Monetary Policy Delegation with Operational Rules and Sanctioned Promises*

Lukasz A. Drozd

Federal Reserve Bank of Philadelphia

lukasz.drozd@phil.frb.org

September 2026

Abstract

How can society achieve the commitment outcome when delegating monetary policy to a central bank with private information? We study standing targeting rules supplemented by discretionary forward guidance, where promises are disciplined by bounded sanctions. We show that partial codification can make delegation more demanding of the designer's knowledge. Under a coarse reactive rule, which prescribes an incomplete response, a smaller promise lowers expectations of both the promised departure and the rule's response. Yet the rule responds to the shock that arrives, so only the promised accommodation is forgone. This distortion shrinks promises and, for intermediate rules, leads to uninformative announcements for ordinary forecasts. Correcting it requires knowledge of the stabilization tradeoffs. With delivery assured, a strict, flat inflation target eliminates the game, allowing accountable guidance to convert discretion into commitment without a reporting correction calibrated to the model.

Keywords: monetary policy delegation, inflation targeting, forward guidance, central bank communication, signaling, accountability.

JEL codes: D82, D83, D86, E52, E58, E61.

*I thank Alessandro Dovis for helpful comments. All errors are my own. 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

Through research, forecasting, and judgment, central banks produce private information about the economy that often determines the course of future policy.[1] At the same time, constraints on the conduct of policy are known to be desirable: policy under commitment manages inflation expectations better and attains stabilization outcomes unavailable under discretion, and credible rules address the time-inconsistency problem (Kydland and Prescott, 1977; Barro and Gordon, 1983). Modern delegation pairs a standing inflation-targeting regime with forward guidance that uses that private information. How the mix should be designed is less clear. The need for guidance is immediate: a complete targeting rule, with every contingency written in, is not something society can prescribe.

The private information at issue is often not the size of a shock but its nature—how long it will last, what is driving it, whether it will pass through to wages and prices. That is what forward guidance is for: to move expectations away from what the standing rule alone would imply. Pandemic-era monetary response is one illustration of its use. Facing rising inflation in a weak recovery, central banks had to judge how much of it would persist. The judgment was communicated rather than codified. Across advanced economies the message was that the shock was transient, and that policy would look through temporary inflation to protect the recovery. The question we ask here is theoretical: whether such an arrangement gives the central bank the incentive to report what it knows, or to understate the forecast.

We study the problem as a delegation design (Walsh, 1995). Society has three instruments. The first is a standing targeting rule, which maps the realized shock that calls for accommodation into the inflation policy center. The second is a communication channel: society may let the central bank announce a promised departure from that center. The third is sanctioning power, exercised by an auditor who sees the announcement and the realized inflation miss and imposes a bounded penalty when they do not line up. The audit could also assess the announcement itself, and Section 4.2.6 characterizes the audit that would do so under a coarse rule; it requires society to know the stabilization tradeoffs it delegated in the first place, and a version calibrated to the wrong model leaves the reporting incentive uncorrected. Our main object is therefore the simple audit, which asks only whether a promise was kept, and the question is what guidance survives under it. Society chooses all three once, before the central bank observes anything, and cannot revisit them.

The setting is a standard New Keynesian economy. Cost-push shocks create a contem-

[1] Romer and Romer (2000) document the Federal Reserve's informational advantage in forecasting inflation; Faust and Wright (2009) compare Greenbook forecasts with real-time forecasting models. From the perspective of the current paper, the forecast can concern information that influences actual future policy.

poraneous tradeoff between inflation and the output gap, so some accommodation of price pressure is desirable (Clarida et al., 1999; Galí, 2015). Society's first choice is whether to install a standing rule at all. A central bank left to reoptimize cannot commit not to push output above what society wants, and the public knows it, which leaves a steady-state inflation bias of 13 percentage points in the leading numerical example.[2] An audit whose sanction probability rises with the output gap eliminates that bias, but it does not by itself set stabilization policy. Society chooses a standing rule whose slope ϕ determines that accommodation. A strict target sets $\phi = 0$. An ideal rule would set ϕ equal to the response that is optimal once the shock arrives, 0.95 in that example. A coarse reactive rule has a slope between these two benchmarks. It also stands for flexible inflation targeting as practised: the public forms its picture of the reaction function from years of policy, and that picture is as hard to change as a rule that is written down. The central bank privately observes a forecast of tomorrow's shock and announces a promised departure from the rule. If the forecast stays private, tomorrow's response is ϕ of the realized shock: 0.95 under the ideal rule, 0.70 under the baseline coarse rule, and 0 under a strict target. The last two miss the part of the shock the central bank already sees. Forward guidance is how those rules can be corrected. Once the forecast is public the planner's center is 0.90, not 0.95. A promise can implement 0.90; 0.95 is no longer optimal once the forecast is public.

Forward guidance poses two incentive problems. The central bank must deliver its promise and choose the promise society would prefer. An audit can discipline delivery, subject to the precision of monitoring and the available sanction. Reputation can also support delivery when future guidance depends on a record of fulfilled promises. Neither a fulfilled promise nor that record establishes that the right promise was chosen. The paper focuses on this reporting problem, which can arise under a coarse rule even when every promise is delivered.

An informative announcement changes expectations of both the promised departure and the rule's response to the forecast. Only the departure is fixed by the announcement: the rule responds to the shock that actually arrives. In the baseline calibration, the planner wants expected inflation of 0.90 per unit of forecast once that forecast is public. A rule that accommodates 0.70 therefore leaves a promise of 0.20. Suppose the public initially interprets promises according to the planner's schedule. A promise associated with a forecast one unit lower reduces expected inflation by 0.90: 0.70 through the anticipated rule response and 0.20 through the promise. The central bank gives up only the latter when

[2]Inflation and the output gap are measured in decimals throughout, so 0.01 is one percentage point. The number is $\lambda x^*/\kappa$, the fixed point in a repeated static economy, at $\lambda = 0.4$, $\kappa = 0.15$ and an output-gap target $x^* = 0.05$. The two-period analogue in the paper is 7.8 percentage points.

it delivers policy. Under a strict target, the same reduction in expected inflation requires forgoing 0.90 of accommodation, so the central bank faces society's tradeoff.

The public anticipates this incentive and interprets promises accordingly. An informative promise must deter imitation by a central bank with a higher forecast. The cost of imitation is the accommodation forgone by delivering that promise. The largest forecast can announce the planner's promise; lower positive forecasts announce less than the planner would to deter imitation. The amount left for guidance to supply— 0.20 per unit of forecast in the example—is its signaling capacity. For a range of intermediate rules, separation reduces the promise so much that ordinary forecasts prefer the uninformative announcement of zero departure. This occurs before the separating promise reaches zero. At the baseline, forecasts within 1.72 of zero share that announcement, accounting for 92 percent of forecasts. The promised departure is delivered, but the response the rule omits remains unsupplied. Section 4.2.1, Section 4.2.2, and Section 4.2.3 develop the numerical example.

The rule creates a tradeoff. A better automatic response reduces the need to reveal forecasts, which can itself be costly because revelation changes current expectations. In the pure-forecast case, the ideal rule supplies the desired response without guidance, and revealing the forecast only worsens current stabilization. A strict target leaves the entire response to the promise. With delivery assured, expected inflation then equals the promised departure regardless of what the public infers about the forecast, and the central bank chooses the planner's promise.[3] Under an intermediate rule, the benefit of correcting the response can exceed the cost of revelation. The reporting distortion can nevertheless prevent the central bank from providing that valuable guidance for ordinary forecasts.

The audit that repairs a coarse rule has two parts: an announcement correction that restores the planner's reporting incentives, and a restriction on admissible announcements that determines the public's reading. Both are computed from the stabilization tradeoffs society left to the central bank, and a correction computed in one model leaves a reporting wedge in another. Partial codification therefore raises what the designer must know before delegating, while the strict target requires nothing of the kind.

The welfare comparison also depends on the inflation index. Our numerical analysis holds that index and its cost-push shocks fixed. In richer economies, the same underlying disturbances can imply different inflation–output tradeoffs under different indices. A suitable index can accommodate relative-price adjustment while reducing the residual need for discretionary stabilization (Rubbo, 2023; La'O and Tahbaz-Salehi, 2022). The

[3]This choice is optimal under every public belief, including beliefs following announcements that are not made in equilibrium. Under a coarse rule in the revelation region, the central bank's preferred promise depends on the public's reading (Proposition 4).

advantage of the ideal rule over a strict target is therefore conditional on the chosen index. Our analysis characterizes the communication problem that remains after this choice; it does not quantify the gains from choosing a better index.

2 Related Literature

The paper builds most directly on Kartik (2009) and Halac and Yared (2022). Kartik (2009) shows how communication with lying costs can involve distorted reports and incomplete separation. Here the message is a policy promise, and its cost comes from delivering the promised response rather than from misrepresenting private information. The standing rule determines both the incentive to distort the announcement and the policy adjustment available to sustain separation: information conveyed by a promise changes expectations of the rule's response, while only the promised departure binds future policy. The reporting distortion arises from the interaction between the rule and discretionary guidance.[4] Halac and Yared (2022) study the design of fiscal rules when compliance must be sustained through bounded penalties that are costly to both the government and society. We adapt their limited-enforcement approach to a setting in which the central bank supplements a standing rule with its own policy promises. However, our penalties affect the delegated authority's incentives but carry no direct social cost. The design problem also requires distinguishing delivery from reporting: sanctions can make the central bank deliver its announcement without inducing it to choose the promise that society would prefer.

More broadly, the paper relates to the literature on how private central bank information shapes monetary policy communication. Bassetto (2019) shows that private information can give forward guidance an informational role even when announcements do not directly constrain future policy.[5] Here delegation adds promise enforcement: an announcement conveys information and specifies a departure from the standing rule whose delivery is disciplined by an audit. We study how the design of that rule affects both the social value of disclosure and the central bank's incentive to reveal its forecast.

The paper also relates to delegation through restrictions on feasible policy (Holmström, 1984; Alonso and Matouschek, 2008).[6] Here the standing rule provides an accountability benchmark: the central bank can announce departures from it, and an audit disciplines their

[4]The characterization draws on the logic of costly signaling (Mailath, 1987; Kartik, 2009). Pooling among central types with separation among extreme types also arises in Bernheim and Severinov (2003).

[5]Related work studies optimal policy and disclosure when the central bank has imperfect information (Iovino et al., 2022) and forecast communication when the public is uncertain about the precision of that information (Gáti and Handlan, 2025).

[6]Applications to the tradeoff between commitment and informed discretion include Athey et al. (2005), Amador et al. (2006), and Waki et al. (2018).

delivery through bounded sanctions. This use of payoff incentives connects the institution to central-bank contracts (Walsh, 1995), costly political override (Lohmann, 1992), and forecast and forward-guidance contracts (Gersbach and Hahn, 2014; Gersbach et al., 2019).

A related question is how policy rules affect public inference. Dovis and Kirpalani (2021) study how rule complexity affects monitoring when commitment capacity is uncertain. Elkina (2026) studies the tradeoff between discretion and reputation building when policy responds to private conditions and also signals a persistent hawk or dove type. In our setting, the central bank's stabilization preferences and enforcement arrangement are known. The private information concerns economic conditions, and the reporting problem arises because an announcement changes expectations of the standing rule's response to future shocks. It therefore remains even without uncertainty about the central bank's preferences or commitment capacity.

Section 4.2 contains the main results. The Online Appendix contains all omitted proofs and supporting results, the numerical example, and the infinite-horizon extension with recursive delivery; a supplementary appendix contains the operating-rule implementation and material the text cites in passing.

3 Environment

We study delegation within the canonical "three-equation" New Keynesian framework. There are three consumption periods, $t = 1, 2, 3$, preceded by a preliminary period 0 in which institutions are designed. Prices are sticky in periods 1 and 2 and flexible in period 3. The period of interest in this policy episode model is period 1, which we refer to as the *current stabilization period*; periods 2 and 3 close the model. The three-period structure is chosen for clarity of exposition; Online Appendix D extends the Ramsey decomposition and the conditional delivery analysis to an infinite horizon.[7] The steady-state monopolistic distortion is removed here by the standard subsidy of monopolistic producers, focusing attention on the stabilization distortion.

Phillips curves. Since the setup is standard, we directly state the log-linearized equilibrium conditions to describe the model economy.[8] The supply constraints are summarized

[7]The cost-push disturbances summarize the residual inflation–output tradeoff for a given inflation index. In multisector economies, their response to underlying disturbances depends on the index (Rubbo, 2023; La'O and Tahbaz-Salehi, 2022). We hold this choice fixed and take the resulting tradeoff as given.

[8]A detailed description of the model economy can be found in Galí (2015, chs. 3–5) and Woodford (2003, chs. 3, 4, and 6).

by the New Keynesian Phillips curves,

$$\pi_1 = \beta \mathbb{E}[\pi_2 \mid \mathcal{I}] + \kappa x_1 + v_1 + \varepsilon_1, \tag{1}$$

$$\pi_2 = \kappa x_2 + v_2 + \varepsilon_2. \tag{2}$$

Here π_t is the inflation rate realized in period t, $\mathbb{E}[\pi_2 \mid \mathcal{I}]$ is the expectation of future inflation formed by the public under its period-1 information set $\mathcal{I}$, x_t is the output gap relative to the flexible-price level, v_t is the realized cost-push shock at t, and $\varepsilon_t \sim N(0, \sigma_\pi^2)$ is independent noise.[9] The parameter $\kappa > 0$ is the slope of the Phillips curve and $\beta \in (0, 1)$ is the discount factor. Expected future inflation enters the period-1 equation but not the period-2 equation, since prices are flexible and inflation is at target in period 3 by assumption. The demand side is the usual IS equation, which we omit: with an unrestricted nominal interest rate it determines the rate that supports a given output-gap allocation and its expectations, and it adds no further layer of incentives that the delegation design would need to address.[10]

Up to second order, society evaluates a policy $\big(x_1(\mathcal{T}, q), x_2(\mathcal{T}, v_2, q)\big)$, whose arguments are defined below, according to a loss function given by

$$\mathcal{L} = \frac{1}{2}\mathbb{E}\Big[\theta_1^2 + \lambda x_1^2 + \beta\big(\theta_2^2 + \lambda x_2^2\big)\Big], \tag{3}$$

where $\lambda > 0$ is the fixed relative weight on output-gap stabilization and θ_t is the inflation policy center, inflation net of the implementation noise, $\pi_t = \theta_t + \varepsilon_t$. The loss is written on the policy centers θ_t rather than on realized inflation π_t. This is without loss: integrating out the implementation noise only adds a policy-independent constant. The central bank's objective differs in one respect. It values an output gap above potential, so that

$$\mathcal{L}^B = \frac{1}{2}\mathbb{E}\Big[\theta_1^2 + \lambda(x_1 - x^*)^2 + \beta\big(\theta_2^2 + \lambda(x_2 - x^*)^2\big)\Big], \tag{4}$$

with $x^* > 0$ a primitive known to society when it writes the mandate. Setting $x^* = 0$ recovers a central bank whose objective coincides with society's. The divergence is in the level society would like the gap to take, not in the weight it places on stabilization, and the audit-induced sanction is added to (4).

When we refer to optimal policy we mean period-0 Ramsey policy that minimizes the

[9] Throughout, $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and standard deviation σ.

[10] Determinate implementation is a separate question, discussed briefly in the supplementary appendix; the effective lower bound is outside the analysis, and Supplementary Appendix 2 discusses its implications. An instrument observed before expectations are formed would be an additional public signal and would enter the conditioning information in (9); we omit this consideration.

social loss subject to the Phillips curves—as discussed in more detail below.

Shocks and information structure. Conditional on information at the start of the policy episode, the cost-push disturbances (v_1, v_2) have mean zero and finite second moments, but

$$\mathbb{E}[v_2 \mid v_1] = \eta v_1, \qquad \eta \geq 0. \tag{5}$$

When $\mathrm{Var}(v_1) > 0$, the coefficient $\eta = \mathrm{Cov}(v_1, v_2)/\mathrm{Var}(v_1)$ thus measures the future-news content of the stabilization state. The public knows η, which bundles the two kinds of information in the model; that is, revealing the nowcast v_1 also reveals part of a forecast of v_2. (Joint normality is sufficient for this linear structure; what follows uses only the linear structure parameterized by η.)

At the onset of period 1 the central bank observes the realized current cost-push shock v_1, its nowcast, and receives a private signal s about the future cost-push shock, its forecast. The central bank's information type as of period 1 is thus the pair $\mathcal{T} \equiv (v_1, s)$. The signal s contains pure information about the future shock. We define this residual forecast as the central bank's update to the expected future shock after accounting for the predictable component: $s \equiv \mathbb{E}[v_2 \mid \mathcal{T}] - \eta v_1$, where s is independent of v_1. The central bank's forecast structure thus takes the following form:

$$v_2 = \eta v_1 + s + \varepsilon, \qquad \mathbb{E}[s \mid v_1] = 0, \qquad \mathbb{E}[\varepsilon \mid \mathcal{T}] = 0, \tag{6}$$

where $\varepsilon \equiv v_2 - \mathbb{E}[v_2 \mid \mathcal{T}]$ is the innovation—unknown to everyone in period 1—and the stated independence of s ensures that a separate forecast disclosure reveals no information about current conditions through higher moments.[11] The main results concern the forecast s; the nowcast, which carries news about the future shock, is taken up in Online Appendix B.3.1. The public observes neither v_1 nor s when they are useful in decisions (each firm observes and prices on its own idiosyncratic draw of the cost-push disturbance). The support of private information is bounded and common knowledge: $v_1 \in [-\bar{v}, \bar{v}]$ and $s \in [-S, S]$, with symmetric distributions. The calibration sets $S = 3\sigma_s$.

The divergence has a consequence that motivates the institution. Left to itself in period 2, the central bank accommodates more than society would, and by an amount that does not scale with the shock: its preferred response carries an intercept proportional to x^*. Write $\bar{\pi}$ for the inflation the economy is anchored at beyond the episode, so that the period-2 center is $\beta\bar{\pi} + \kappa x_2 + v_2$. Under any standing rule the anchor is zero, because the rule

[11] While the audit is evaluated under risk neutrality, the social loss function is convex, so higher moments matter.

prescribes no intercept and promises are symmetric; we therefore set $\bar{\pi} = 0$ throughout and it plays no part in what follows. Under unconstrained discretion it is instead a fixed point, the central bank's intercept feeding the anchor and the anchor feeding back into its own problem, which multiplies the distortion by $1/(1 - \beta C)$—about eleven at the baseline calibration. The public anticipates it, expected inflation is positive even when no shock is expected, and period-1 policy must lean against inflation nobody wanted. This is why the mandate contains a rule at all, and it gives the audit work to do even when the central bank knows nothing the public does not.

Write $b \equiv C\kappa x^*$ for the amount by which the central bank's preferred response exceeds society's. The mandate below removes it in every period, so the delegated central bank delivers what society would and the announcement a zero forecast makes is zero.

The central bank and the public know the structural model when they make decisions and form expectations. Society need not share that knowledge when it designs the institution in period 0. Section 3.2 specifies this separate source of uncertainty.

Timing and communication. In period 0, once and for all, society sets up the central bank as an institution, writes its mandate, and appoints a public auditor as the automatic, state-dependent executor of its will. The mandate comprises a standing rule, a format for its forward guidance announcements, and an audit-based oversight. The audit comprises a probability function, executed by the auditor, and a fixed sanction imposed on the central bank in the event of an audit; the sanction reduces the central bank's payoff when an audit takes place and has no effect otherwise. The central bank is risk neutral and minimizes (4), which differs from the social loss only through x^*. Accordingly, the audit-induced sanction adds a bounded loss that modifies the central bank's incentives relative to the baseline without sanctions.

