\theta - d(e,u)] + \varrho_j(\theta - \Theta^H) = 0, \tag{139}$$

whose unique solution is (67). □

7.6 Low-noise fragility

Proof of Proposition 5. Write the ordinary-mode honor probability as

$$\chi_{\sigma_F}(g) = \frac{1}{1 + \exp\{[\Psi_G g^2/2 - \bar{F}_G]/\sigma_F\}}. \tag{140}$$

For $|g| < g^\star$, the exponent tends to $-\infty$ and $\chi_{\sigma_F}(g) \to 1$; for $|g| > g^\star$, it tends to $+\infty$ and $\chi_{\sigma_F}(g) \to 0$. This proves the pointwise cliff.

Let $m_I(g) = g[1-q(g)]$ and let $\bar{g}_I$ be its unique maximizer. If $\bar{g}_I < g^\star$, then on a compact neighborhood of $\bar{g}_I$ the honor probability and its derivative converge uniformly to one and zero, respectively. The first-order condition for the full frontier therefore converges to the inference-only condition, so $\bar{g}(\sigma_F) \to \bar{g}_I$.

Suppose instead that $g^\star < \bar{g}_I$. For every $\varepsilon > 0$, the inference-only marginal effect is strictly positive on $[0, g^\star - \varepsilon]$. On that compact set, the temptation term converges uniformly to zero, so the full marginal effect is positive for sufficiently small σ_F. For $g \geq g^\star + \varepsilon$, $\chi_{\sigma_F}(g) \to 0$ uniformly on compact sets and the implemented effect vanishes. The unique maximizer must therefore lie in $(g^\star - \varepsilon, g^\star + \varepsilon)$ for sufficiently small σ_F. Letting $\varepsilon \downarrow 0$ gives convergence to $g^\star$. The equality case follows by the same argument, proving (69). □

7.7 Speech design, policy mix, and implementation

Proof of Proposition 6. Substitute (70) into the delegated loss (71). When speech is allowed, the first-order conditions are

$$\alpha\kappa\pi + \vartheta(x - \bar{x}) = 0, \qquad \alpha\xi\pi + \Gamma_s g = 0. \tag{141}$$

Solving the two conditions with the constraint gives

$$\pi^1 = \frac{u + \kappa\bar{x}}{D + h_s}, \qquad g^1 = -\frac{\alpha\xi}{\Gamma_s}\pi^1, \tag{142}$$

where $D + h_s$ is the speech-augmented denominator. Under a ban, $g = 0$ and the remaining first-order condition gives

$$\pi^0 = \frac{u + \kappa\bar{x}}{D}. \tag{143}$$

Using the corresponding output gaps in the principal's loss $L^P = \frac{1}{2}(\alpha\pi^2 + \vartheta x^2 + \Gamma_s g^2)$ and simplifying yields

$$L^P_{\text{allow}}(u) - L^P_{\text{silent}}(u) = -\frac{\alpha h_s}{2D(D + h_s)}(u - \kappa\bar{x})(u + \kappa\bar{x}), \tag{144}$$

which is (72). Taking expectations establishes the criterion. If objectives are aligned, $\bar{x} = 0$ and the difference is weakly negative state by state, strictly so whenever $u \neq 0$. □

Proof of Proposition 7. Substituting (74) into the objective in (75) and differentiating with

respect to x and g gives

$$\alpha\kappa\pi + \vartheta x = 0, \qquad \alpha\beta\bar{\tau}\pi = 0, \tag{145}$$

which is (76). The second condition forces $\pi = 0$ (since $\alpha\beta\bar{\tau} \neq 0$), and the first then gives $x = 0$; the constraint delivers $g = -u/(\beta\bar{\tau})$. The announcement alone attains the first best for a cost-push shock. A natural-rate disturbance enters (5); when the current rate is unrestricted, it can be offset directly by i_t without altering the Phillips-curve allocation. □

Proof of Proposition 8. Let $y_t = (\pi_t, x_t)^\top$. From (77),

$$\mathbb{E}_t \pi_{t+1} = \frac{1}{\beta\delta}\pi_t - \frac{\kappa}{\beta\delta}x_t. \tag{146}$$

Using this expression, the feedback rule $i_t = \phi_\pi \pi_t$, and (78) gives

$$\mathbb{E}_t y_{t+1} = A y_t, \qquad A = \begin{pmatrix} \dfrac{1}{\beta\delta} & -\dfrac{\kappa}{\beta\delta} \\ \dfrac{\phi_\pi - \beta^{-1}}{\sigma} & 1 + \dfrac{\kappa}{\sigma\beta} \end{pmatrix}. \tag{147}$$

Both variables are nonpredetermined, so determinacy requires both eigenvalues of A to lie outside the unit circle, equivalently both eigenvalues of A^{-1} to lie inside. The second-order Jury conditions for A^{-1} are

$$\det A > 1, \qquad \det(A - I) > 0, \qquad \det(A + I) > 0. \tag{148}$$

Direct calculation yields

$$\det A = \frac{1 + \kappa\phi_\pi/\sigma}{\beta\delta} > 1, \tag{149}$$

$$\det(A - I) = \frac{\kappa}{\sigma\beta}\left(\frac{\phi_\pi}{\delta} - 1\right), \tag{150}$$

$$\det(A + I) = \frac{2\beta\delta\sigma + \delta\kappa + \kappa\phi_\pi + 2\sigma}{\beta\delta\sigma} > 0. \tag{151}$$

The first and third inequalities hold under the maintained parameter restrictions. The second holds if and only if $\phi_\pi > \delta = 1 - \bar{\tau}$. At equality one eigenvalue is one; below it one eigenvalue lies inside the unit circle. This proves Proposition 8. □

8 Numerical models for application section

The post-pandemic application solves the model forward along a fixed annual timeline. It features a general public state, persistent private information, and equilibrium-generated message distributions; the calibrated objects are collected below.