At the onset of period 1, after observing its type, the central bank chooses a message $q = q(\mathcal{T})$ to communicate to the public. The public forms its expectation of future inflation, firms set prices, and the central bank chooses its output-gap policy for the first period, $x_1(\mathcal{T}, q)$, and later, in the second period, $x_2(\mathcal{T}, v_2, q)$; in expectation the central bank thereby chooses the inflation policy centers $\theta_1(\mathcal{T}, q)$ and $\theta_2(\mathcal{T}, v_2, q)$, as implied by the mean Phillips curves

$$\theta_1(\mathcal{T}, q) = v_1 + \kappa x_1(\mathcal{T}, q) + \beta e_1(q), \tag{7}$$

$$\theta_2(\mathcal{T}, v_2, q) = v_2 + \kappa x_2(\mathcal{T}, v_2, q), \tag{8}$$

where

$$e_1(q) \equiv \mathbb{E}\left[\theta_2(\mathcal{T}, v_2, q) \mid q\right] \tag{9}$$

is the public's message-induced expectation of the period-2 inflation policy center. The expectation operator averages over the conditional distribution of $(\mathcal{T}, v_2)$ given q. Because the noise has mean zero, $e_1(q) = \mathbb{E}[\pi_2 \mid q]$. The message is at this point an arbitrary object (a function, scalar, list) and can convey any part of the central bank's private information, including no information. The disclosure comparison concerns information available when e_1 is formed (more information or even full revelation can be assumed afterwards).

The public knows the communication strategy of the central bank and forms Bayesian posterior means of the shocks from the message, $\widehat{v}_i(q) \equiv \mathbb{E}[v_i \mid q]$ for $i \in \{1, 2\}$. Because q is a function of the central bank's type $\mathcal{T}$, iterated expectations and (6) give $\widehat{v}_2(q) = \eta\widehat{v}_1(q) + \mathbb{E}[s \mid q]$. For any message generated from v_1 alone, therefore,

$$\widehat{v}_2(q) = \eta\widehat{v}_1(q). \tag{10}$$

The auditor observes the announcement, realized inflation, and an exact measurement of the period-2 state. It evaluates delivery against the announced center and may also condition sanctions on the realized state to shape reporting incentives (Section 3.2 specifies the audit in full). Online Appendix Figure A.1 summarizes the timing and who knows what.

3.1 Ramsey Benchmark — What Society Desires

Before introducing the frictions of delegation, we characterize the first-best commitment outcome. This unrestricted Ramsey benchmark defines the allocation society wishes to attain and reveals the precise role of communication.

It is convenient to restate the period-2 loss in a form that exposes the tradeoffs. Eliminating x_2 from the period-2 loss using (8), we obtain $\frac{1}{2}[\theta_2^2 + \frac{\lambda}{\kappa^2}(\theta_2 - v_2)^2]$, and minimizing over θ_2 gives $\theta_2 = Cv_2$ with $C \equiv \lambda/(\kappa^2 + \lambda)$. The composite parameter C is the share of a cost-push shock that discretionary policy accommodates ex post (looks through and passes on to inflation). Completing the square around this discretionary benchmark gives the following reformulation of the loss function as of period 2:

$$\frac{1}{2}\left[\theta_2^2 + \frac{\lambda}{\kappa^2}(\theta_2 - v_2)^2\right] = \frac{A}{2}(\theta_2 - Cv_2)^2 + \frac{C}{2}v_2^2, \qquad A \equiv \left(\lambda + \kappa^2\right)/\kappa^2. \tag{11}$$

The formula shows the value of commitment and its ex post cost. The loss from

committing ahead of time to any center other than the ex-post preferred Cv_2 is quadratic in the deviation, but such a commitment can be useful because the Phillips curve links current inflation π_1 to expected future inflation $\mathbb{E}[\pi_2 \mid \mathcal{I}]$. The weight A on the deviation is tied to the accommodation share by $A(1-C)=1$: the stronger the preference for accommodation ex post (higher C), the costlier a given deviation from it.

We refer to it as the *unrestricted* Ramsey policy, because we later also consider a *rule-constrained* Ramsey problem, which takes the rule as given and manages forward guidance under commitment. These benchmarks describe the desired allocation in each model; the information available when society designs the institution is specified in Section 3.2.

In period 0, before any private information is observed, the unrestricted Ramsey planner commits to a policy vector given by

$$\mathcal{P}_0 \coloneqq \big(q(\mathcal{T}),\, x_1(\mathcal{T},q),\, x_2(\mathcal{T},v_2,q),\, \theta_1(\mathcal{T},q),\, \theta_2(\mathcal{T},v_2,q)\big)\,. \tag{12}$$

The policy vector comprises an announced communication strategy $q(\cdot)$ generating message q, and output-gap and inflation-center schedules $x_t(\cdot)$, $\theta_t(\cdot)$ that depend on everything the central bank knows when it acts: its type, the message, and in period 2 the realized shock. All of these objects are known to the public. The ex ante desired optimum is the $\mathcal{P}_0$ that minimizes the expected discounted social loss:

$$\frac{1}{2}\mathbb{E}\Big[\theta_1(\mathcal{T},q)^2+\lambda x_1(\mathcal{T},q)^2+\beta\big(\theta_2(\mathcal{T},v_2,q)^2+\lambda x_2(\mathcal{T},v_2,q)^2\big)\Big], \tag{13}$$

subject to

$$\begin{aligned} q &= q\,(\mathcal{T})\,,\\ e_1(q) &= \mathbb{E}[\theta_2(\mathcal{T},v_2,q)\mid q],\\ \theta_1(\mathcal{T},q) &= v_1+\kappa x_1(\mathcal{T},q)+\beta e_1(q),\\ \theta_2(\mathcal{T},v_2,q) &= v_2+\kappa x_2(\mathcal{T},v_2,q). \end{aligned} \tag{14}$$

The allocation is the familiar optimal commitment solution, written in the form delegation requires.

Lemma 1 (Message compression)**.** *(i) Fix a public message q and an expected period-2 center $e_1(q)$. Among square-integrable schedules $\theta_2(\mathcal{T},v_2,q)$ satisfying*

$$\mathbb{E}\left[\theta_2(\mathcal{T},v_2,q)\mid q\right]=e_1(q), \tag{15}$$

the unique continuation-loss-minimizing schedule, up to almost-sure equality, is

$$\theta_2(\mathcal{T}, v_2, q) = Cv_2 + g(q), \qquad \text{where} \qquad g(q) = e_1(q) - C\widehat{v}_2(q). \tag{16}$$

(ii) Fix any communication strategy that yields a message q, and write $M(q) \equiv \widehat{v}_1(q) + \beta C\widehat{v}_2(q)$ and $B \equiv 1/(A+\beta C)$. In the Ramsey program (13)–(14), the optimal shift attached to message q and the expectation it induces are

$$g^R(q) = -BC\, M(q), \qquad e_1^R(q) = BC\left[A\widehat{v}_2(q) - \widehat{v}_1(q)\right], \tag{17}$$

and the minimized period-0 loss under the strategy is

$$\mathcal{L}(q) = Z + \frac{C}{2}\mathbb{E}\left[\mathrm{Var}(v_1 \mid q)\right] + \frac{1}{2}ABC\,\mathbb{E}\left[M(q)^2\right], \qquad Z \equiv \frac{\beta C}{2}\mathbb{E}[v_2^2], \tag{18}$$

where Z does not depend on the strategy.

Proof. See Online Appendix A.1.

Any variation in the shift beyond the scalar leaves the public's expectation unchanged and raises the quadratic loss, as in Fujiwara and Waki (2022), so optimal policy is an ex-post response plus a promised departure from it, which is a standing rule and forward guidance. The inflation policy centers are

$$\theta_1^R(v_1, q) = C\left[v_1 + \beta e_1^R(q)\right], \qquad \theta_2^R(v_2, q) = Cv_2 + g^R(q), \tag{19}$$

the current center under the induced expectation and the ex-post response shifted by the promise. Society weighs today's benefit of lower expected inflation, at weight C through the Phillips curve, against tomorrow's cost of departing from the ex-post ideal Cv_2, at weight A, and leans against the pressure $M(q)$ the revealed information creates by the fraction BC.

That pair of centers is the object a standing rule and a sanctioned promise can implement. The next subsection specifies the mandate that can implement objects of that form.

3.2 Delegation Design Problem and Setup

The central bank chooses inflation and output gap policy sequentially, and without a mandate an announcement cannot commit it to anything. What it does instead is the benchmark the rest of the paper is measured against.

Lemma 2 (Discretion without sanctions)**.** *Consider the scalar benchmark $s = 0$ almost surely. Suppose the central bank chooses sequentially and no sanction is imposed, and let $\eta > 0$. In every equilibrium: (i) in period 2 the central bank sets its static best response $\theta_2 = C(v_2 + \kappa x^*)$ whatever it announced during the current stabilization period; (ii) its announcements convey no policy-relevant posterior-mean information about the current state, $\widehat{v}_1(q) = 0$ almost surely, implying $e_1(q) = C\kappa x^*$; (iii) in period 1 the central bank sets*

$$\theta_1 = Cv_1 + C\kappa x^*(1 + \beta C), \qquad x_1 = -\frac{\kappa v_1}{\kappa^2 + \lambda} + Cx^*\big[1 - \beta(1 - C)\big]. \tag{20}$$

The equilibrium allocation is the static discretionary response, displaced by an intercept proportional to x^* that does not scale with the shock. At $x^* = 0$ it is the static discretionary response itself.

Proof. See Online Appendix A.1.

The lemma shows what the bias costs and that discretion cannot shed it. Under discretion the intercept is present in period 2, is inherited by expectations because announcements reveal nothing, and is inherited again by period 1 through the Phillips curve. A standing rule alone does not remove it: the central bank carries the intercept into its promise, which the audit then enforces (Corollary 1), and into period-1 policy, which no promise covers. The gap component of the audit removes it in every period (Lemma 5 in Online Appendix A.5). The case for the rule itself is made in Section 4: it fixes how much of the response is codified, and therefore what guidance must do and whether it can.

Corollary 1 (Promise without a rule)**.** *Suppose there is no standing rule and no charge, the central bank announces a period-2 center g in period 1 and delivers it, and the public expects it, $e_1 = g$. In the scalar benchmark $s = 0$ the announcement is*

$$g = BC\big[A\widehat{v}_2 - v_1\big] + ABC^2\kappa x^*, \tag{21}$$

the Ramsey promise of Lemma 1(ii) plus ABC times the discretionary intercept $C\kappa x^*$ of Lemma 2. At $\beta = 1$ the factor is $AC/(A + C)$.

Proof. See Online Appendix A.1.

An obeyed promise commits delivery one period ahead. It does not change what the central bank wants to deliver: period 2 is the last period, no expectation remains to be managed, and the central bank's own loss is minimized at $C(v_2 + \kappa x^*)$. So it promises the intercept and delivers it, and commitment through the announcement removes only the period-1 leak, the factor $ABC = 0.903$. The same holds in the repeated economy. In its stationary Nash equilibrium, each period's promise chosen taking the next period's as given, inflation is 4.8 percentage points against 7.8 under discretion at the baseline discount factor, and 6.5

against 13.3 at $\beta = 1$, so the bias is not a discounting effect (Online Appendix A.1). Full commitment with the same objective would leave 0.5 percentage points at the baseline and none at $\beta = 1$; full commitment is what delegation cannot supply. A standing rule pins the center when nothing is announced, whatever its slope, and the gap component of the audit removes the intercept from the objective. The two do different work, and the strict target is the rule of slope zero, not the absence of a rule.

To recover the gains from commitment through forward guidance, society must align incentives on two margins: the central bank must choose the promise the planner would choose, and it must deliver that promise ex post. Enforcing delivery alone need not align the choice of promise. Society designs the mandate in period 0, before the central bank observes its private information. It chooses a standing rule, an announcement format, and an audit with bounded sanctioning power. These instruments are fixed thereafter. Society's objective is to minimize the stabilization loss (3); sanctions affect the central bank's incentives but carry no social cost. The designer's objective is to preserve the gains from commitment while allowing policy to use the central bank's information, subject to public verifiability and finite sanctioning power.

The designer lacks knowledge of the structural model of the economy. Let $\mathcal{M}$ denote the set of models society cannot distinguish in period 0. Each $\Omega \in \mathcal{M}$ specifies the stabilization parameters and the shock distribution, subject to the information structure above; the discount factor β is common to these models. Society must choose one mandate for all models in $\mathcal{M}$. The central bank and the public know the true model when they act, so this uncertainty concerns institutional design only. The public also does not observe the central bank's type $\mathcal{T}$, since it is revealed after period 0.

The mandate (delegated design) involves the following three restricted instruments.

1) Standing targeting rule. The standing rule prescribes an inflation policy center as a function of period-2 cost-push shock:

$$\bar{\theta}_\phi(v_2) = \phi v_2, \qquad \phi \geq 0. \tag{22}$$

We refer to $|C - \phi|$ as the rule's *coarseness*. The *ideal* rule is $\phi = C$, the *strict target* is $\phi = 0$, and rules with $0 < \phi < C$ are *coarse reactive rules*, abbreviated to *coarse rules*. The strict target is also coarse under the distance measure $|C - \phi|$, but prescribes no response to the realized shock. The ideal rule prescribes the response that minimizes society's period-2 loss after the shock arrives, for the specified index and model. Its slope coincides with the static discretionary response once the inflation bias is removed. The complete Ramsey

policy also includes the commitment adjustment $g^R(q)$. We assume that rules of this class, expressed as functions of the observed state, are the part of policy that the mandate can codify, communicate, and verify.

2) Announcements. The central bank announces a departure $g \in \mathcal{G}$ from the standing rule, where $\mathcal{G} = \mathbb{R}$ unless the mandate restricts the admissible announcements. The format admits only this policy promise: the central bank communicates its private information through the departure it chooses, rather than through a separate forecast report. Within the equilibrium class of Section 4.2.3 for $0 \leq \phi \leq C$, adding a separate message to which neither the rule nor the sanction responds does not improve equilibrium welfare.[12]

Under *assured delivery*, the delivered center equals the announced one (Section 4.1.2):

$$\theta_2(v_2, g) = \bar{\theta}_\phi(v_2) + g = \phi v_2 + g. \tag{23}$$

The rule responds to the shock realized in period 2, while the promise is fixed in period 1. Under the ideal rule, the planner's optimal continuation policy conditional on a public message has precisely this form: a scalar departure from Cv_2 (Lemma 1; Corollary 4 in Online Appendix A.1).[13] For other rule slopes, a promise chosen before the shock arrives can repair only the forecastable part of the rule's error.

We call the planner's optimum subject to (23) the *rule-constrained* Ramsey benchmark. At $\phi = C$, this continuation-policy restriction is without loss conditional on the planner's public message. The delegated central bank faces the same rule and announcement format but chooses sequentially. Delivery of the prescribed center must therefore be supported by incentives.

3) Audits. The central bank announces after observing its type $\mathcal{T}$ and chooses its period-2 policy center $\theta_2(\mathcal{T}, v_2, g)$ after observing v_2. The rule and the announcement prescribe the center $\phi v_2 + g$, but the central bank need not find it optimal to deliver that center ex post. An *obedience audit* disciplines adherence to the promise.

[12] Lemma 10 in Online Appendix A.15 enlarges the announcement to a promise plus such a separate message. The enlarged format supports no fully revealing outcome that the promise alone does not, and within the equilibrium class of Section 4.2.3, its best equilibrium is the one the promise alone produces. Remark 5 covers the strict target and the ideal rule.

[13] A promise is a number for two reasons. Lemma 1 shows that one number per message is all the planner wants, and one number is what an audit can verify: the auditor compares realized inflation with an announced center and never observes beliefs. A message that carries other content as well reduces to the promise within the equilibrium class of Section 4.2.3 (the supplementary message of Lemma 10); which centers may be announced is the format, an instrument of Proposition 6. The strict target is the rule of slope zero, not the absence of a rule: Corollary 1 prices the promise without one.

The auditor does not observe the central bank's private forecast, but observes the announcement, realized inflation, and the period-2 shock, which it measures exactly. Actual inflation satisfies $\pi_2 = \theta_2 + \varepsilon_2$, so implementation noise remains even when the prescribed center is delivered. The audit is a pair (Ψ, ρ): a finite sanction stake $\Psi > 0$ and a publicly specified probability $\rho_t \in [0,1]$ in each period that may depend on all the auditor's observations. The probability has two parts. In the delivery period it depends on the announcement and the residual,

$$\rho(g, z, v_2), \qquad z \equiv \pi_2 - [\phi v_2 + g], \tag{24}$$

where z is the audit residual, the inflation miss relative to the announced center. In every period it also rises linearly with the output gap, $(\lambda x^*/\Psi)(x_t - \underline{x})$, with $\underline{x}$ a constant and the probability capped at one; the cap binds only at gaps far outside anything the model produces, so it plays no role below. The expected sanction is bounded by Ψ. The sanction is an institutional cost—a governance, legal, or career consequence, or public censure—rather than a transfer.

The gap component is the counterpart of a linear contract in the spirit of Walsh (1995), levied through the audit rather than as a transfer. With a risk-neutral central bank its expected cost is $\lambda x^* x_t$ plus a constant, so it restores society's objective itself and not only the static first-order condition: the central bank's loss (4) plus the expected sanction of the gap component equals the social loss plus a constant, so the announcement and delivery problems below are those of a central bank with $x^* = 0$, its preferred response is society's in both periods, and a zero forecast announces exactly zero (Lemma 5 in Online Appendix A.5). Society must know x^*; the component does not depend on the forecast, and since the Phillips curves give $\kappa x_2 = \theta_2 - v_2$ and $\kappa x_1 = \theta_1 - v_1 - \beta e_1$, it is levied on $\pi_2 - v_2$ and on $\pi_1 - v_1 - \beta e_1$, both of which the auditor observes. The gap is the right base. A component on inflation would restore the response but not the announcement game: it would leave $\beta\mu^*$, $\mu^* = \lambda x^*/\kappa$, in the marginal cost of a promise, and the central bank would under-promise to reduce what it pays (Remark 3).

The central bank is risk neutral, so what enters its objective is the expected sanction:

$$\mathcal{A}(g, \theta_2, v_2) \equiv \Psi\, \mathbb{E}\big[\rho(g, z, v_2) \mid g, \theta_2, v_2\big]. \tag{25}$$

It bears this expected sanction in addition to its stabilization loss.

We use the audit in two steps. First, the obedience audit conditions its probability only on the delivery residual, $\rho_C(z)$. This lets us study which promises the central bank chooses when accountability disciplines delivery alone. Second, Section 4.2.6 adds an

announcement-dependent component to correct reporting incentives under coarse rules. The resulting audit, constructed in Online Appendix A.17, has the form $\rho(g, z)$: one component disciplines the choice of promise, and the other its delivery.

The specification (24) also permits separate dependence on the realized shock. Under the ideal rule, Lemma 7 in Online Appendix A.8 shows that this additional dependence is unnecessary for attaining the benchmark. For coarse rules, we construct the announcement-and-delivery audit described above rather than characterize the optimal audit over the entire class $\rho(g, z, v_2)$.

Policy and announcement choices. Given the mandate, the central bank chooses policy sequentially. After observing the period-2 shock, it chooses the output gap and inflation policy center to solve

$$\begin{aligned} \mathcal{V}_2(g, v_2) \equiv \min_{x_2, \theta_2} \quad & \frac{1}{2}\left[\theta_2^2 + \lambda x_2^2\right] + \mathcal{A}(g, \theta_2, v_2) \\ \text{subject to} \quad & \theta_2 = \kappa x_2 + v_2. \end{aligned} \tag{26}$$

The constraint is the period-2 Phillips curve (8). The announced center $\phi v_2 + g$ enters the sanction through the audit residual; it is not imposed as a constraint on the central bank's choice.

In period 1, given its type $\mathcal{T} = (v_1, s)$ and the expectation $e_1(g)$ induced by an announcement g, the central bank chooses current policy to solve

$$\begin{aligned} \mathcal{J}(v_1, s, g) \equiv \min_{x_1, \theta_1} \quad & \frac{1}{2}\left[\theta_1^2 + \lambda x_1^2\right] + \beta\, \mathbb{E}\left[\mathcal{V}_2(g, v_2) \mid \mathcal{T}\right] \\ \text{subject to} \quad & \theta_1 = \beta e_1(g) + \kappa x_1 + v_1. \end{aligned} \tag{27}$$

The constraint is the period-1 Phillips curve (7). Anticipating these policy choices, the central bank selects its announcement to solve

$$\min_{g \in \mathcal{G}} \mathcal{J}(v_1, s, g). \tag{28}$$

These problems separate the two incentive margins: the audit affects whether the central bank delivers its promise and, through the resulting continuation loss, which promise it chooses to make.

Design across models. For each model, we characterize the delegated outcome and the accountability needed to attain the relevant commitment benchmark. Under model un-

certainty, we then ask whether the same accountability provisions align incentives across $\mathcal{M}$ without recalibration. The comparison in Section 4.2.6 focuses on reporting: holding the standing rule, report interval, candidate announcement schedule, and obedience audit fixed, can one announcement correction align reporting incentives in every model? This is a compatibility requirement across models and requires no prior over them. Assured delivery isolates this reporting question; the information and sanctioning capacity needed to enforce delivery remain part of the design problem.

3.3 Equilibrium Definition

Fix a model $\Omega \in \mathcal{M}$ and the mandate chosen in period 0: the standing rule, the admissible announcement set $\mathcal{G}$, and the audit (Ψ, ρ), where $\Psi > 0$ is the sanction and $\rho(g, z, v_2)$ is its probability of imposition, as stated in (24). A perfect Bayesian equilibrium consists of an announcement strategy $G(\mathcal{T}) \in \mathcal{G}$, policy functions under delegation

$$\mathcal{P}_1^C(\mathcal{T}, g) = \big(x_1^C(\mathcal{T}, g), \theta_1^C(\mathcal{T}, g)\big), \qquad \mathcal{P}_2^C(\mathcal{T}, v_2, g) = \big(x_2^C(\mathcal{T}, v_2, g), \theta_2^C(\mathcal{T}, v_2, g)\big),$$

a public posterior probability measure $\mu(\mathrm{d}\mathcal{T} \mid g)$ over the central bank's type, and an expectations function $e_1(g)$, such that the following conditions are satisfied:

1. *Optimal delivery.* For every type $\mathcal{T}$, announcement g, and realized shock v_2, the policy pair $\mathcal{P}_2^C(\mathcal{T}, v_2, g)$ solves (26).

2. *Optimal current policy and announcement.* For every type $\mathcal{T} = (v_1, s)$ and announcement g, the policy pair $\mathcal{P}_1^C(\mathcal{T}, g)$ solves (27), given $e_1(g)$. The announcement $G(\mathcal{T})$ solves (28).

3. *Consistent beliefs.* Let F denote the prior probability measure over types and F_G the probability measure over announcements induced by G:

$$F_G(B) = \int \mathbf{1}\{G(\mathcal{T}) \in B\}\, F(\mathrm{d}\mathcal{T}).$$

   The public's belief $\mu(\cdot \mid g)$ is a probability measure over types. Bayes' rule requires that, for every measurable set of types A and announcements B,

$$\int_B \mu(A \mid g)\, F_G(\mathrm{d}g) = \int_A \mathbf{1}\{G(\mathcal{T}) \in B\}\, F(\mathrm{d}\mathcal{T}).$$

In particular, at an announcement made with positive probability,

$$\mu(A \mid g) = \frac{\int_A \mathbf{1}\{G(\mathcal{T}) = g\}\, F(\mathrm{d}\mathcal{T})}{\int \mathbf{1}\{G(\mathcal{T}) = g\}\, F(\mathrm{d}\mathcal{T})}.$$

At every announcement, including those off path, beliefs assign probability only to types in the support of F. Conditional on each type, the distribution of v_2 is given by the information structure.

4. *Consistent expectations.* At every announcement, public expectations equal expected delivered policy:

$$e_1(g) = \int \mathbb{E}\left[\theta_2^C(\mathcal{T}, v_2, g) \mid \mathcal{T}\right] \mu(\mathrm{d}\mathcal{T} \mid g).$$

Thus expectations reflect optimal delivery, including any departure from the announced center. Under assured delivery, $\theta_2^C(\mathcal{T}, v_2, g) = \phi v_2 + g$, and this condition reduces to

$$e_1(g) = \phi\, \mathbb{E}_\mu[v_2 \mid g] + g,$$

where the expectation integrates over the public posterior and the conditional distribution of the future shock.

To characterize the uninformative announcement, we restrict attention to symmetric monotone equilibria satisfying criterion D1 (Cho and Kreps, 1987; Banks and Sobel, 1987). These restrictions are defined when the announcement game is analyzed in Section 4.2.3; their formal implications are derived in Online Appendix A.15.

4 Results

This section compares standing rules by asking three questions of each: what communication about the forecast society wants; what the delegated central bank does on its own, with only the obedience audit in place; and what the mandate must add to close the gap. Section 4.1 answers the three questions for the ideal rule, separating the value of revelation from the enforcement of promises. It also introduces the audit used in the subsequent analysis. Section 4.2 answers them for coarse rules and the strict target and contains the main results: the reporting distortion a coarse rule creates, the equilibrium of the announcement game it induces, the welfare comparison across rules, and the mandate that restores guid-

ance. Online Appendix B.3.1 collects what each result says about guidance about current conditions.

The central bank's private information is the pair (v_1, s), and the results treat its two components one at a time. Section 4.1 and Section 4.2 concern the forecast s; except where a statement says otherwise, they hold the current state public and equal to zero, so that the news content η plays no role (a public ηv_1 is a known shift of the future shock), and the reader may set $\eta = 0$ throughout. Where a statement's coda turns to the nowcast, the separate forecast is absent and the private information is v_1 alone. Its news content ηv_1 makes guidance about current conditions informative about tomorrow's shock as well. Each statement specifies the case it covers.

4.1 Ideal Rule Communication

A standing rule determines both the response left for guidance to supply and the response the public anticipates when guidance reveals a forecast. Under the ideal rule, $\phi = C$, there is no omitted response to repair. Revelation still changes expectations, but the forecast enters the central bank's announcement loss only as an additive constant. We first establish why the forecast remains private, then examine how sanctions enforce the promises that policy requires.

4.1.1 Promises Carry No Forecast

Under the ideal rule tomorrow's response to the shock is Cv_2 whatever is announced, and by Lemma 1 the promise $g^R(q) = -BCM(q)$ is the planner's only instrument for lowering expected inflation today. Revealing the forecast does nothing for that instrument: the public expects the rule to accommodate C of any shock it learns about, so learning $\widehat{s}$ changes the rule's contribution to expected inflation by $C\widehat{s}$ without improving its response tomorrow. The following corollary of Lemma 1 measures the resulting cost and recovers the disclosure-cost mechanism of Fujiwara and Waki (2022).

Corollary 2 (Cost of forecast revelation under the ideal rule)**.** *Fix any communication strategy about the current state, $q_1 = q_1(v_1)$, and let the planner append a separate pure-forecast message $q_f = q_f(s)$, so that the message is $q = (q_1, q_f)$. Relative to a constant auxiliary message q_f^0, the increase in ex-ante Ramsey loss is*

$$\mathcal{L}(q_f) - \mathcal{L}(q_f^0) = \frac{1}{2} ABC(\beta C)^2 \mathbb{E}\big[\mathbb{E}[s \mid q]^2\big] \geq 0. \tag{29}$$

Revelation is optimal if and only if it leaves the posterior mean of s unchanged, $\mathbb{E}[s \mid q] = 0$ almost

surely. With no current information ($v_1 = 0$ almost surely), a constant message is optimal and the allocation is $e_1^R(q) = g^R(q) = 0$, $\theta_1^R(0, q) = 0$, $\theta_2^R(v_2, q) = Cv_2$.

Proof. See Online Appendix A.2.

Holding the nowcast message fixed, each unit of variance in the posterior mean of the forecast raises loss by $\frac{1}{2}ABC(\beta C)^2$. Revelation adds $\beta C\,\mathbb{E}[s \mid q]$ to the inflationary pressure $M(q)$, while the rule would deliver Cv_2 whether or not the forecast were disclosed. The planner therefore keeps the pure forecast private.

Under delegation, the central bank also keeps the pure forecast private when the audit depends only on the announcement and the delivery residual. To see why, write the delivered center as $Cv_2 + g + a$, where a is the departure from the announced center. The deviation from the ex-post ideal is then $g + a$, and the audit residual is $a + \varepsilon_2$. Neither expression depends on the shock. Consequently, the pure forecast enters the central bank's optimized loss only through an additive constant, and central banks with the same nowcast rank announcements identically. In every equilibrium with unique best responses, the announcement therefore depends on the nowcast alone (Lemma 6 in Online Appendix A.8).

Conditioning the sanction separately on the realized shock could make reporting incentives depend on the forecast, as (24) allows. Within the class of mandates that communicate a nowcast message together with a separate forecast message,[14] however, society gains nothing from doing so. Under the implementation conditions of Lemma 7, an audit depending only on the announcement and the inflation miss attains the Ramsey benchmark for this class, and every equilibrium that moves the posterior mean of the forecast loses at least the cost in Corollary 2 (Online Appendix A.8). The auditor uses the realized shock to construct that miss, but needs no separate shock-dependent sanction. Guidance about current conditions is different: the nowcast carries news ηv_1 about the shock to which the rule responds. Under the ideal rule, revealing it is optimal only for a small news content, $\eta < \eta^\star = 0.026$ at the baseline calibration (Proposition 1).

4.1.2 Delivery and Traction

Implementing the planner's promise requires discipline over delivery. Once the public has formed its expectations, the central bank is tempted to move policy toward its ex-post ideal. The audit limits that adjustment and determines how much of an announcement is delivered. The construction applies to any rule slope after expressing the announced center relative to the ex-post ideal.

[14] A separate forecast message is one generated from s alone, and hence independent of v_1. Messages that jointly garble the two components are outside this comparison.

Let g be the departure announced in period 1 under a rule of slope ϕ, so that the announced center in period 2 is $\phi v_2 + g$. The central bank's action can be measured in two ways: the *delivery shortfall* $a \equiv \theta_2 - (\phi v_2 + g)$, how much it misses its own announced center, and the *delivered deviation* $\Delta \equiv \theta_2 - C v_2$ from its ex-post ideal. They are linked by the *pre-shortfall deviation*, the distance from the ex-post ideal that the rule and the promise would produce if obeyed exactly,

$$\Delta^a \equiv \phi v_2 + g - C v_2, \qquad \text{so that} \qquad \Delta = \Delta^a + a. \tag{30}$$

Under the ideal rule $\Delta^a = g$; under the strict target $\Delta^a = g - C v_2$.

The auditor observes the shortfall only through realized inflation, $\pi_2 = \phi v_2 + g + a + \varepsilon_2$, and forms the residual $z \equiv \pi_2 - [\phi v_2 + g] = a + \varepsilon_2$. The mandate specifies the sanction $\Psi > 0$ and a tolerance $\tau > 0$. Both are chosen in period 0, and τ says how large a residual the auditor lets pass.

We use the *graduated parameterization* of the delivery audit, ρ_C, defined in (A.58) in Online Appendix A.11. An announcement component, ρ_I, is added in Section 4.2.6. Integrating over implementation noise gives the expected delivery sanction $\mathcal{A}_C(a)$ in (A.59). It is smooth, symmetric, increasing in $|a|$, and bounded by Ψ. The magnitude of its marginal sanction peaks at $|a| = \sqrt{\tau^2 + \sigma_\pi^2}$ and declines beyond it. We call the interval of that half-width around zero the *responsive range*. Choosing the tolerance τ trades off the width of this range against the curvature of the expected sanction at zero shortfall.

In period 2 the central bank takes Δ^a as given and chooses the shortfall that balances the macroeconomic cost of staying away from its ideal against the expected sanction,

$$\mathcal{V}_C(\Delta^a) \equiv \min_a \left\{ \frac{A}{2}(\Delta^a + a)^2 + \mathcal{A}_C(a) \right\}. \tag{31}$$

At an interior optimum, the marginal stabilization cost balances the marginal sanction. The peak of $|\mathcal{A}_C'|$ bounds the deviation that the audit can sustain. The curvature at zero shortfall, $\mathcal{A}_C''(0) = \Psi\tau/(\tau^2 + \sigma_\pi^2)^{3/2}$, determines the response to a small announcement. Define the *traction function* as $\Delta(\Delta^a) \equiv \Delta^a + a^*(\Delta^a)$, the deviation actually delivered. Its slope at the origin is the *traction* of guidance, $\omega \equiv \mathcal{A}_C''(0)/(A + \mathcal{A}_C''(0))$ (Online Appendix A.11). Thus the central bank delivers a fraction ω of a small pre-shortfall deviation and undoes the rest.

The public anticipates actual delivery. An announcement can therefore be calibrated to implement a desired deviation within a bounded range, provided the responsive solution is globally optimal. Condition (32) below ensures global strict convexity and is satisfied at

the baseline calibration.[15]

Lemma 3 (Exact delivery over a bounded domain)**.** *Suppose the sanction stake and the audit tolerance satisfy*

$$\Psi\tau < \frac{A}{2}\exp(3/2)\,(\tau^2+\sigma_\pi^2)^{3/2}. \tag{32}$$

Define the delivered-deviation ceiling

$$\bar{\Delta}_C \equiv \frac{\Psi\tau}{A\exp(1/2)\,(\tau^2+\sigma_\pi^2)}. \tag{33}$$

For every target deviation with $|\Delta^\star| \le \bar{\Delta}_C$ *there is a unique shortfall* $a^* \in [-\sqrt{\tau^2+\sigma_\pi^2}, \sqrt{\tau^2+\sigma_\pi^2}]$ *solving*

$$\mathcal{A}_C'(a^*) = -A\Delta^\star, \tag{34}$$

of the opposite sign to $\Delta^\star$ *and zero at* $\Delta^\star = 0$*. The pre-shortfall deviation*

$$\Delta^a(\Delta^\star) \equiv \Delta^\star - a^* \tag{35}$$

has a^* *as its unique delivery optimum, so that* $\Delta = \Delta^\star$ *is delivered; on this domain* Δ^a *is continuous, odd, and strictly increasing.*

Proof. See Online Appendix A.11.

The ceiling rises with the sanction stake and falls with the temptation curvature A. On the responsive range, the marginal sanction $\mathcal{A}_C'$ is strictly increasing, so (34) identifies a unique shortfall for each target within the ceiling. Condition (32) ensures that this shortfall is also the global optimum.[16] If the condition fails, the responsive solution need not be globally optimal: a larger miss may be less costly once the marginal sanction has declined.

To deliver $\Delta^\star$, the announcement must overstate its magnitude by the amount the central bank will subsequently undo. This predictable overstatement is *optimal hyperbole*. Online Appendix Figure A.2 plots the traction function: the horizontal distance to the 45-degree line is the overstatement required to implement a given deviation. Beyond the responsive range, delivery declines with the marginal sanction. Under the ideal rule the pre-shortfall deviation equals the promise. Under other rules it also reflects the response the rule omits, so an unanticipated innovation changes the delivered deviation (Online Appendix A.12

[15]The minimum curvature of the expected sanction is $-2\exp(-3/2)\mathcal{A}_C''(0)$. Adding the stabilization term leaves the delivery objective strictly convex if $A > 2\exp(-3/2)\mathcal{A}_C''(0)$. At the baseline, the two sides are 18.78 and 9.39, respectively.

[16]In the leading example $\Psi\tau = 9.9$ against a bound of 19.8, equivalently $\mathcal{A}_C''(0) = 21.0$ against $(A/2)\exp(3/2) = 42.1$.

derives the state-by-state delivered center and the conditional-mean calibration). The tolerance $\tau = \sigma_\pi$ maximizes the ceiling, and the minimum stake meeting it for every desired deviation up to $\bar{\Delta}$ is $2\exp(1/2)\,A\sigma_\pi\bar{\Delta}$ (Online Appendix A.11); the supplementary appendix establishes a ceiling for any bounded outcome-based sanction. As the stake grows, the shortfall tends to zero and traction tends to one. The *assured-delivery* benchmark sets the shortfall to zero: the delivered center equals the announced one.

Under a coarse rule, the promise can repair part of the omitted response, but it also reveals information about the rule's own future response. Section 4.2 studies the resulting announcement incentives under assured delivery, isolating the reporting problem from the delivery problem. Section 4.2.4 returns to finite traction.

4.2 Coarse Reactive Rules versus the Strict Target

A coarse rule carries only part of the response to tomorrow's shock, $0 < \phi < C$, leaving an omission of $(C - \phi)v_2$. A promise made before the shock arrives can repair its forecastable part.[17] We study the case used in the introductory example: the current state is public and equal to zero, $v_1 = 0$, and the private information is the forecast s. Section 4.2.1 shows when society wants the forecast revealed and what promise it attaches; Section 4.2.2 derives the reporting distortion; Section 4.2.3 solves the announcement game; Section 4.2.4 compares the loss across rules; Section 4.2.6 asks what a mandate that repairs a coarse rule must know.

4.2.1 Revelation under Coarse Rules

Consider the rule-constrained Ramsey planner, who takes a rule of slope ϕ as given and commits to a communication strategy about the forecast as of period 0. Under the rule-plus-promise form (23) the expected period-2 deviation from the ex-post ideal is $(\phi - C)s + g$ conditional on the central bank's forecast s, and the public's expectation is $e_1 = \phi\widehat{s} + g$, the rule's anticipated response to the forecast it has learned plus the promise. Revealing the forecast leaks into today's expectations at the same cost as under the ideal rule, and it lets the promise repair the rule's omission $(C - \phi)s$. The result covers every rule slope at once and reduces to Corollary 2 at the ideal rule.

Proposition 1 (Revelation of the forecast)**.** *Hold the current state public and let q_f be a forecast message with posterior mean $\widehat{s} \equiv \mathbb{E}[s \mid q_f, v_1]$; let q_f^0 convey no forecast information. Under the*

[17] The strict target is $\phi = 0$. The revelation and reporting results below also cover $\phi > C$, where guidance corrects an excessive response. Online Appendix B.2 develops the corresponding announcement game.

rule (22) *with* $\phi \geq 0$*, the change in rule-constrained Ramsey loss is*

$$\mathcal{L}(q_f) - \mathcal{L}(q_f^0) = \left[\frac{1}{2}ABC(\beta C)^2 - \frac{\beta A}{2}(C-\phi)^2\right] \mathrm{Var}(\widehat{s}). \tag{36}$$

(i) Ideal rule. *At* $\phi = C$ *the bracket is the cost of Corollary 2: every message that moves the posterior mean of the forecast raises the loss, and a constant message is optimal.*

(ii) Rules away from the ideal response. *For a message with* $\mathrm{Var}(\widehat{s}) > 0$*, revelation is strictly preferred to opacity if and only if* $|\phi - C| > \bar{c} \equiv C\sqrt{B\beta C}$*. Under full revelation the promise attached to the forecast is*

$$g_\phi^R(s) = \gamma^R s, \qquad \gamma^R \equiv B\big[A(C-\phi) - \beta C\phi\big] = ABC - \phi, \tag{37}$$

the expectation it induces is $ABCs$*, and the conditional expected deviation from the ex-post ideal is* $\mathbb{E}[\theta_2 - Cv_2 \mid s] = (\phi - C)s + g_\phi^R(s) = -B\beta C^2 s$*; neither conditional mean depends on* ϕ*.*

For guidance about current conditions, maintain $0 \leq \phi \leq C$*. With the nowcast* v_1 *private and the forecast absent, the rule-constrained Ramsey loss is affine in* $\mathrm{Var}(\widehat{v}_1(q))$ *with a coefficient proportional to the revelation coefficient* $\Gamma_\phi(\eta)$ *of* (B.18)*: full revelation of the nowcast is strictly preferred to opacity if and only if* $\Gamma_\phi(\eta) < 0$*. Under the ideal rule this confines revelation to a small news content,* $\eta < \eta^\star$ *of* (B.15)*. A coarse rule makes the region strictly wider, because guidance now also supplies the response the rule omits to the predictable part of tomorrow's shock, and under the strict target revelation is optimal at every news content (Lemmas 15 and 16 in Online Appendix B.3).*

Proof. See Online Appendix A.13.