8.1 Post-pandemic application

The post-pandemic exercise is not a stationary policy but a single path solved forward from 2020 to 2027; the calibrated objects are collected in Table 2. Each year carries the outstanding announcement g_{t-1} and the override prior λ_t, together with the realized cost-push u_t. The realized shock peaks in 2021 and decays geometrically, $u_t = U_{2021}\,\rho_u^{t-2021}$, alongside a separate small 2020 disturbance; the central bank privately forecasts $\rho^{CB}u_t$ for the following year and discloses it, while the public prices the promise at its own forecast of u_{t+1}.

Within a year the central bank chooses the new announcement g_t on a grid over the authorization range $[\underline{g}, \overline{g}]$ to minimize

$$\begin{aligned}\mathcal{J}(g) = {} & \ell(\theta_t, e_t, u_t) + \beta p\, \ell(\theta^H, e_t, u_{t+1}) \\ & + \beta(1-p)\left\{\ell(\theta^W, e_t, u_{t+1}) + \frac{\varrho_t}{2}(\theta^W - \theta^H)^2\right\} \\ & + \beta\xi_\lambda[\, q - \lambda_t\,]_+,\end{aligned} \tag{152}$$

where ℓ is the period quadratic loss, p the prospective plan survival, $\varrho_t = \bar{\varrho}(1-\lambda_t)$ the imposed reputational-tether curvature, $\theta^H = \phi_{IT}u_{t+1} + g$ and θ^W the honor and waiver centers, and the last term is a reduced-form credibility penalty proportional to the message-stage belief revision $q - \lambda_t$—the discrete counterpart of the continuation cost in Section 4. In the bound years (2020–21) the honor branch keeps the central bank at the lower bound rather than tightening; a waiver is an early liftoff onto the reset path. Honoring is a conditional plan on the stance: the delivered activity gap is the one the announced center implied on the central bank's own forecast, and supply and expectational surprises pass through prices, so a miss driven by a shock the disclosed forecast did not contain is a forecast error, not a breach. Prospective survival and the fragility rows are priced on the letter of the center evaluated at the state the public believes will arrive.

The honor probability is priced from the optimized branch gap rather than from current loss. Around the target reference state, the code computes a finite-difference Hessian of the continuation-adjusted objective after jointly reoptimizing the reset center and the new

announcement. The implied Schur curvature is $\widehat{\mathcal{A}} = 11.69$ and the local gap is

$$\widehat{\Delta\mathcal{V}}_t = \frac{\widehat{\Psi}_t}{2}\left[\theta_t^H - d(e_{t+1}, u_{t+1})\right]^2, \qquad \widehat{\Psi}_{\text{ref}} = 11.55, \tag{153}$$

with the state dependence of $\widehat{\Psi}_t$ inherited from the imposed reputational tether. Logistic honoring probabilities are therefore $\chi_j = \text{expit}[(\bar{F}_j - \widehat{\Delta\mathcal{V}}_t)/\sigma_F]$. This differs materially from the static current-loss approximation: $\widehat{\Psi}_{\text{ref}}/A_\theta = 0.62$. Given the plan-consistent continuation expectation e_{t+1}, current expected inflation is evaluated once as $e_t = p\theta^H + (1-p)\theta^W$; it is not a self-referential fixed point. The message posterior q comes from the frozen Gaussian message densities with dispersions $\sigma_{\mathsf{G}}, \sigma_{\mathsf{R}}$; after inflation is realized, the override prior is updated from words and outcomes, $\lambda_{t+1} = (1-\rho_\lambda)\bar{\lambda} + \rho_\lambda q^D$. The 2020 announcement is imposed at +0.50 and the central bank reoptimizes every later year on its own information. The counterfactuals in Table 4 hold that baseline sequence of announcements and disclosed forecasts fixed from 2021 onward and re-price it under alternative public readings—varying the forecast credence ϖ, public persistence ρ^{pub}, threshold dispersion σ_F, a discretionary bias, the override prior, and the diffuseness of message densities. They are path-pricing exercises, not equilibria in which the central bank observes the altered interpretation and chooses a new path.

8.2 Data construction

One-year expectations and the target-weight regression. Quarterly observations span 2017Q1 through 2025Q4. Monthly year-over-year CPI inflation and the Federal Reserve Bank of Cleveland's model-based one-year expected-inflation measure are averaged within quarters. The headline regression sample ends in 2024Q4 so that every included quarter contains three monthly observations.

Taking four-quarter changes of the target-weight forecasting equation gives

$$\Delta_4\pi^e_{t,1y} = \alpha_\Delta + \tau\Delta_4\pi_t + u_t, \qquad \widehat{w}_0 = 1 - \widehat{\tau}. \tag{154}$$

The covariance estimator is Newey–West with four quarterly lags. The baseline contains 28 regression observations and gives $\widehat{\tau} = 0.244$ with standard error 0.031, hence $\widehat{w}_0 = 0.756$ with a 95-percent interval $[0.696, 0.817]$.

Meeting-day news. The market exhibit uses the two-year fitted breakeven inflation rate from the Federal Reserve Board's TIPS yield-curve dataset, the two-year constant-maturity Treasury yield from the H.15 release, and scheduled FOMC announcement dates from the

Board's meeting calendars. Spring-2020 crisis days are excluded. Bootstrap 90-percent intervals resample announcement and ordinary days separately using 2,000 draws.

The pre period is 2012–19, the surge period is 2021–23, and the post period is 2023–25. The overlap between the last two windows is intentional: the post window asks whether the relative-news statistic recovered as inflation declined, while the surge window preserves the full high-inflation episode.