The bracket in (36) compares the cost of disclosure with the value of repairing the rule's omission. The former is the cost in Corollary 2; the latter grows with $(C-\phi)^2$. They balance at $|\phi - C| = \bar{c} = 0.203$. The *revelation region* consists of $\phi < C - \bar{c}$ and $\phi > C + \bar{c}$: guidance can repair either an insufficient or an excessive response. Between these boundaries, opacity is optimal, with indifference at the boundaries. At the baseline, the boundaries are approximately 0.743 and 1.150.

Under truthful guidance, the public expects ABC per unit of forecast under every rule. The planner sets $\gamma^R = ABC - \phi$, so the promise supplies whatever the rule omits from that response. Both the allocation on the predictable component and the expected deviation from the ex-post ideal, $-B\beta C^2 s$ or -0.044 per unit of forecast, are therefore independent of the rule.

What varies is the size of the promise: $ABC = 0.903$ per unit of forecast under the strict target, compared with $\gamma^R = 0.203$ under the rule in the introduction. The next two

subsections show why this matters. Separating from higher forecasts may require giving up part of the promise. We call the amount available for this purpose its *signaling capacity*.

Leading example. Suppose there is no cost-push shock today ($t = 1$), but the central bank privately observes a forecast, denoted by s, of tomorrow's ($t = 2$) shock v_2. The forecast need not be perfect: the realized shock can differ from it, with forecast errors having mean zero conditional on the central bank's information. Forecasts are symmetric around zero and, for this example, range from -3 to 3 percentage points. A period is a year. Without guidance, period 1 has zero inflation and a zero output gap, and the standing rule determines tomorrow's policy. At the baseline parameters, $\beta = 0.96$, $\kappa = 0.15$, and $\lambda = 0.4$, the ideal rule accommodates approximately 0.95 of the realized shock. Keeping the forecast private then achieves the commitment allocation. A coarse rule of 0.70 supplies less accommodation, while a strict target supplies none without guidance.

Once the forecast is revealed, it affects today's inflation expectations. The planner then chooses slightly less accommodation tomorrow: approximately 0.90 per unit of forecast. A promise can implement this conditional expected center. Revelation is desirable when the gain from correcting the rule exceeds its cost through current expectations (Proposition 1).

Let delivery be assured and the rule slope be 0.70. Its shortfall relative to the ideal response is approximately 0.25, above the revelation threshold 0.203. Society therefore wants the forecast communicated. For a forecast of one, the rule supplies 0.70 and the planner adds a promise of approximately 0.20:

$$\gamma^R = ABC - \phi, \tag{38}$$

where ABC is the rule-constrained Ramsey accommodation per unit of forecast under full revelation. The example holds the inflation index fixed; Section 4.2.5 discusses how index choice affects the comparison across rules.

4.2.2 Reporting Distortion

Now consider which promise the delegated central bank chooses when the obedience audit of Section 4.1.2 assures delivery but imposes no separate reporting incentive. Under a coarse rule, delivering the promise g leaves the policy center short of the ex-post ideal by $(C-\phi)v_2 - g$. The central bank's forecast therefore affects the expected cost of each promise: a higher forecast makes additional accommodation more valuable. Unlike under the ideal rule, central banks with different forecasts need not rank announcements identically.

The announcement also affects current expectations. If the public infers a forecast $\widehat{s}$

from the promise g, it expects inflation of $\phi\widehat{s} + g$: the anticipated response of the rule plus the promised departure. A smaller promise interpreted as a lower forecast therefore reduces expected inflation through both components. Yet the central bank gives up only the promised accommodation, because the rule will respond to the shock that actually arrives, regardless of the announcement. This difference creates the reporting distortion.

Take the candidate policy in which the central bank announces the planner's promise, $g^R_\phi(s) = \gamma^R s$, and the public reads it at face value. Write r for the forecast the public infers from an observed promise, $r = g/\gamma^R$, a relabeling of the announcement that we call the *report*. The expectation attached to r is $\phi r + \gamma^R r = ABCr$, and a central bank with forecast s that announces the promise indexed by r incurs

$$J_\phi(s, r) = \frac{C}{2}\left[\beta ABC\, r\right]^2 + \frac{\beta A}{2}\left[(\phi - C)s + \gamma^R r\right]^2 \tag{39}$$

up to terms independent of r: the current stabilization loss after expectations have moved, plus the discounted cost of the delivered deviation from the ex-post ideal, which contains the rule's unrepaired response to the true forecast.

Proposition 2 (Reporting distortion)**.** *Let the current state be public and equal to zero, $\phi > 0$, $|\phi - C| > \bar{c}$, and delivery assured. At the truthful report, the reduction in the central bank's loss per unit reduction in the reported forecast is*

$$\frac{\partial J_\phi}{\partial r}(s, s) = \phi\, AB(\beta C)^2\, s, \tag{40}$$

positive for a positive forecast. Let

$$\mathcal{A}^\phi_R(r) \equiv -\frac{1}{2}\phi\, AB(\beta C)^2\, r^2. \tag{41}$$

If the expected audit burden associated with the report r varies according to $\mathcal{A}^\phi_R$, then $r = s$ is the unique global minimizer of $J_\phi(s, r) + \mathcal{A}^\phi_R(r)$ for every forecast. Under the strict target, $\phi = 0$, the distortion is zero for every forecast and the central bank announces the planner's promise with no correction.

For guidance about current conditions, maintain $0 \le \phi \le C$, with the nowcast v_1 private, the forecast absent, and $\Gamma_\phi(\eta) < 0$, the distortion at the truthful report is $ABC\beta\eta\phi\, c_\eta\, v_1$ with $c_\eta \equiv 1 + \beta C\eta$: the rule's loading on the news the report carries, $\phi\eta$, times the value of a unit reduction in expected inflation. For $\eta\phi > 0$ it has the sign of the nowcast; it vanishes when $\eta\phi = 0$. The correction $-\frac{1}{2}ABC\beta\eta\phi c_\eta r^2$ makes the truthful report the unique global minimizer throughout the nowcast revelation region (Lemma 17 in Online Appendix B.3).

Proof. See Online Appendix A.14.

At the planner's promise, the marginal benefit of changing the promise itself exactly offsets its delivery cost. The remaining incentive comes from the rule's anticipated response to the information revealed. The *reporting distortion* is the product of the rule's loading on the forecast, ϕ, and the value of a unit reduction in expected inflation, $AB(\beta C)^2 s$. For a positive forecast it favors a smaller report. Below $C - \bar{c}$, this means less promised accommodation; above $C + \bar{c}$, it means a less negative promise and hence less restraint. The *announcement correction*, $\mathcal{A}_R^{\phi}$, removes this incentive in both regions.

The next subsection studies what happens without the correction: the central bank gives up accommodation to sustain informative announcements, incurring a separation cost. Section 4.2.6 then shows how society can design the correction and announcement format to restore the planner's outcome. Both depend on the model. Under the strict target, the reporting distortion is absent and neither is needed. Under the ideal rule, the forecast does not affect the ranking of announcements and society prefers opacity. For $\phi < C$ close to the ideal rule, implementing opacity requires a public formula restricting announcements; a smooth announcement-dependent burden cannot induce the same announcement from every forecast type (supplementary appendix).

Leading example (continued). Such a coarse rule gives the central bank an incentive to promise less than the Ramsey planner would. To understand why, suppose initially that the public reads announcements as if the central bank were following the planner's schedule. A smaller promise is interpreted as a smaller forecast, so the public expects less accommodation from both the promise and the rule. For every unit by which the central bank understates its forecast, expected inflation falls by $0.70 + 0.20 = 0.90$: by 0.70 through the anticipated response of the rule and by 0.20 through the smaller promise. But tomorrow the rule responds to the shock that actually arrives, whatever was announced. The central bank obtains the full reduction in expected inflation today while giving up only 0.20 of accommodation tomorrow. It therefore pays for only the fraction $\gamma^R/ABC = (ABC - \phi)/ABC$, or $0.20/0.90 \approx 0.22$, of the expectation it moves. The remaining 78 percent is supplied by the rule and costs the central bank nothing. This is what makes understatement attractive. The marginal gain from being read as a smaller forecast is the rule's loading on the revealed forecast times the value of a unit reduction in expected inflation, $\phi\, AB(\beta C)^2 s$, which is zero at $\phi = 0$ and grows with the rule's reactivity (Proposition 2). Under a strict target, by contrast, the rule contributes nothing, so the same understatement lowers expected inflation by 0.90 and forgoes the full 0.90 of accommodation: the central bank's tradeoff is society's. Panel A of Figure 1 sets the two regimes side by side.

A. Gain and cost of understating

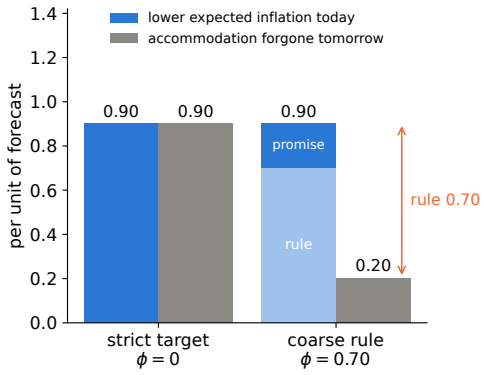

B. Announcements by forecast

		Separating schedule	
Forecast s	Optimal	3 forecasts	6 forecasts
3.0	0.61	0.61	0.61
2.5	0.51	—	0.34
2.0	0.41	0.25	0.21
continuum limit: collapse below $s_b = 1.72$			
1.5	0.30	—	0.12
1.0	0.20	0.09	0.05
0.5	0.10	—	0.00

Figure 1: Leading example: Reporting distortion and the separating schedule

Notes: $\phi = 0.70$, assured delivery. Panel A: reduction in expected inflation and accommodation forgone per unit of forecast understatement. Panel B: Ramsey promises and the largest separating promises that deter imitation from the next higher forecast, on grids of three and six types. The marked $s_b = 1.72$ boundary applies only to the continuum limit. Both finite grids separate; the lowest six-type promise is positive but rounds to 0.00.

4.2.3 Guidance without a Correction

We now study the announcement game when the public anticipates the reporting distortion. The mandate enforces delivery but provides no separate reporting incentive. The question is which announcements can convey the forecast and how much of its value they preserve. This subsection focuses on $\phi < C - \bar{c}$; Online Appendix B.2 constructs an equilibrium for rules above C.

Assumption 1 (Announcement game without a correction)**.** *(a) Delivery is assured: the delivered center equals the announced center* $\phi v_2 + g$*. (b) The mandate carries no announcement component, and sanctions are socially free, so the social loss is the planner's loss evaluated at the equilibrium allocation. (c) The forecast has a bounded, symmetric support that is common knowledge,* $s \in [-S, S]$*; the calibration sets* $S = 3\sigma_s$*. (d) Announcements are unrestricted,* $g \in \mathbb{R}$*, unless an announcement format is specified. Throughout, the current state is public and equal to zero.*

To interpret an informative announcement, the public must understand how the central bank's forecast maps into its promise. Consider a strictly monotone candidate schedule $G(s)$. If the public believes that the central bank follows this schedule, it interprets an announcement g in its range as the forecast $G^{-1}(g)$. We call this interpretation a *reading*. Under assured delivery, it implies

$$e_1(g) = \phi G^{-1}(g) + g.$$

A reading is *self-confirming* when each type optimally chooses the promise assigned by G given that interpretation. A separating equilibrium additionally requires beliefs outside the schedule's range that satisfy the equilibrium conditions in Section 3.3. The next two results characterize, first, self-confirming readings of a single slope, which need no bound on the forecast, and then, using the bounded support, equilibrium separation and the uninformative announcement.

Readings of a single slope. Consider first a reading of a single slope, $G(s) = \gamma s$ with $\gamma \neq 0$: the public decodes an announcement g as the forecast g/γ and expects $e_1(g) = \phi g/\gamma + g$, one unit of promise moving expected inflation by one unit directly and by ϕ/γ units through the rule's anticipated response to the forecast it reveals. The central bank of type s chooses g to minimize

$$J(s, g) = \frac{C}{2}\Big[\beta\big(1 + \tfrac{\phi}{\gamma}\big)g\Big]^2 + \frac{\beta A}{2}\big[(\phi - C)s + g\big]^2, \tag{42}$$

the current stabilization loss after expectations have moved plus the discounted cost of the delivered deviation from the ex-post ideal. The minimizer is $g = b(\gamma)s$ with

$$b(\gamma) = \frac{A(C - \phi)}{A + \beta C(1 + \phi/\gamma)^2}, \tag{43}$$

and the reading is self-confirming when $b(\gamma) = \gamma$. The comparison point is the rule-constrained Ramsey promise of Proposition 1, of slope $\gamma^R = B[A(C - \phi) - \beta C\phi]$.

Lemma 4 (Self-confirming readings of a single slope)**.** *Under Assumption 1(a), (b), and (d), let $0 < \phi < C - \bar{c}$ and let a reading of slope γ be applied to every announcement. The reading is self-confirming if and only if γ solves*

$$(A + \beta C)\gamma^2 - \big[A(C - \phi) - 2\beta C\phi\big]\gamma + \beta C\phi^2 = 0. \tag{44}$$

Such readings exist if and only if $\phi \leq \hat{\phi}$, where

$$\hat{\phi} \equiv \frac{AC}{\big[\sqrt{A + \beta C} + \sqrt{\beta C}\big]^2}; \tag{45}$$

Below the threshold there are two distinct roots, $0 < \gamma_1 < \gamma_2 < \gamma^R$; at the threshold they coincide. We call γ_2 the near-Ramsey reading*. The social loss under a self-confirming reading of slope γ exceeds the rule-constrained Ramsey loss by $\frac{1}{2}(\beta A + \beta^2 C)(\gamma - \gamma^R)^2$ per unit of variance of the forecast.*

For guidance about current conditions, with the nowcast private and $\Gamma_\phi(\eta) < 0$, readings of a single slope are self-confirming at the roots of the nowcast counterpart of (44), and for $\eta > 0$, such readings exist if and only if the rule slope lies at or below the threshold $\hat{\phi}(\eta)$ in (B.25), which is V-shaped in the news content: it equals C at $\eta = \eta^\star/2$, falls to zero at $\eta = 1/A$, where the planner's promise under the strict target vanishes, and rises thereafter toward the pure-forecast threshold $\hat{\phi}$ (Lemma 18 in Online Appendix B.3).

Proof. See Online Appendix A.15.

Both self-confirming slopes lie below the planner's slope. The loss is quadratic in this departure, with no first-order term. The *collapse band*, $(\hat{\phi}, C - \bar{c})$, consists of rules under which society wants the forecast communicated but no reading of a single slope is self-confirming.

The fixed-point equation has two roots below $\hat{\phi}$, a repeated root at the threshold, and none above it. Its left side can be written as $(A+\beta C)\gamma(\gamma-\gamma^R)+\beta C\phi(\phi+\gamma)$. At the planner's slope the first term is zero and the second is positive. This additional expectations cost pushes both self-confirming slopes below γ^R.

The threshold is where the two quantities of the introduction's example cross. The signaling capacity, what the rule leaves for the promise to do, is $A(C-\phi)$ per unit of forecast in value terms, the marginal value of repairing the rule's omission, and the planner's promise $\gamma^R s$ in units of policy. The *separation cost*, what a forecast must give up to be believed, is the minimum marginal cost of promising along a reading of a single slope, the delivery cost of the promise plus the expectations cost of the forecast it reveals; per unit of forecast it is $\phi[2\beta C + 2\sqrt{\beta C(A+\beta C)}]$. The cost grows with ϕ and the capacity shrinks with it, and $\hat{\phi}$ is the rule at which they are equal.

For less reactive rules, the central bank can announce nearly the planner's promise. At $\phi = 0.5$, the near-Ramsey slope is 0.347 against $\gamma^R = 0.403$, and the loss is 2.1% of the value of the information. As the rule becomes more reactive, the promise shrinks and the incentive to understate the forecast grows. Above $\hat{\phi} = 0.612$, a positive forecast prefers a smaller promise under every candidate reading of a single slope.

Figure 2 shows how the two readings merge at the threshold. There the best-response curve is tangent to the 45-degree line, and its slope reaches one. This slope also governs the sensitivity of the chosen promise to a misjudged reading (Remark 6 in Online Appendix A.15).

Bounded forecasts. Whether a reading of a single slope is self-confirming depends on the rule. When none is, the public still needs a rule for decoding announcements, and the bound on the forecast, part (c) of Assumption 1, supplies it. The public cannot believe a

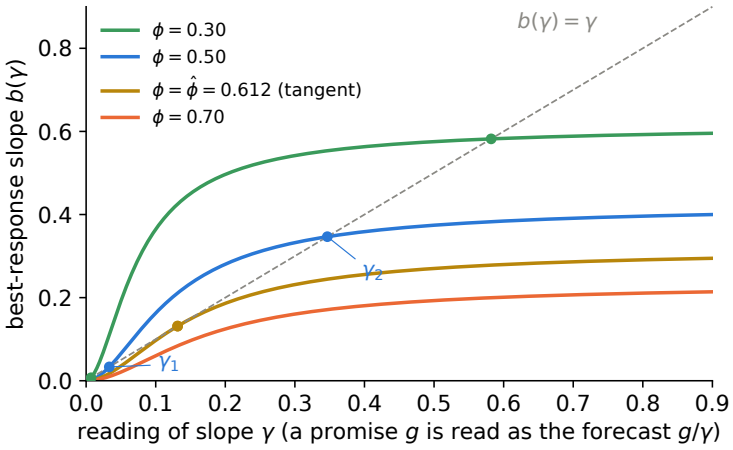

Figure 2: Self-confirming readings as fixed points of the best response

Notes: Pure forecast, assured delivery, obedience audit only; baseline calibration. Curves show the best-response slope $b(\gamma)$ in (43); intersections with the 45-degree line are self-confirming slopes (dots). At $\phi = \hat{\phi} = 0.612$ the roots coincide; at $\phi = 0.70$ none exists. At the near-Ramsey root, b' equals 0.07, 0.26, and 0.71 for $\phi = 0.3$, 0.5, and 0.6, respectively, and reaches one at tangency (Remark 6).

forecast larger than S, so the central bank with that forecast is read at worst as itself. When its forecast is revealed, this pins its announcement; the incentive conditions then determine the separating schedule below it, and where separation fails the schedule must end in an uninformative announcement shared by a set of forecasts. To characterize that outcome, we restrict attention to symmetric monotone equilibria satisfying criterion D1 (Cho and Kreps, 1987; Banks and Sobel, 1987). Monotonicity requires $G(s)$ to be nondecreasing. Symmetry requires opposite forecasts to induce opposite announcements and opposite announcements to induce opposite posterior means: $G(-s) = -G(s)$ and $\widehat{s}(-g) = -\widehat{s}(g)$, where $\widehat{s}(g) \equiv \int s\, \mu(\mathrm{d}s \mid g)$. D1 disciplines how the public interprets an announcement that no type makes in equilibrium. It excludes a type if every public response making that announcement weakly profitable for that type makes it strictly profitable for another. The formal implications of these restrictions are derived in Online Appendix A.15.

Proposition 3 (The collapse of informative guidance)**.** *Under Assumption 1, let* $0 < \phi < C - \bar{c}$.

(i) Boundary condition. *In every equilibrium in which the largest forecast is read as itself, it announces its complete-information promise,* $G(S) = \gamma^R S$.

(ii) Below the threshold. *If* $\phi \le \hat{\phi}$, *there is exactly one separating equilibrium announcement strategy. It is odd and increasing, lies between the near-Ramsey reading and the rule-constrained Ramsey promise,* $\gamma_2 s \le G(s) \le \gamma^R s$ *for* $s \ge 0$, *and converges to the near-Ramsey*

reading as $S \to \infty$. No symmetric monotone equilibrium satisfying criterion D1 pools forecasts around zero.

(iii) In the collapse band. *If $\hat{\phi} < \phi < C - \bar{c}$, no separating equilibrium exists. Exactly one symmetric monotone equilibrium outcome satisfies criterion D1: forecasts with $|s| \leq s_b$ issue the uninformative announcement $g = 0$, and forecasts beyond $s_b \in (0, S)$ announce along the separating schedule that starts from the boundary condition of part (i). The ratio s_b/S depends on ϕ alone.*

For guidance about current conditions, with the nowcast private and bounded and $\Gamma_\phi(\eta) < 0$, the boundary condition of part (i) and the differential equation behind parts (ii)–(iii) carry over, with the threshold $\hat{\phi}(\eta)$ of Lemma 4 in place of $\hat{\phi}$: below it the separating schedule has the structure of part (ii), its global incentive compatibility verified numerically (the nowcast part of the proof in Online Appendix A.15).

Proof. See Online Appendix A.15.

Two properties determine the separating schedule. First, the public's belief enters every type's loss through the same expectations term. The largest forecast therefore cannot be mistaken for a still larger one, which pins its announcement at the complete-information promise. Second, a higher forecast values additional promised accommodation more. This single-crossing property turns the incentive conditions into a differential equation with a unique solution through the top announcement.[18]

Below the threshold, the schedule reaches the origin and separates all forecasts. Above it, continuing the separating schedule downward would produce a zero promise at a positive forecast. It therefore cannot describe a fully separating equilibrium. D1 determines the remaining pooling outcome: it rules out pools with a nonzero mean, leaving a symmetric group that announces zero. The pooling boundary s_b is determined by indifference between the informative promise and this uninformative announcement. It lies strictly above the forecast at which the continued separating schedule would reach zero. Thus the failure of full separation follows from the incentive conditions; the refinement identifies the central pool and its boundary.

The economic cost of separation is the accommodation forgone to deter imitation by a higher forecast.[19] Under a more reactive rule, a smaller promise reduces expected inflation substantially while giving up relatively little accommodation. Deterring imitation then requires a larger departure from the planner's promise. In the collapse band, the remaining

[18] In (42), $\partial^2 J/\partial s\, \partial g = \beta A(\phi - C) < 0$. Only the delivery term depends on the central bank's own forecast.

[19] This is the costly-signaling logic of Spence (1973); Mailath (1987); the costly action here is a smaller promise, delivered in full.

promise is too small to justify informative guidance for ordinary forecasts. They announce zero, leaving the rule's omitted response unrepaired. At $\phi = 0.70$, $s_b/S = 0.57$. With $S = 3\sigma_s$, the pool covers 92% of forecasts and guidance recovers only 14% of the value of the information; at $\phi = 0.65$, it recovers 66%.

Figure 3 shows the schedules. At $\phi = 0.5$ the promise stays within the narrow band between the near-Ramsey reading and the planner's promise for every forecast. At $\phi = 0.70$ the promise falls from 0.609 at the largest forecast to 0.169 at a forecast of 2 and to 0.115 at $s_b = 1.72$, where the schedule stops: most of the signaling capacity is spent separating the largest forecasts from one another. Without a correction, the bound and the equilibrium restrictions select this outcome; the correction of Proposition 2 instead changes the reporting incentives, and under assured delivery the strict target needs neither (Proposition 4). Figure 4 traces the loss across rules; Supplementary Appendix 2.1 gives its dependence on the support.

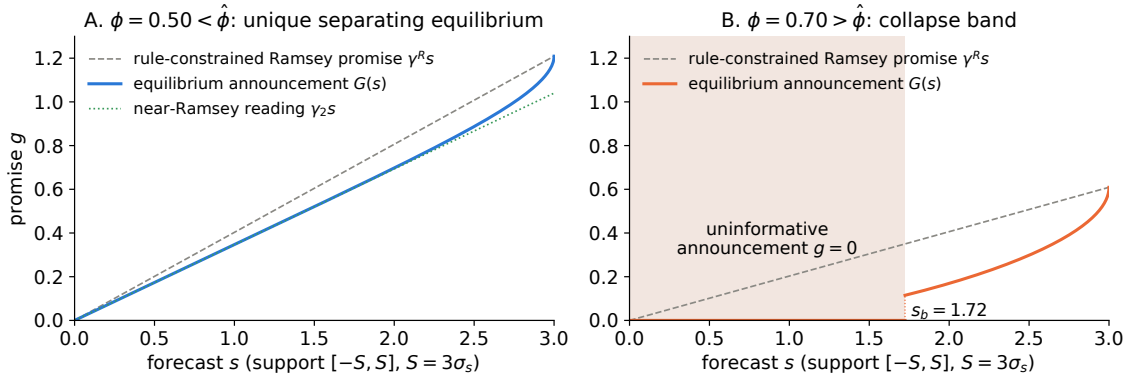

Figure 3: Equilibrium announcement schedules with bounded forecasts

Notes: Pure forecast, assured delivery, obedience audit only; baseline calibration, $S = 3\sigma_s$. Positive forecasts are shown; the schedules are odd. Solid: equilibrium $G(s)$. Dashed: Ramsey promise $\gamma^R s$. Panel A, $\phi = 0.5$: full separation between the near-Ramsey line (dotted, $\gamma_2 = 0.347$) and the Ramsey line ($\gamma^R = 0.403$). Panel B, $\phi = 0.70$: the informative schedule falls from 0.609 to 0.115 at $s_b = 1.72$; forecasts in the shaded interval announce zero (Proposition 3).

The strict target. Under the strict target the public expects from the rule nothing that depends on the forecast, so an announcement moves expectations by exactly the promise, whatever the public infers.

Proposition 4 (Strict targeting: a promise means what it says)**.** *Under Assumption 1(a)–(b) and the strict target $\phi = 0$, the rule-constrained Ramsey promise, $g = ABC\,s$, is each type's best announcement under every belief the public could hold, and hence under every reading. This promise is therefore the unique equilibrium announcement strategy, for any support of the private information and without any announcement correction or format, and the truthful reading is the only self-confirming reading. In the revelation region the strict target is the only rule with this property: for every $0 < \phi < C - \bar{c}$ the central bank's best promise depends on the reading.*

For guidance about current conditions the same holds for every news content $\eta \neq 1/A$: under the strict target the best response to every belief is the planner's promise $g_0^R(v_1) = BC(A\eta - 1)v_1$, so the truthful reading is the only self-confirming reading and the central bank chooses the planner's promise with no reporting correction. At $\eta = 1/A$ that promise is zero for every nowcast, and revelation and opacity yield the same loss (Lemma 18 in Online Appendix B.3).

Proof. See Online Appendix A.15.

Under the strict target, $e_1(g) = g$ regardless of the public's belief. The reading therefore drops out of the central bank's loss, and every type has a unique optimal promise. The entire response is carried by the promise, so separation requires no sacrifice of accommodation. With delivery assured, truthful guidance needs neither a forecast bound nor an announcement restriction or reporting correction calibrated to the model.[20]

Leading example (continued). The public anticipates the incentive to understate the forecast under coarse rules. To be believed, a forecast must be associated with a promise that a central bank holding a slightly higher forecast would not prefer to send. What deters the central bank with the higher forecast is the accommodation tomorrow, at $t = 2$, that it would forgo by delivering the smaller promise. The issue is that society's desired schedule ex ante—the Ramsey schedule—does not provide enough deterrence: a unit of understatement lowers expected inflation by 0.90 while forgoing only 0.20 of accommodation.

The separating schedule is pinned at the top: nobody can be read as having a forecast above 3, so the central bank with the largest forecast announces its full promise, $3 \times 0.20 \approx 0.61$. Below the top, each forecast must promise enough less to deter imitation by the forecast just above it.[21] Panel B of Figure 1 illustrates this requirement on two grids: the promise per unit of forecast falls at every step, from 0.20 at the top to 0.09 at a forecast of

[20] Under a coarse rule, the cost of misjudging the public's reading is second order in the error. Its coefficient depends on the slope of the best response, which rises from 0.07 at $\phi = 0.3$ to one at $\hat{\phi}$. Under the strict target that slope is zero (Remark 6 in Online Appendix A.15).

[21] Only the constraint from the adjacent forecast binds. Being read as a much smaller forecast requires announcing a much smaller promise and then delivering it, and the accommodation forgone grows with the distance, while the reduction in expected inflation it buys is bounded. At this calibration the largest forecast is deterred from imitating any forecast below 0.73: with forecasts 3 and 0 alone, both announce the Ramsey

one with three forecasts, and to 0.05 with six. Both finite grids separate; the continuum calculation below determines where separation gives way to pooling.

With a continuum of forecasts, the incentive constraint binds at every step and the reductions accumulate. Under truthful guidance, expected inflation rises by 0.90 per unit of forecast under every rule in the revelation region. The coarse rule changes the promise that can be sustained in equilibrium. That promise falls below the planner's 0.20 per unit of forecast, reducing the accommodation that makes revelation worthwhile. At a forecast of 1.72, the promise is 0.11, and the central bank is indifferent between that informative announcement and the uninformative one. Smaller positive forecasts announce no departure; larger forecasts remain informative (Proposition 3).

Take a forecast of $s = 1$. The central bank announces no departure, and because the forecasts sharing that announcement are symmetric around zero, the public expects a forecast of zero. Conditional on the central bank's own forecast, the rule provides expected accommodation of 0.70 tomorrow, but the 0.20 the planner would have promised is never supplied. Forecasts within 1.72 of zero are about 92% of the distribution, so fewer than one forecast in ten is announced informatively. Promises are kept; the 0.90 center is not reached.

The failure is therefore in communication, even though promises are delivered.

4.2.4 Welfare across Standing Rules

We compare three losses per unit of forecast variance: truthful guidance, equilibrium guidance under delivery enforcement alone, and no guidance. The strict target attains the first without a reporting correction; a coarse rule attains it with the correction and format of Section 4.2.6. The second is the equilibrium of Proposition 3. Under the third, the rule responds to the realized shock while the forecast stays private. Parts (i) and (ii) compare outcomes on the predictable component; part (iii) adds the unforecastable innovation.

Proposition 5 (Welfare across standing rules)**.** *Under Assumption 1, let* $L^T \equiv \frac{1}{2}ABC(\beta C)^2$. *Per unit of forecast variance, with the innovation set to zero in parts (i) and (ii):*

(i) Truthful guidance. *For every* $0 \le \phi < C - \bar{c}$ *the loss with truthful guidance is* L^T*, the same under every rule in the revelation region and the lowest attainable there. The strict target attains it with no correction; without a correction, the equilibrium of Proposition 3 has loss strictly above* L^T *for every* $0 < \phi < C - \bar{c}$.

promise and there is no distortion, and with a single forecast there is nothing to conceal. Adding forecasts therefore adds binding links rather than a new force, and each link rests on a promise the link above it has already lowered.

(ii) No guidance. *The loss with no guidance is $\frac{1}{2}\beta A(C-\phi)^2$; it equals L^T at $\phi = C - \bar{c}$, lies below L^T for $\phi \in (C-\bar{c}, C]$, and is zero at the ideal rule.*

(iii) The innovation. *Let $\sigma^2_{2|1} \equiv \mathrm{Var}(\varepsilon)$ be the variance of the part of the shock no forecast covers. The innovation adds $\frac{1}{2}\beta A(C-\phi)^2\sigma^2_{2|1}$ to the loss under every rule with assured delivery. Among $0 \le \phi \le C$, it is largest under the strict target and zero under the ideal rule.*

For guidance about current conditions the same comparison holds on the predictable part ηv_1 of the shock. Among the rules under which society wants the nowcast revealed, the loss with truthful guidance is again common to all of them, $\frac{1}{2}ABC\, c^2_\eta$ per unit of nowcast variance; the loss with no guidance, $\frac{C}{2} + \frac{1}{2}\beta A(C-\phi)^2\eta^2$, again falls as the rule becomes more responsive; and the innovation again adds the term of part (iii).

Proof. See Online Appendix A.16.

On the predictable component of the shock, truthful guidance produces the same allocation under every rule in the revelation region. The identity $\phi + \gamma^R = ABC$ in Proposition 1 explains why: the promise supplies what the rule omits, so expected accommodation is ABC per unit of forecast. The resulting loss is $L^T = 0.373$ at the baseline calibration. Under assured delivery, the strict target attains this baseline without a reporting correction; a corrected coarse rule attains it as well. With delivery enforcement alone, however, coarse rules pay a separation cost. It is small when guidance stays close to the planner's schedule, 1.1% of L^T at $\phi = 0.3$ and 7.5% at $\phi = 0.5$, and rises steeply in the collapse band, reaching 38% at $\phi = 0.65$ and 40% at $\phi = 0.70$ (Figure 4).

The comparison with no guidance explains why this problem is concentrated among intermediate rules. At $\phi = C - \bar{c}$, no guidance and truthful guidance have the same loss. Above that threshold, withholding the forecast is preferable: the rule supplies enough of the desired response that revelation costs more through current expectations than guidance gains through future policy. The no-guidance loss falls to zero at the ideal rule under the proposition's normalization, although the ideal rule and rules near it require the designer to know the optimal response C. In this pure-forecast case, the current shock is zero and the episode carries no inherited promise. With the forecast kept private, discretion therefore delivers the same allocation as the ideal rule. The gains from commitment arise when a promise improves current stabilization but would not be honored later; the infinite-horizon extension in Online Appendix D also incorporates promises inherited from earlier periods.

Less reactive rules can perform well because guidance largely completes their response. Rules close to the ideal response can perform well without guidance. Between them, a rule can leave a valuable role for guidance while preventing ordinary forecasts from being communicated. With delivery assured, the strict target avoids this reporting problem

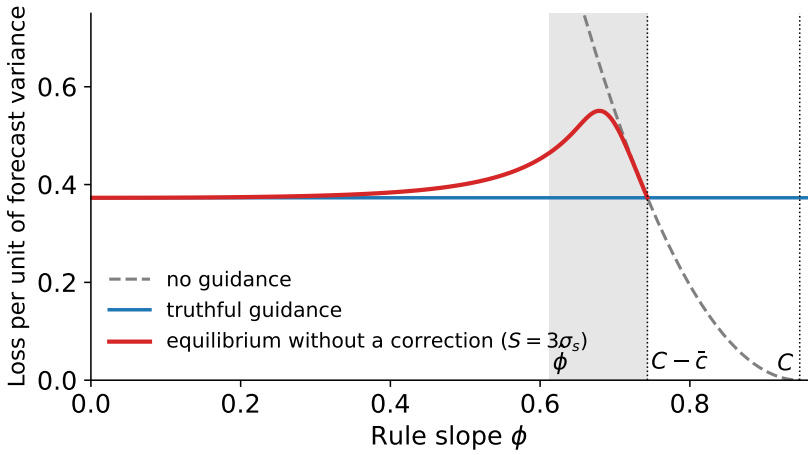

Figure 4: Loss across standing rules on the predictable part of the shock

Notes: Loss per unit of forecast variance; pure forecast, assured delivery, socially free sanctions, no innovation; baseline calibration and $S = 3\sigma_s$. Horizontal: truthful guidance, $L^T = 0.373$, attained by the strict target or a corrected coarse rule. Red: equilibrium without a reporting correction; its excess over L^T is the separation cost. Dashed: no guidance, $\frac{1}{2}\beta A(C - \phi)^2$. Shaded: collapse band $(\hat{\phi}, C - \bar{c})$. Beyond $C - \bar{c}$, society prefers opacity.

without a model-calibrated correction. For the fixed index studied here, it incurs a larger loss on unforecastable innovations than any rule with $0 < \phi \leq C$. The following discussion separates this stabilization cost from the choice of inflation index.

Noisy delivery. The collapse arises even when delivery is assured. To examine what changes with a finite sanction stake, consider the graduated parameterization. Under this audit the rule-constrained planner attains L^T when the target lies below the delivery ceiling and the responsive delivery choice is globally optimal, by announcing more and delivering the same (the optimal hyperbole of Section 4.1.2; Lemma 12 in Online Appendix B.1). What traction changes is the alternative to guidance: with a finite stake the central bank repairs part of the rule's omission at delivery through its accountable discretion, so the loss without guidance is $\frac{1}{2}\beta A\,\Delta((\phi - C)s)^2$ in place of $\frac{1}{2}\beta A(C - \phi)^2 s^2$, and the value of guidance under a coarse rule shrinks as the stake falls. The collapse threshold falls as well, to $\hat{\phi}(\omega) = 0.46$ at the baseline stake $\Psi = 18$ (Online Appendix C), because only a fraction ω of a promise is delivered while the public also anticipates discretionary repair of the rule's omission (Online Appendix B.1 gives the closed form). Below the traction $\omega^\star = 0.16$ no admissible rule has a positive separating slope at the origin. This is a local threshold; bounded-support outcomes also depend on global incentives.

Unforecastable shocks and monitoring. Part (iii) of Proposition 5 is the stabilization advantage of reactive rules: on the part of the shock no forecast covers, a rule of slope ϕ delivers ϕ of the response automatically, and under assured delivery a promise fixed before the shock arrives can add nothing to it. With a finite stake the central bank accommodates part of the innovation through its accountable discretion—locally, the fraction $1-\omega$ of the rule's omission—and the price then moves to the monitoring side.

The standing rule is also the baseline against which performance is judged. For this comparison let the private information be the nowcast alone, the case of Supplementary Appendix 2.4 (with the forecast the argument is the same, with ε as the innovation), reported truthfully as $r = v_1$; let the central bank add a hidden departure b from honest policy, and let the innovation $v_2 - \eta v_1$ be Gaussian with variance $\sigma^2_{2|1} > 0$, independent of v_1, b, and implementation noise. If the central bank fully repairs the rule's response to the innovation, the audit residual is

$$z_\phi = b + (C-\phi)(v_2 - \eta r) + \varepsilon_2, \qquad \operatorname{Var}(z_\phi \mid b, r) = \sigma^2_\pi + (C-\phi)^2\sigma^2_{2|1}. \tag{46}$$

The residual combines legitimate stabilization with nonperformance. The ideal rule removes the innovation term; among $0 \le \phi \le C$, the strict target leaves the most variation. Its residual is the ideal rule's plus independent noise, a Blackwell garbling. The total-variation distance between honest and opportunistic outcomes is strictly smaller, reducing the maximum expected sanction difference for a given hidden deviation (Lemma S1 of the supplementary appendix, in Supplementary Appendix 2.4).[22] Thus a strict target can accommodate innovations through accountable discretion, but this accommodation makes delivery harder to monitor under the prescribed audit.

For a fixed index and a coarse rule in the revelation region, the comparison depends on innovation variance relative to forecast variance. With assured delivery, a sufficiently small innovation variance makes the separation cost decisive in favor of the strict target; a larger variance can reverse that ranking. Choosing the inflation index can change these variances and the stabilization tradeoff itself. The next example distinguishes this choice from uncertainty about the rule's response to a given shock process.

[22]The comparison holds the audit statistic fixed. Replacing the standing rule in that statistic with a state-contingent benchmark recovers the ideal rule's residual. In the local linear approximation, accountable discretion supplies the fraction $1-\omega$ of the omitted response, adding variance $(1-\omega)^2(C-\phi)^2\sigma^2_{2|1}$; Supplementary Appendix 2.4 plots the residual's standard deviation.

4.2.5 Leading Example: A Rule of Uncertain Quality

The preceding comparisons evaluate a known rule slope. We now allow the rule's response to fall short of or exceed the ideal response across episodes. Let its *relative response*, $\rho \equiv \phi/C$, follow

$$\rho = \exp(-\sigma^2/2 + \sigma Z), \qquad Z \sim N(0,1), \tag{47}$$

with $\sigma = 0.60$. The mean relative response is one, the median is 0.835, and the central 95 percent of draws lie between 0.26 and 2.71 times the ideal response. Centering the coefficient on C gives the reactive rule a favorable calibration: uncertainty about the ideal response need not produce a rule centered on it. This assumption concerns the mean coefficient, not the optimality of average policy.

The central bank and the public observe the draw before the announcement, so each episode is the announcement game at the realized slope. We hold C, the inflation index, and the shock process fixed and vary ϕ. The exercise therefore measures the consequences of uncertain relative response for a given index. Losses concern the predictable component, with assured delivery and no reporting correction. Below the ideal response we use Proposition 3; above it we use the equilibrium constructed in Lemma 14.

Figure 5 shows three regions. Within $\bar{c}$ of the ideal response, society prefers opacity. These draws account for 27 percent of episodes and have an average loss of $0.33L^T$. In the two collapse bands, society wants revelation, but ordinary forecasts share an uninformative announcement. Average loss is $1.52L^T$ and the maximum is $2.07L^T$. These bands contain 24 percent of draws and account for 66 percent of the expected loss that a reporting correction would remove. Outside these regions, forecasts separate and average loss is $1.13L^T$; its excess over L^T is the separation cost. Across all draws, median loss is $1.07L^T$ and the ninetieth percentile is $1.52L^T$.

Expected loss is $1.006L^T$, close to the strict target's. This average combines the gains from rules near the ideal response with the reporting costs of rules farther away. It is conditional on the chosen index and shock process.[23]

Index choice changes the comparison because cost-push shocks are endogenous to the measure of inflation in richer models. The same underlying disturbances can generate a smaller residual tradeoff under a different index. In Rubbo (2023), for example, a "divine coincidence" index eliminates the productivity-driven cost-push term from the aggregate Phillips curve. Stabilizing that index also stabilizes the output gap, although it need not

[23] A rule centered on a weaker response would place more mass below C. Mean relative responses from 0.7 to 1.3 yield expected losses between $1.00L^T$ and $1.04L^T$. The exercise reports the predictable component. On unforecastable innovations, average loss under the same distribution of rule slopes is 43.3 percent of the strict target's loss (Proposition 5(iii)).

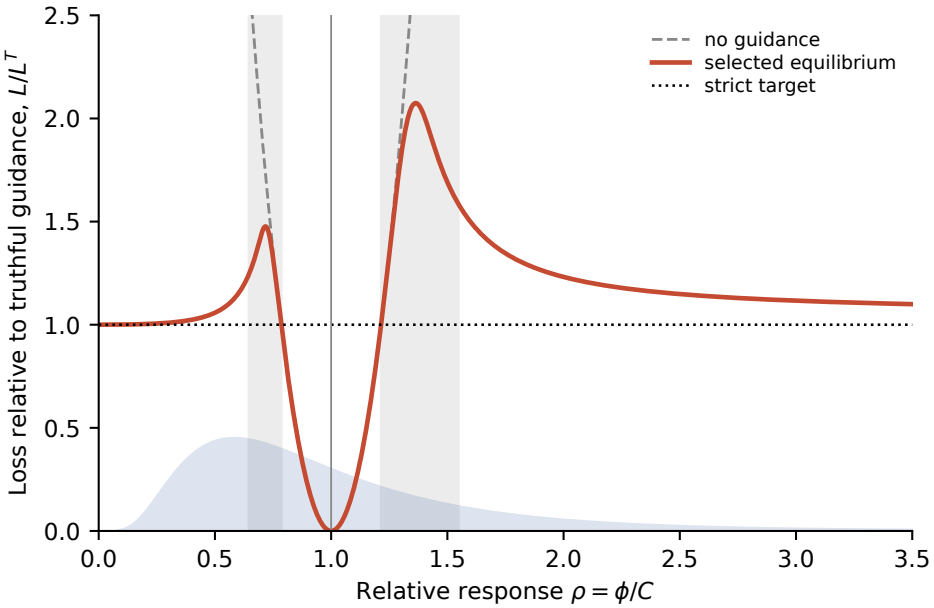

Figure 5: A standing rule with an uncertain relative response

Notes: Fixed inflation index, baseline calibration, $S = 3\sigma_s$, assured delivery, no reporting correction. Loss per unit of forecast variance, divided by the strict target's truthful-guidance loss L^T. Red: selected equilibrium; dashed: no guidance; dotted: strict target. Shaded: collapse bands $(\hat{\phi}, C - \bar{c})$ and $(C + \bar{c}, \hat{\phi}_+)$. Blue shading: lognormal density, rescaled for display, with mean one and log dispersion 0.60. Equilibria follow Proposition 3 below C and Lemma 14 above it; opacity is optimal on $[C - \bar{c}, C + \bar{c}]$. Computed in `code/simulate_mismatch.py`.

eliminate all welfare losses in the multisector economy. La'O and Tahbaz-Salehi (2022) likewise show how production linkages and nominal rigidities affect the index suited to optimal stabilization.

A strict target on a better-chosen index can therefore leave less for discretionary guidance to supply. The size of the ideal rule's advantage over such a target depends on the remaining tradeoff, which this exercise takes as given. Choosing the index also requires economic knowledge, and the model does not quantify that design problem. What it establishes is conditional: for any chosen index in this framework, a strict target aligns the choice of promise with the rule-constrained Ramsey benchmark when delivery is assured. A reactive rule can improve the response to residual shocks, but may require a reporting correction to preserve the value of guidance.

4.2.6 Mandates and the Designer's Knowledge

A mandate can restore guidance through two instruments. The announcement correction in Proposition 2 changes reporting incentives. The announcement format restricts the

largest admissible promise and supplies the boundary condition that selects a separating schedule. The following proposition distinguishes their roles.

Proposition 6 (Repairing a coarse rule)**.** *Under Assumption 1(a)–(c), let the announcement format admit exactly* $[-\bar{g}, \bar{g}]$.

(i) Correction and format. *With the announcement correction* $\mathcal{A}_R^{\phi}$ *of Proposition 2 in place and* $\bar{g} = \gamma^R S$*, the truthful reading is the unique separating equilibrium for every* ϕ *in the revelation region, the collapse band included.*

(ii) The format selects the reading. *Without the correction and* $0 < \phi \leq \hat{\phi}$*, every* $\bar{g} \in [\gamma_1 S, \gamma^R S]$ *selects as the unique separating equilibrium the schedule that solves the first-order condition from the point at which the largest forecast announces* $\bar{g}$*; for* $\bar{g} < \gamma_1 S$ *no separating equilibrium exists.*

Proof. See Online Appendix A.15.

Together, the correction and the format restore truthful guidance in the collapse band. The correction makes the truthful schedule satisfy each type's first-order condition. By admitting exactly its image, the format supplies the boundary condition that makes this the unique separating equilibrium. Designing both instruments requires model knowledge: the correction depends on stabilization tradeoffs, and the cap $\gamma^R S$ also requires the forecast bound.

With noisy delivery, the expected delivery sanction itself affects announcement incentives. Online Appendix A.17 constructs a *coarse-rule audit* with two components: one depends on the announcement and corrects reporting incentives; the other depends on the inflation miss and disciplines delivery. Once the announcement is made, the first component is fixed and leaves marginal delivery incentives unchanged.

Lemma 11 shows that this audit supports truthful reporting when the corrected reporting loss is strictly convex and sanctioning capacity covers both components. With no innovation, a candidate schedule calibrated to the Ramsey center delivers it exactly. Under the strict target with assured delivery, the announcement component is constant and only the obedience audit is needed.

The designer must know the candidate schedule, the traction function, the corrected loss over the admissible types and reports, and the required stake. Convexity must be checked separately from delivery: at $\phi = 0.2$, $\eta = 0.2$, and the baseline stake $\Psi = 18$, the corrected reporting loss is locally concave at truth. Raising the stake to $\Psi = 35$ restores positive local curvature; the appendix establishes global strict convexity for all sufficiently large stakes.

How far can that knowledge be economized on? To make the question precise, return to the period-0 design question of Section 3.2: can society leave the choice of the commitment promise to the central bank without knowing the stabilization tradeoffs that determine it? Across the models $\Omega \in \mathcal{M}$, hold the standing rule, report interval, candidate announcement schedule g_ϕ, and obedience audit fixed, and ask whether one announcement correction can align reporting incentives in all of them. In model Ω, with expectations formed from actual delivery and the delivery-valued reporting loss $\mathcal{J}_\phi^{\text{del}}$ of Online Appendix A.17, write $w_\phi(v;\Omega) \equiv \partial_r \mathcal{J}_\phi^{\text{del}}(v, r \mid \Omega)|_{r=v}$ for the reporting slope at truth, where v is the private component the report concerns. The answer follows from the form of the correction and from Proposition 4.

Corollary 3 (The designer's knowledge)**.** *(i)* Misspecification. *Let the reporting losses be continuously differentiable and the truthful reports interior. A correction computed from model $\widehat{\Omega}$ leaves, in model Ω, the marginal reporting incentive*

$$\mathcal{W}_\phi(v;\Omega,\widehat{\Omega}) = w_\phi(v;\Omega) - w_\phi(v;\widehat{\Omega}) \tag{48}$$

at truth. One correction makes truth stationary in every model of $\mathcal{M}$ if and only if $w_\phi(v;\Omega)$ is the same function of v across $\mathcal{M}$.

(ii) The strict target. *Under the strict target with assured delivery, for every model and every (v_1, s) the central bank's preferred promise is the rule-constrained Ramsey promise*

$$g_0^R(v_1, s) = BC\left[A(\eta v_1 + s) - v_1\right], \tag{49}$$

so truth is stationary in every model with no correction.

Proof. See Online Appendix A.19.

For a fixed candidate schedule and obedience audit, the relevant knowledge is the reporting slope at truth. Model errors that leave this slope unchanged preserve stationarity. If plausible models require different slopes at the same announcement, no common correction can make truth stationary in all of them. Increasing its stake cannot resolve the conflict: the residual incentive in at least one model is at least half the spread of the required slopes (Online Appendix A.19).

The knowledge requirement can be quantified. Under the pure forecast, the marginal correction is $-\phi\, AB(\beta C)^2 r$. The factor $AB(\beta C)^2$ is 0.86 at $\kappa = 0.10$, 0.79 at $\kappa = 0.15$, and 0.62 at $\kappa = 0.25$. A designer using the baseline slope when the true slope is 0.25 overstates this factor by 27%. For a positive forecast, the incentive at truth becomes $\phi[0.62 - 0.79]s < 0$

by (48), so the central bank has an incentive to over-promise. The parameters that size the correction are the parameters that move the collapse band in Supplementary Appendix Figure 3.

Under the strict target with assured delivery, the promise can instead be chosen by the central bank. Its preferred promise is the planner's in each model, even though the promise itself changes with the stabilization tradeoffs. The designer therefore need not know those tradeoffs to correct reporting incentives: there is no correction to specify. One promise transmits the relevant composite $A(\eta v_1 + s) - v_1$.[24] This is the informational advantage of delegation in the model: the institution aligns the choice of promise while leaving its calculation to the informed central bank. Partial codification can undermine that advantage by creating reporting incentives whose correction requires the designer to know the economic tradeoffs it left to discretion.

5 Conclusion

Coarse reactive rules create a tradeoff. They can improve stabilization and reduce the need to reveal forecasts, whose effect on current expectations can itself be costly. When a rule leaves enough of the desired response unsupplied, the benefit of corrective guidance exceeds this revelation cost. Yet the rule also affects the central bank's reporting incentives. For a range of intermediate rules, those incentives prevent ordinary forecasts from being communicated even though society would benefit from revelation. Promises are delivered, but the response that guidance should supply remains absent.

Delegation can support a coarse rule. The audit that does so assesses the announcement as well as its delivery, and we characterize it: the correction and the announcement format are computed from the stabilization tradeoffs that delegation was meant to leave to the central bank, and a correction set in one model fails in another. It is difficult to design upfront. A strict target needs no such correction: with delivery assured, the central bank chooses the promise the rule-constrained Ramsey planner would choose. The underlying incentive problem can also arise when a stable public understanding of policy supplies the standing response, even if that response is not written as a formal rule.

The welfare comparison depends on what remains to be stabilized after the inflation index is chosen. For a fixed index, reactive rules can improve the response to residual shocks. A better index can reduce that residual by accommodating relative-price adjust-

[24] This result concerns the choice of promise conditional on assured delivery. With noisy delivery, the expected audit burden can depend on the hidden state even under the strict target, and a reporting correction may be needed (Online Appendix A.19).

ment within the target itself. Our numerical comparison measures the first margin while holding the second fixed. It therefore leaves open how much an ideal state-contingent rule would improve on a strict target applied to an economically appropriate index. The strict target's institutional advantage is that, conditional on assured delivery, society can leave the commitment promise to the informed central bank without a model-calibrated reporting correction.

Coarse reactive rules can improve stabilization when they are accurate, but they obstruct the guidance needed to complete their response when they are not.

References

Ricardo Alonso and Niko Matouschek. Optimal delegation. *Review of Economic Studies*, 75(1): 259–293, 2008.

Manuel Amador, Iván Werning, and George-Marios Angeletos. Commitment versus flexibility. *Econometrica*, 74(2):365–396, 2006. doi: 10.1111/j.1468-0262.2006.00666.x.

Susan Athey, Andrew Atkeson, and Patrick J. Kehoe. The optimal degree of discretion in monetary policy. *Econometrica*, 73(5):1431–1475, 2005. doi: 10.1111/j.1468-0262.2005.00626.x.

Jeffrey S. Banks and Joel Sobel. Equilibrium selection in signaling games. *Econometrica*, 55(3): 647–661, 1987.

Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121, 1983.

Marco Bassetto. Forward guidance: Communication, commitment, or both? *Journal of Monetary Economics*, 108:69–86, 2019. doi: 10.1016/j.jmoneco.2019.08.015.

B. Douglas Bernheim and Sergei Severinov. Bequests as signals: An explanation for the equal division puzzle. *Journal of Political Economy*, 111(4):733–764, 2003. doi: 10.1086/375383.

In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221, 1987.

Richard Clarida, Jordi Galí, and Mark Gertler. The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707, 1999.

Alessandro Dovis and Rishabh Kirpalani. Rules without commitment: Reputation and incentives. *The Review of Economic Studies*, 88(6):2833–2856, 2021. doi: 10.1093/restud/rdab006.

Mariia Elkina. Reputation-driven design of monetary policy rules. Working paper, 2026.

Jon Faust and Jonathan H. Wright. Comparing Greenbook and reduced form forecasts using a large realtime dataset. *Journal of Business & Economic Statistics*, 27(4):468–479, 2009.

Ippei Fujiwara and Yuichiro Waki. The delphic forward guidance puzzle in New Keynesian models. *Review of Economic Dynamics*, 46:280–301, 2022. doi: 10.1016/j.red.2021.10.002.

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

Laura Gáti and Amy Handlan. Reputation for confidence. ECB Working Paper Series 3141, European Central Bank, 2025.

Hans Gersbach and Volker Hahn. Inflation forecast contracts. *Journal of Economic Dynamics and Control*, 48:26–40, 2014.

Hans Gersbach, Volker Hahn, and Yulin Liu. Forward guidance contracts. *Macroeconomic Dynamics*, 23(8):3386–3423, 2019.

Marina Halac and Pierre Yared. Fiscal rules and discretion under limited enforcement. *Econometrica*, 90(5):2093–2127, 2022.

Tom D. Holden. Robust real rate rules. *Econometrica*, 92(5):1521–1551, 2024. doi: 10.3982/ECTA21069.

Bengt Holmström. On the theory of delegation. In Marcel Boyer and Richard E. Kihlstrom, editors, *Bayesian Models in Economic Theory*, pages 115–141. North-Holland, Amsterdam, 1984.

Luigi Iovino, Jennifer La'O, and Rui Mascarenhas. Optimal monetary policy and disclosure with an informationally-constrained central banker. *Journal of Monetary Economics*, 125:151–172, 2022.

Navin Kartik. Strategic communication with lying costs. *The Review of Economic Studies*, 76(4): 1359–1395, 2009. doi: 10.1111/j.1467-937X.2009.00559.x.

Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491, 1977.

Jennifer La'O and Alireza Tahbaz-Salehi. Optimal monetary policy in production networks. *Econometrica*, 90(3):1295–1336, 2022. doi: 10.3982/ECTA18627.

Susanne Lohmann. Optimal commitment in monetary policy: Credibility versus flexibility. *American Economic Review*, 82(1):273–286, March 1992.

George J. Mailath. Incentive compatibility in signaling games with a continuum of types. *Econometrica*, 55(6):1349–1365, 1987.

Christina D. Romer and David H. Romer. Federal reserve information and the behavior of interest rates. *American Economic Review*, 90(3):429–457, 2000.

Elisa Rubbo. Networks, Phillips curves, and monetary policy. *Econometrica*, 91(4):1417–1455, 2023. doi: 10.3982/ECTA18654.

Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973.

Yuichiro Waki, Richard Dennis, and Ippei Fujiwara. The optimal degree of monetary discretion in a New Keynesian model with private information. *Theoretical Economics*, 13(3):1319–1367, 2018. doi: 10.3982/TE2369.

Carl E. Walsh. Optimal contracts for central bankers. *American Economic Review*, 85(1):150–167, March 1995.

Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ, 2003.

Online Appendix — Extended (Not for Publication)

Online Appendix to "Monetary Policy Delegation with Operational Rules and Sanctioned Promises" (Lukasz A. Drozd). It is intended for online publication and contains the omitted proofs in an extended version, the supporting results the main text cites, and the numerical details referenced in the main text. The infinite-horizon extension is in a separate supplementary appendix.

Conventions. Throughout we use the optimized schedules $\theta_1(v_1, q)$, $x_1(v_1, q)$, and $\theta_2(v_2, q)$ of Section 3. The revelation comparisons for the current state restrict messages to $q = q(v_1)$. Proofs for the pure forecast hold the current state public and index types by s; reporting calculations for guidance about current conditions set $s = 0$ almost surely and index types by v_1. Neither restriction is imposed on the proof of delegated selective opacity (Appendix A.9), which allows both private components. Appendix A proves the results in the order of the main text. Appendix B develops noisy delivery, rules above the ideal response, and guidance about current conditions (Section B.3.1). Appendix C records the numerical example, and Appendix D extends the results to an infinite horizon. A separate supplementary appendix collects supporting material the main text cites in passing (Supplementary Appendix 2) and the operating implementation through interest rates (Supplementary Appendix 1.1).

A Omitted Proofs

This appendix proves the results of Sections 3 and 4 in the order of the main text: the Ramsey benchmark and discretion (A.1–A.7), the ideal rule (A.8–A.12), and coarse rules (A.13–A.19). Statements the main text cites without proof are restated where they are proved. Online Appendix Figure A.1 summarizes the timing and who knows what.

A.1 Message Compression

Proof of Lemma 1. Fix a public message q and an expected period-2 center $e_1(q)$.

Part (i): the optimal continuation schedule. Consider all square-integrable schedules $\theta_2(\mathcal{T}, v_2, q)$ satisfying $\mathbb{E}[\theta_2(\mathcal{T}, v_2, q) \mid q] = e_1(q)$. We show that $\theta_2 = Cv_2 + g(q)$ uniquely minimizes continuation loss, where $g(q) = e_1(q) - C\widehat{v}_2(q)$.

By (11), which completes the square using the period-2 Phillips curve (8), the expected

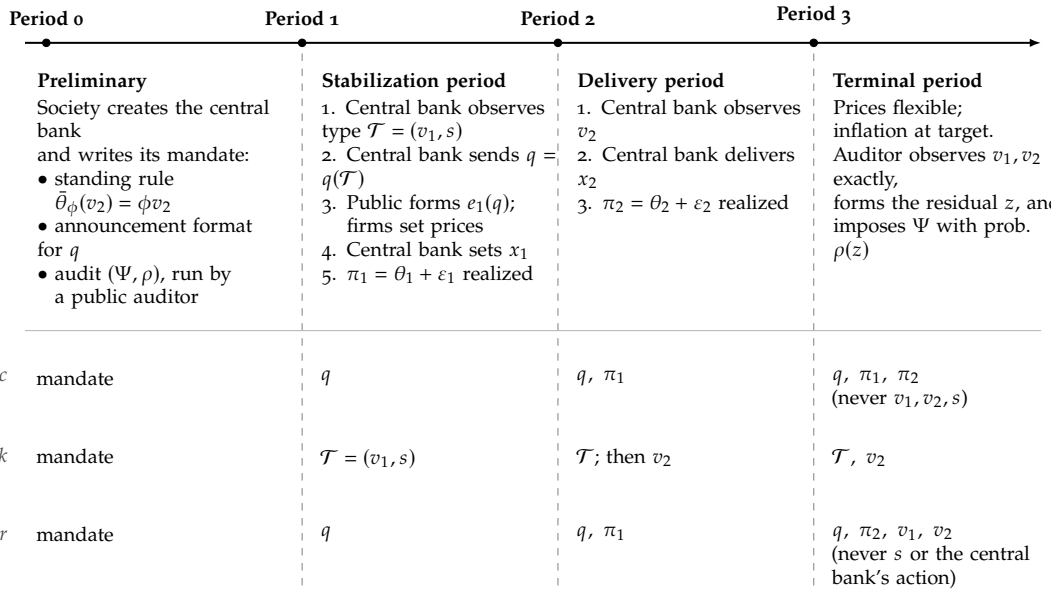

Figure A.1: Timing, information, and institutions in the policy episode

Notes: The four periods of the policy episode (top) and the episode-specific information available to each party at each date (bottom). The structural model is known to the central bank and the public; Section 3.2 specifies the designer's information. In the preliminary period society writes the mandate: the standing rule $\bar{\theta}_\phi(v_2) = \phi v_2$ with slope ϕ (the strict target is $\phi = 0$, the ideal rule $\phi = C$), the format of the announcement q, and the audit, a sanction Ψ imposed with probability $\rho(z)$ by a public auditor. In the stabilization period the central bank observes its type $\mathcal{T} = (v_1, s)$, the nowcast and the forecast, sends the announcement q, the public forms the expectation $e_1(q)$ and firms set prices, and the central bank sets x_1. In the delivery period the central bank observes v_2 and delivers x_2. In the terminal period prices are flexible, the auditor observes v_1 and v_2 exactly, forms the residual z of realized inflation from the rule and the announcement, and imposes the sanction. The public never observes the central bank's private information; the auditor observes the shocks after the fact but never the forecast s or the central bank's action, so the audit reads realized inflation alone.

period-2 loss conditional on q is

$$\mathbb{E}\left[\frac{A}{2}(\theta_2 - Cv_2)^2 + \frac{C}{2}v_2^2 \,\middle|\, q\right]. \tag{A.1}$$

Because the term $\frac{C}{2}v_2^2$ is independent of policy, minimizing this conditional loss reduces to minimizing the policy-relevant expected squared deviation from the ex-post ideal:

$$\frac{A}{2}\,\mathbb{E}\big[(\theta_2 - Cv_2)^2 \mid q\big]. \tag{A.2}$$

Let θ_2 be any admissible schedule satisfying $\mathbb{E}[\theta_2 \mid q] = e_1(q)$. Define the residual u by

$$\theta_2 = Cv_2 + g(q) + u, \tag{A.3}$$

where $g(q) \equiv e_1(q) - C\widehat{v}_2(q)$ is a fixed constant given q. This decomposition defines u for each admissible schedule. Taking conditional expectations gives

$$\mathbb{E}[\theta_2 \mid q] = C\mathbb{E}[v_2 \mid q] + g(q) + \mathbb{E}[u \mid q]. \tag{A.4}$$

Substituting the constraint $\mathbb{E}[\theta_2 \mid q] = e_1(q)$ and the definition $\mathbb{E}[v_2 \mid q] = \widehat{v}_2(q)$, we obtain

$$e_1(q) = C\widehat{v}_2(q) + \left[e_1(q) - C\widehat{v}_2(q)\right] + \mathbb{E}[u \mid q], \tag{A.5}$$

which simplifies to $\mathbb{E}[u \mid q] = 0$. This is the constraint now rewritten in terms of u.

The optimization problem thus reduces to choosing a schedule u to minimize the objective, subject only to the constraint $\mathbb{E}[u \mid q] = 0$. Substituting $\theta_2 - Cv_2 = g(q) + u$ into the objective (A.2) gives:

$$\frac{A}{2}\,\mathbb{E}\left[(g(q) + u)^2 \mid q\right] = \frac{A}{2}\,\mathbb{E}\left[g(q)^2 + 2g(q)u + u^2 \mid q\right]. \tag{A.6}$$

Using the linearity of expectations and the fact that $g(q)$ is constant given q, this expands to

$$\frac{A}{2}\Big(g(q)^2 + 2g(q)\mathbb{E}[u \mid q] + \mathbb{E}[u^2 \mid q]\Big). \tag{A.7}$$

Applying the constraint $\mathbb{E}[u \mid q] = 0$, the cross term vanishes, leaving

$$\frac{A}{2}g(q)^2 + \frac{A}{2}\mathbb{E}[u^2 \mid q]. \tag{A.8}$$

The first term is fixed, and the second is nonnegative. The feasible choice $u = 0$ almost surely attains the minimum. Conversely, $\mathbb{E}[u^2 \mid q] = 0$ requires $u = 0$ almost surely, proving uniqueness. The associated output gap is $x_2 = (\theta_2 - v_2)/\kappa$ by (8).

Part (ii): policy and expected loss for a fixed communication strategy. The strategy $q(\cdot)$ is fixed and known to the public. Write its posterior means as $\widehat{v}_i(q) = \mathbb{E}[v_i \mid q]$ and let $B \equiv 1/(A + \beta C)$.

The planner's objective is to minimize the expected discounted loss:

$$\frac{1}{2}\mathbb{E}\left[\theta_1^2 + \lambda x_1^2 + \beta(\theta_2^2 + \lambda x_2^2)\right], \tag{A.9}$$

subject to the expectation constraint $e_1(q) = \mathbb{E}[\theta_2 \mid q]$ and the period-1 and period-2 Phillips curves:

$$\theta_1 = v_1 + \kappa x_1 + \beta e_1(q), \qquad \theta_2 = v_2 + \kappa x_2. \tag{A.10}$$

Because the period-1 loss depends on period 2 only through the scalar $e_1(q)$, we can solve the problem in stages.

Step 1: Period-1 optimization for a given $e_1(q)$. The period-1 central bank chooses x_1 knowing v_1 and $e_1(q)$. It minimizes $\frac{1}{2}(\theta_1^2 + \lambda x_1^2)$ subject to $\theta_1 = (v_1 + \beta e_1(q)) + \kappa x_1$. This is strictly convex in x_1, yielding the first-order condition $\kappa\theta_1 + \lambda x_1 = 0$.

Substituting the constraint into this first-order condition gives $x_1 = -\frac{\kappa}{\kappa^2+\lambda}(v_1 + \beta e_1(q))$, and therefore $\theta_1 = C(v_1 + \beta e_1(q))$, where $C \equiv \lambda/(\kappa^2 + \lambda)$. Substituting these optimal choices back into the period-1 objective yields the minimized current loss:

$$\frac{C}{2}\big(v_1 + \beta e_1(q)\big)^2. \tag{A.11}$$

Step 2: Period-2 optimization for a given $e_1(q)$. By Part (i), the optimal continuation schedule satisfying the expectation constraint is $\theta_2 = Cv_2 + g(q)$, where $g(q) = e_1(q) - C\widehat{v}_2(q)$. The expected period-2 loss conditional on q is therefore the variance-free minimum:

$$\frac{A}{2}g(q)^2 + \frac{C}{2}\,\mathbb{E}[v_2^2 \mid q], \tag{A.12}$$

as we have shown in part (i).

Step 3: Choice of the expectation schedule $e_1(\cdot)$. By the law of iterated expectations, the total period-0 loss is $\mathbb{E}[\ell(q)]$, where the conditional loss $\ell(q)$ combines the minimized results from Steps 1 and 2 above:

$$\ell(q) = \frac{C}{2}\,\mathbb{E}\Big[\big(v_1 + \beta e_1(q)\big)^2 \Big| q\Big] + \beta\left[\frac{A}{2}g(q)^2 + \frac{C}{2}\,\mathbb{E}[v_2^2 \mid q]\right]. \tag{A.13}$$

The planner minimizes this loss for each possible message q. First, expand the expectation in the period-1 term: $\mathbb{E}[(v_1 + \beta e_1(q))^2 \mid q] = (\widehat{v}_1(q) + \beta e_1(q))^2 + \mathrm{Var}(v_1 \mid q)$. Next, substitute the definition $g(q) = e_1(q) - C\widehat{v}_2(q)$ to write $\ell(q)$ entirely in terms of $e_1(q)$:

$$\ell(q) = \frac{C}{2}\big(\widehat{v}_1(q) + \beta e_1(q)\big)^2 + \frac{\beta A}{2}\big(e_1(q) - C\widehat{v}_2(q)\big)^2 + \frac{C}{2}\,\mathrm{Var}(v_1 \mid q) + \frac{\beta C}{2}\mathbb{E}[v_2^2 \mid q]. \tag{A.14}$$

The last two terms are independent of $e_1(q)$. The first two terms are a strictly convex quadratic function in $e_1(q)$. The first-order condition with respect to $e_1(q)$ is:

$$\beta C\big(\widehat{v}_1(q) + \beta e_1(q)\big) + \beta A\big(e_1(q) - C\widehat{v}_2(q)\big) = 0. \tag{A.15}$$

Dividing by β and grouping the $e_1(q)$ terms yields $(A + \beta C)e_1(q) = AC\widehat{v}_2(q) - C\widehat{v}_1(q)$.

Multiplying by $B \equiv 1/(A+\beta C)$ gives the optimal expectation:

$$e_1^R(q) = BC\big[A\widehat{v}_2(q) - \widehat{v}_1(q)\big]. \tag{A.16}$$

To find the optimal promise $g^R(q)$, substitute $e_1^R(q)$ back into $g(q) = e_1(q) - C\widehat{v}_2(q)$:

$$\begin{aligned} g^R(q) &= BC\big[A\widehat{v}_2(q) - \widehat{v}_1(q)\big] - C\widehat{v}_2(q) \\ &= -BC\widehat{v}_1(q) + C\widehat{v}_2(q)(AB-1). \end{aligned} \tag{A.17}$$

Because $AB - 1 = \frac{A-(A+\beta C)}{A+\beta C} = -\beta CB$, this becomes

$$g^R(q) = -BC\big[\widehat{v}_1(q) + \beta C\widehat{v}_2(q)\big] \equiv -BC\, M(q). \tag{A.18}$$

Substituting $e_1^R(q)$ into the Step-1 center gives $\theta_1^R = C[v_1 + \beta e_1^R(q)]$, and substituting $g^R(q)$ into the Step-2 schedule gives $\theta_2^R = Cv_2 + g^R(q)$.

Step 4: The value of the strategy. To compute the minimized loss, we evaluate the first two terms of $\ell(q)$ at the optimum. Notice that the period-1 composite becomes $\widehat{v}_1(q) + \beta e_1^R(q) = \widehat{v}_1(q) + B\beta C[A\widehat{v}_2(q) - \widehat{v}_1(q)] = AB[\widehat{v}_1(q) + \beta C\widehat{v}_2(q)] = AB\, M(q)$.

Substituting this and $g^R(q) = -BC\, M(q)$ into the conditional loss yields:

$$\ell(q) = \left[\frac{C}{2}A^2B^2 + \frac{\beta A}{2}C^2B^2\right] M(q)^2 + \frac{C}{2}\operatorname{Var}(v_1 \mid q) + \frac{\beta C}{2}\mathbb{E}[v_2^2 \mid q]. \tag{A.19}$$

The bracket simplifies to $\frac{1}{2}AB^2C(A+\beta C) = \frac{1}{2}ABC$. Taking the unconditional expectation over q gives the final minimized period-0 loss for the strategy:

$$\mathbb{E}[\ell(q)] = \frac{1}{2}ABC\,\mathbb{E}\big[M(q)^2\big] + \frac{C}{2}\,\mathbb{E}\big[\operatorname{Var}(v_1 \mid q)\big] + \frac{\beta C}{2}\mathbb{E}[v_2^2]. \tag{A.20}$$

This establishes (18) and completes the proof of Lemma 1.

A.2 Selective Opacity

Proof of Corollary 2. The conventions of the proof of Lemma 1(ii) apply. Fix a communication strategy about the current state, $q_1 = q_1(v_1)$, and let $q_f = q_f(s)$ be an auxiliary message generated from the pure forecast s, so the combined message is $q = (q_1, q_f)$. Let q_f^0 denote a constant, uninformative auxiliary message. Because q_f^0 conveys no information, the information set under (q_1, q_f^0) is identical to q_1 alone.

Let $\mathcal{L}(q)$ denote the minimized period-0 loss $\mathbb{E}[\ell(q)]$ in (A.20). We compare $\mathcal{L}(q_1, q_f)$

with $\mathcal{L}(q_1, q_f^0)$.

Step 1: What the auxiliary message changes in (A.20). First, recall the definition of the inflationary pressure term from Lemma 1(ii):

$$M(q) \equiv \widehat{v}_1(q) + \beta C \widehat{v}_2(q). \tag{A.21}$$

From the forecast structure, $v_2 = \eta v_1 + s + \varepsilon$. Taking the conditional expectation given q gives $\widehat{v}_2(q) = \eta \widehat{v}_1(q) + \mathbb{E}[s \mid q]$. Substituting this into the definition of $M(q)$ gives

$$M(q) = (1 + \beta C \eta)\widehat{v}_1(q) + \beta C \, \mathbb{E}[s \mid q]. \tag{A.22}$$

Under the baseline strategy (q_1, q_f^0), the public only learns from q_1. Because q_1 is a function of v_1 and s is independent of v_1, observing q_1 conveys no information about s. Thus, $\mathbb{E}[s \mid q_1, q_f^0] = 0$. Substituting this into the expression above gives

$$M(q_1, q_f^0) = (1 + \beta C \eta)\widehat{v}_1(q_1). \tag{A.23}$$

Now evaluate the combined message $q = (q_1, q_f)$ (forecast-augmented message). Because q_f is generated from s, and s is independent of v_1, appending q_f tells the public nothing new about v_1. This implies $\mathbb{E}[v_1 \mid q] = \mathbb{E}[v_1 \mid q_1]$, so $\widehat{v}_1(q) = \widehat{v}_1(q_1)$. It also implies $\mathrm{Var}(v_1 \mid q) = \mathrm{Var}(v_1 \mid q_1)$, meaning the conditional variance term in (A.20) is identical under q and (q_1, q_f^0).[25]

Substituting $\widehat{v}_1(q) = \widehat{v}_1(q_1)$ into the expression for $M(q)$, we can express the inflationary pressure of the combined message exactly in terms of the uninformative baseline message:

$$M(q) = M(q_1, q_f^0) + \beta C \, \mathbb{E}[s \mid q]. \tag{A.24}$$

Step 2: The cross term and recombination. To substitute $M(q)$ back into the loss function (A.20), we square the identity from Step 1. Taking unconditional expectations gives:

$$\mathbb{E}\big[M(q)^2\big] = \mathbb{E}\Big[M(q_1, q_f^0)^2\Big] + 2\beta C \, \mathbb{E}\big[M(q_1, q_f^0)\mathbb{E}[s \mid q]\big] + (\beta C)^2 \, \mathbb{E}\big[\mathbb{E}[s \mid q]^2\big]. \tag{A.25}$$

We evaluate the cross term using the law of iterated expectations. Because $M(q_1, q_f^0)$ is exclusively a function of v_1, we condition on v_1 to factor it out. Since s is independent of

[25] Independence matters only when the current-state message pools types. In the leading case of full revelation, $q_1 = v_1$, the step is trivial ($\mathbb{E}[v_1 \mid v_1, q_f] = v_1$) and only $\mathbb{E}[s \mid v_1] = 0$ is used for the cross term below. For a coarser q_1, mean independence alone would not suffice: if the distribution of s depended on v_1 through higher moments, a message about s could separate the pooled types, changing $\widehat{v}_1(q)$ and $\mathrm{Var}(v_1 \mid q)$.

v_1, this gives:

$$\mathbb{E}\big[M(q_1, q_f^0)s\big] = \mathbb{E}\big[M(q_1, q_f^0)\mathbb{E}[s \mid v_1]\big] = 0. \tag{A.26}$$

With the cross term vanishing, we substitute $\mathbb{E}[M(q)^2]$ back into the full loss function (A.20):

$$\mathcal{L}(q) = \frac{1}{2}ABC\Big(\mathbb{E}\Big[M(q_1, q_f^0)^2\Big] + (\beta C)^2\,\mathbb{E}\big[\mathbb{E}[s \mid q]^2\big]\Big) + \frac{C}{2}\,\mathbb{E}\big[\mathrm{Var}(v_1 \mid q_1)\big] + \frac{\beta C}{2}\mathbb{E}[v_2^2]. \tag{A.27}$$

Regrouping the terms recovers the exact loss under the uninformative baseline:

$$\mathcal{L}(q) = \mathcal{L}(q_1, q_f^0) + \frac{1}{2}ABC(\beta C)^2\mathbb{E}\big[\mathbb{E}[s \mid q]^2\big]. \tag{A.28}$$

Thus, $\mathcal{L}(q) - \mathcal{L}(q_1, q_f^0) = \frac{1}{2}ABC(\beta C)^2\mathbb{E}\big[\mathbb{E}[s \mid q]^2\big] \geq 0$. Revelation is optimal if and only if the excess loss is zero, requiring $\mathbb{E}[s \mid q] = 0$ almost surely.

Consider now the special case with no current information ($v_1 = 0$ almost surely). Here, $\widehat{v}_1(q) = 0$ and $\mathrm{Var}(v_1 \mid q) = 0$ for every message. The uninformative baseline loss $\mathcal{L}(q_1, q_f^0)$ reduces to a constant $\frac{\beta C}{2}\mathbb{E}[v_2^2]$, making a constant message q_f^0 optimal. Under a constant message, $\widehat{v}_1 = \widehat{v}_2 = 0$, and Lemma 1(ii) implies $e_1^R(q) = g^R(q) = 0$, $\theta_1^R(0, q) = 0$, and $\theta_2^R(v_2, q) = Cv_2$.

Two implications are useful for the analysis below.

Remark 1 (Cost of full revelation). Under full revelation of v_1 and full revelation of the forecast, $\mathbb{E}[s \mid q] = s$, so the additional loss is $\frac{1}{2}ABC(\beta C)^2\,\mathrm{Var}(s)$: a larger share of the conditional future-shock variance known privately in period 1 raises the welfare cost of publishing that information before current prices are set.

Remark 2 (Pure-forecast benchmark). The pure-forecast case of Corollary 2 carries the following bound, which also holds under delegation. Let $v_1 = 0$ almost surely, with v_2 mean zero and of finite variance. For any policy and any communication strategy, $\mathcal{L} \geq \frac{\beta C}{2}\mathbb{E}[v_2^2]$, and, under a constant message, the allocation $e_1 = 0$, $\theta_1 = x_1 = 0$, $\theta_2 = Cv_2$ attains this bound.[26]

[26] This bound holds because current loss is nonnegative and, by (11), period-2 loss is $\frac{A}{2}(\theta_2 - Cv_2)^2 + \frac{C}{2}v_2^2$; the stated allocation makes the current loss and the first period-2 term zero.

A.3 Auxiliary Results: Representation and Discretion

Corollary 4 (Representation)**.** *Every solution of the unrestricted Ramsey program* (13)–(14) *can be represented, up to almost-sure equality, by the compressed policy vector*

$$\mathcal{P}_0 = \big(q(\mathcal{T}),\ x_1(v_1, q),\ x_2(v_2, q),\ \theta_1(v_1, q),\ \theta_2(v_2, q)\big)\,, \tag{A.29}$$

with $\theta_2(v_2, q) = Cv_2 + g^R(q)$, $\theta_1(v_1, q) = C[v_1 + \beta e_1^R(q)]$, *and the output gaps given by the mean Phillips curves. The type enters current policy only through* v_1 *and continuation policy only through the message.*

Proof of Corollary 4. By Lemma 1(i), for any message the loss-minimizing continuation center is $Cv_2 + g(q)$, which depends on $(\mathcal{T}, v_2, q)$ only through (v_2, q); by part (ii) the optimal shift is $g^R(q)$ and the optimal current center is $C[v_1 + \beta e_1^R(q)]$, which depends on $\mathcal{T}$ only through v_1 (Step 1 of that proof). The output gaps follow from the mean Phillips curves in (14). For the same communication strategy, any other solution differs from this one on a null set by the uniqueness statements in the lemma. Uniqueness holds given the message; across communication strategies several may be optimal, as at the revelation threshold, where all are payoff-equivalent.

The allocation without a mandate, cited in Section 3.2, is the following.

A.4 Discretion without a Charge

Proof of Lemma 2. Part (i): the period-2 best response. With no sanction the period-2 central bank minimizes $\frac{1}{2}[\theta_2^2 + \lambda(x_2 - x^*)^2]$ subject to the period-2 Phillips curve (8) alone. The gap target enters this objective exactly as a shock of size κx^* would, so the square-completion identity (11) applies with $v_2 + \kappa x^*$ in place of v_2 and gives

$$\theta_2 = C(v_2 + \kappa x^*) \tag{A.30}$$

for every announcement and every type.

Part (ii): what announcements reveal. By part (i) the delivered center is $Cv_2 + C\kappa x^*$ whatever was announced, so period-2 loss does not depend on the message, and the expectation the public forms is

$$e_1(q) = \mathbb{E}[\theta_2 \mid q] = C\eta\mathbb{E}[v_1 \mid q] + C\kappa x^*, \tag{A.31}$$

using $\mathbb{E}[v_2 \mid q] = \eta\mathbb{E}[v_1 \mid q]$ in the scalar benchmark. Write $m = \mathbb{E}[v_1 \mid q]$ for the posterior mean induced by the equilibrium message of type v_1, and let (v_1', m') be an independent

copy. Minimizing the period-1 objective as in Step 1 of the proof of Lemma 1(ii), now with the bias carried through the Phillips curve, leaves the message-dependent part of the period-1 loss equal to

$$\frac{C}{2}\big(v_1 + \beta C\eta m + \kappa x^*(1+\beta C)\big)^2. \tag{A.32}$$

Sequential optimality implies that type v_1 weakly prefers its own equilibrium message, with posterior mean m, to the equilibrium message of type v_1', with posterior mean m'. Interchanging the two types gives the reverse incentive inequality. Adding the two inequalities cancels the bias intercept $\kappa x^*(1+\beta C)$ along with the symmetric term in $m + m'$, leaving

$$(v_1 - v_1')(m - m') \le 0, \tag{A.33}$$

where we use $\eta > 0$. The bias therefore leaves the reporting incentive untouched. Taking expectations over the independent pair gives

$$\mathrm{Cov}(v_1, m) \le 0. \tag{A.34}$$

Bayesian consistency gives the opposite inequality. Since $m = \mathbb{E}[v_1 \mid q]$,

$$\mathrm{Cov}(v_1, m) = \mathbb{E}[v_1 m] - \mathbb{E}[v_1]\mathbb{E}[m] = \mathbb{E}[m^2] - \mathbb{E}[m]^2 = \mathrm{Var}(m) \ge 0. \tag{A.35}$$

Hence $\mathrm{Var}(m) = 0$. Because $\mathbb{E}[m] = \mathbb{E}[v_1] = 0$, it follows that $m = 0$ almost surely and therefore $e_1(q) = C\kappa x^*$. The argument allows any number of on-path messages and concerns the posterior mean alone: higher moments of v_1 may remain informative.

Part (iii): period-1 policy. With expectations anchored at $e_1(q) = C\kappa x^*$ the period-1 central bank minimizes $\frac{1}{2}[\theta_1^2 + \lambda(x_1 - x^*)^2]$ subject to $\theta_1 = v_1 + \kappa x_1 + \beta C\kappa x^*$. Its first-order condition $\kappa\theta_1 + \lambda(x_1 - x^*) = 0$ gives the policies stated in (20).

Proof of Corollary 1. With $\theta_2 = g$ delivered and $e_1 = g$, the period-1 loss minimized over θ_1 is $\frac{C}{2}(v_1 + \beta g + \kappa x^*)^2$, as in the proof of Lemma 2(iii) with $e_1 = g$ in place of $C\kappa x^*$, and the expected period-2 loss is $\frac{A}{2}\big(g - C(\widehat{v}_2 + \kappa x^*)\big)^2$ plus a term independent of g, by (11) with $v_2 + \kappa x^*$ in place of v_2. The first-order condition $C(v_1 + \beta g + \kappa x^*) + A\big(g - C\widehat{v}_2 - C\kappa x^*\big) = 0$ gives $g(A + \beta C) = AC\widehat{v}_2 - Cv_1 + C\kappa x^*(A - 1)$, and $A - 1 = AC$. In the stationary Nash equilibrium of the repeated economy, with the following period's promise $\bar{g}$ taken as given, the period-$(t+1)$ loss is $\frac{1}{2}\big[g^2 + \frac{\lambda}{\kappa^2}(g - \beta\bar{g} - \kappa x^*)^2\big]$, and the same condition at $g = \bar{g}$ gives $\bar{g}\,[1 + \beta C + AC(1-\beta)] = AC^2\kappa x^*$. Full commitment with the same objective chooses the steady state on the long-run Phillips curve $x = \pi(1-\beta)/\kappa$, which is zero at $\beta = 1$; the numbers in the text are computed in `code/promise_no_rule.py`.

A.5 Gap Component of the Audit

The audit of Section 3.2 carries a component whose probability rises linearly with the output gap; with a risk-neutral central bank its expected cost is $T_t = \lambda x^* x_t$ per period up to a constant. The following lemma records what that component does to the central bank's objective and why it is levied on the output gap.

Lemma 5 (The gap component of the audit restores the objective)**.** *Let the central bank's period-t loss be $\frac{1}{2}[\theta_t^2 + \lambda(x_t - x^*)^2]$, as in (4), and add the expected sanction of the gap component, $T_t = \lambda x^* x_t$ up to a constant, in every period.*

(i) The resulting period-t loss equals $\frac{1}{2}[\theta_t^2 + \lambda x_t^2] + \frac{1}{2}\lambda(x^*)^2$ for every (θ_t, x_t). The central bank's announcement problem (28) and delivery problem (26) therefore coincide, for every announcement, belief, and shock, with those of a central bank with $x^* = 0$.

(ii) Its preferred response is society's in both periods, $\theta_2 = Cv_2$ and $\theta_1 = C[v_1 + \beta e_1]$; under an uninformative announcement $e_1 = 0$ and $\theta_1 = Cv_1$; and in every symmetric equilibrium of the announcement game a zero forecast announces exactly zero.

The charge depends on x^*, which society must know, and on nothing the central bank privately observes.

Proof. Part (i). Expanding the square in the period-t loss,

$$\frac{1}{2}\left[\theta_t^2 + \lambda(x_t - x^*)^2\right] + \lambda x^* x_t = \frac{1}{2}\theta_t^2 + \frac{1}{2}\lambda x_t^2 - \lambda x^* x_t + \frac{1}{2}\lambda(x^*)^2 + \lambda x^* x_t = \frac{1}{2}\left[\theta_t^2 + \lambda x_t^2\right] + \frac{1}{2}\lambda(x^*)^2. \tag{A.36}$$

The remainder is a constant that no choice affects, so every problem the central bank solves under the component has the same minimizers as under the social loss (3). The expected sanction (25) is added to both objectives alike, which gives (26) and (27) as written.

Part (ii). With the objective restored, the period-2 first-order condition is that of (11), whose unconstrained minimizer is $\theta_2 = Cv_2$, and the period-1 problem in Step 1 of the proof of Lemma 1(ii) gives $\theta_1 = C[v_1 + \beta e_1]$, which is Cv_1 when $e_1 = 0$. Under an uninformative announcement the public expects the delivered center, $\phi\,\mathbb{E}[v_2] + g = 0$ at $g = 0$. In the pure-forecast announcement game of Appendix A.15, the loss (A.95) is invariant under $(s, g, \widehat{s}) \mapsto (-s, -g, -\widehat{s})$, so in a symmetric equilibrium the announcement schedule is odd and $G(0) = 0$.

Remark 3 (Inflation contracts as a tax on talk). A component levied on inflation, $\mu^* \theta_t$ with $\mu^* = \lambda x^*/\kappa$, offsets the term the divergent target adds to each period's first-order condition and so also delivers $\theta_t = Cv_t$. It does not deliver the announcement game. In period 1 the

Phillips intercept is βe_1, which the announcement moves: a promise raises expected future inflation, and expected future inflation raises current inflation, so a component on inflation prices forward guidance. It leaves $\beta\mu^*$ in the marginal cost of a promise, the announcement schedule is no longer odd, and the central bank under-promises to reduce what it pays. Under the component on the gap the cancellation happens in the objective, so nothing is left to price. A component on the audit residual would instead enter the delivery problem linearly and make the central bank over-deliver at a zero promise.

A.6 Conditional Message Loss and Revelation Threshold

Proof of Lemma 15. Step 1: the conditional loss and the optimal promise. By Lemma 1(i), the ideal-rule continuation is $\theta_2(v_2, q) = Cv_2 + g$. Current-policy optimization gives minimized loss $C[v_1 + \beta e_1(q)]^2/2$, with $e_1(q) = C\widehat{v}_2(q) + g$. Averaging within a posterior cell and adding continuation loss yields the conditional loss of (A.13) as a function of the message and the promise,

$$\ell(q, g) = Z(q) + \frac{C}{2}\operatorname{Var}(v_1 \mid q) + \frac{C}{2}\left[\widehat{v}_1(q) + \beta C\widehat{v}_2(q) + \beta g\right]^2 + \frac{\beta A}{2}g^2, \tag{A.37}$$

where $Z(q) = \beta C\mathbb{E}[v_2^2 \mid q]/2$, with $\mathbb{E}[Z(q)] = Z$, has an ex-ante expectation independent of the communication strategy. With $M(q) = \widehat{v}_1(q) + \beta C\widehat{v}_2(q)$ and $B = 1/(A + \beta C)$, the first-order condition is $C[M(q) + \beta g] + Ag = 0$. Thus $g^R(q) = -BCM(q)$ and $M + \beta g^R = BAM$, giving the minimized mean-loss term $\frac{1}{2}ABCM(q)^2$ in (18).

Step 2: the law of total variance and the affine form. For messages generated from v_1 alone, $\widehat{v}_2(q) = \eta\widehat{v}_1(q)$, and the law of total variance gives $\mathbb{E}[\operatorname{Var}(v_1 \mid q)] = \operatorname{Var}(v_1) - \operatorname{Var}(\widehat{v}_1(q))$. Taking the unconditional expectation of the conditional loss at the optimal promise therefore leaves

$$\mathcal{L}(q) = Z + \frac{C}{2}\operatorname{Var}(v_1) + \frac{C}{2}\Gamma(\eta)\operatorname{Var}\big(\widehat{v}_1(q)\big), \qquad \Gamma(\eta) \equiv AB(1 + \beta C\eta)^2 - 1, \tag{A.38}$$

which is (B.14). The communication strategy enters through one number, the variance of the posterior mean, and its price is $\frac{C}{2}\Gamma(\eta)$.

Step 3: the threshold. Because $\mathcal{L}(q)$ is affine in $\operatorname{Var}(\widehat{v}_1(q))$, the optimum is at one end of the range of that variance. If $\Gamma(\eta) < 0$, the optimum maximizes posterior-mean variance, attaining $\operatorname{Var}(v_1)$ exactly when v_1 is measurable with respect to the message. If $\Gamma(\eta) > 0$, the optimum has zero posterior-mean variance, hence $\widehat{v}_1(q) = 0$ almost surely under the mean-zero normalization. At equality every admissible communication strategy yields the same loss. For $\eta \geq 0$, Γ is strictly increasing, and solving $\Gamma(\eta) = 0$ gives (B.15). The

covariance statement follows from $\eta = \mathrm{Cov}(v_1, v_2)/\mathrm{Var}(v_1)$.

Full-revelation allocation. Under full revelation, write $c_\eta = 1 + \beta C\eta$. Then $g^R(v_1) = -BCc_\eta v_1$ and

$$\theta_1^R(v_1) = ABCc_\eta v_1, \tag{A.39}$$

$$x_1^R(v_1) = -\frac{\kappa}{\kappa^2 + \lambda} ABc_\eta v_1, \tag{A.40}$$

$$\theta_2^R(v_2, v_1) = Cv_2 + g^R(v_1). \tag{A.41}$$

Here the second argument of θ_2^R is the revealed message $q = v_1$. Expected future inflation is

$$e_1^R(v_1) = \left[C\eta - BC(1 + \beta C\eta)\right] v_1. \tag{A.42}$$

An uninformative message sets the posterior means and the optimal promise to zero. Subtracting the full-revelation loss from the loss under the uninformative message gives (B.17).

A.7 Information Boundaries and the Size of the Revelation Region

The following two corollaries record the boundary cases and the structural bound cited in Appendix B.3.

Corollary 5 (Pure current information and pure future forecast)**.** *Suppose a mean-zero scalar signal v enters the current Phillips curve as $\zeta_1 v$ and predicts the future disturbance through $\mathbb{E}[v_2 \mid v] = \zeta_2 v$. For messages generated from v, the coefficient on $\mathrm{Var}(\mathbb{E}[v \mid q])$ in optimized loss is $C\Gamma(\zeta_1, \zeta_2)/2$, where*

$$\Gamma(\zeta_1, \zeta_2) = AB(\zeta_1 + \beta C\zeta_2)^2 - \zeta_1^2, \tag{A.43}$$

so that $\Gamma(\eta) = \Gamma(1, \eta)$. For a nondegenerate signal, full revelation is optimal when this coefficient is negative and policy-relevant opacity when it is positive; at zero all signals are payoff equivalent. Pure current information ($\zeta_1 \neq 0$, $\zeta_2 = 0$) is revealed, whereas pure future forecast ($\zeta_1 = 0$, $\zeta_2 \neq 0$) is kept private.

Proof. Replacing the current-state posterior mean in (A.37) by $\zeta_1 \mathbb{E}[v \mid q]$ gives, up to signal-independent terms,

$$\frac{C}{2}\zeta_1^2 \mathbb{E}[\mathrm{Var}(v \mid q)] + \frac{1}{2} ABC(\zeta_1 + \beta C\zeta_2)^2 \mathbb{E}[\mathbb{E}[v \mid q]^2]. \tag{A.44}$$

Total variance gives the stated coefficient. It equals $-\beta\frac{1}{2}BC^2\zeta_1^2$ for pure current information and $\frac{1}{2}ABC(\beta C)^2\zeta_2^2$ for pure future forecast. When $\zeta_1 \neq 0$, the normalization $v_1 = \zeta_1 v$ gives the effective forecast loading ζ_2/ζ_1; the threshold $\eta^\star$ of Lemma 15 applies when that ratio is nonnegative.

Corollary 6 (Structurally small revelation region)**.** *The threshold $\eta^\star$ of Lemma 15 satisfies*

$$\eta^\star = \frac{1/\sqrt{AB} - 1}{\beta C} = \frac{\sqrt{1 + \beta C/A} - 1}{\beta C} \leq \frac{1}{2A} = \frac{1}{2(1 + \lambda/\kappa^2)}. \tag{A.45}$$

The threshold is unchanged when the central bank observes the current state with independent noise, $v_1^o = v_1 + \varepsilon^o$. If instead the shock itself has a persistent component, $v_1 = p_1 + n_1$ with n_1 independent of (p_1, v_2), $\mathbb{E}[v_2 \mid p_1] = \rho p_1$, and (5) holding, then $\eta = \rho \operatorname{Var}(p_1)/\operatorname{Var}(v_1)$, so for $\rho > 0$ revelation requires $\operatorname{Var}(p_1)/\operatorname{Var}(v_1) < \eta^\star/\rho$.

Proof. The inequality follows from $\sqrt{1+x} - 1 \leq x/2$ for $x \geq 0$. For a noisy observation of the current state, $v_1^o = v_1 + \varepsilon^o$, with ε^o independent of the disturbances, put $m(v_1^o) = \mathbb{E}[v_1 \mid v_1^o]$. Then $\mathbb{E}[v_2 \mid v_1^o] = \eta m(v_1^o)$, and a message generated from v_1^o satisfies $\widehat{v}_2(q) = \eta\mathbb{E}[m(v_1^o) \mid q]$. Optimizing current policy conditional on v_1^o leaves a residual-uncertainty term whose ex-ante expectation is independent of communication. The calculation in the proof of Corollary 5, with $\zeta_1 = 1$ and $\zeta_2 = \eta$, applies to $m(v_1^o)$ with the same coefficient $C\Gamma(\eta)/2$ and hence the same threshold. Finally, $\operatorname{Cov}(v_1, v_2) = \rho \operatorname{Var}(p_1)$ and $\operatorname{Var}(v_1) = \operatorname{Var}(p_1) + \operatorname{Var}(n_1)$ give the component formula.

A.8 Ideal Rule under Delegation: Statements

Section 4.1 states in words that, under the ideal rule, the delegated central bank keeps its forecast private on its own and that society gains nothing from an audit that conditions on the realized shock. The two results are the following; their proofs are in Appendix A.9 and Appendix A.10.

Lemma 6 (Delegated selective opacity)**.** *Suppose the standing rule is Cv_2, the central bank observes both components of private information in (6), and the audit depends only on (g, z). Let $\mathcal{J}(v_1, s, g)$ be the central bank's value function after optimizing current policy, as defined in (27), and suppose optimal delivery choices exist. Then:*

(i) Separability.

$$\mathcal{J}(v_1, s, g) = \mathcal{I}(v_1, g) + Z(v_1, s), \tag{A.46}$$

where $\mathcal{I}$ does not depend on the pure forecast and Z does not depend on the announcement; up to the suppressed control-noise constant, $Z(v_1, s) = \frac{\beta C}{2}\mathbb{E}[v_2^2 \mid v_1, s]$.

(ii) Common best responses. *All types with the same nowcast v_1 rank the available announcements identically and have the same set of optimal announcements.*

(iii) Opacity. *In every equilibrium with unique announcement best responses, the announcement is a function of v_1 alone and conveys no additional information about the pure forecast s.*

Proof. See Appendix A.9.

Part (ii) holds without any curvature assumption; part (iii) needs uniqueness, because several announcements tying for the minimum could be selected differently by central banks with different forecasts. For the ideal-rule case of the coarse-rule audit of Appendix A.17, Appendix A.18 proves this uniqueness directly, without requiring the traction function to be affine. The economics of part (i) is that under the ideal rule, once (g, a) is fixed, neither the deviation from the ex-post ideal, $g + a$, nor the audit residual, $a + \varepsilon_2$, depends on the shock; the term $Z(v_1, s)$ depends on the forecast and on nothing the central bank chooses.

The lemma takes the audit class as given. The mandate could instead condition the sanction on the realized shock, (24), and such an audit could make the central bank's continuation cost vary with the shock and hence with its forecast; the expected sanction $\mathcal{A}$ is the one channel through which s could re-enter the announcement. The next result shows that society, as of period 0, has no use for this freedom.

Lemma 7 (Sufficiency of shock-independent accountability)**.** *Let the standing rule be Cv_2. Write $\mathcal{L}^R(q)$ for the minimized Ramsey loss (18) of the communication strategy q and $\mathcal{L}(\text{mandate})$ for the social loss in an equilibrium of a mandate. Consider mandate–equilibrium pairs whose public announcement information can be represented as a nowcast message $q_1(v_1)$ together with a separate forecast signal $q_f(s)$, and let $q_1^\star$ minimize the Ramsey loss among nowcast messages. Assume that, when revelation of the nowcast is strictly preferred, the Ramsey allocation of $q_1^\star$ is implementable by the ideal-rule case of the coarse-rule audit of Appendix A.17: the announcement format admits exactly the image of the candidate schedule $g_C(r) = \Delta^a(g^R(r))$, every target $g^R(r)$ lies within the ceiling of Lemma 3 on a globally optimal branch, and the sanctioning capacity covers the delivery stake and the announcement correction. Then an optimal mandate in this class can be chosen to depend only on (g, z), and every admissible equilibrium with announcement information q satisfies*

$$\mathcal{L}(\text{mandate}) - \mathcal{L}^R(q_1^\star) \geq \tfrac{1}{2} ABC(\beta C)^2 \, \mathbb{E}\big[\mathbb{E}[s \mid q]^2\big]. \tag{A.47}$$

In particular, every optimal equilibrium leaves the posterior mean of the pure forecast unchanged.

Proof. See Appendix A.10.

The result compares each mandate with the planner's optimum for the information its equilibrium reveals. Within the stated class, a separate forecast message raises that Ramsey loss, while an audit depending only on (g, z) attains the best nowcast-only allocation under the lemma's implementation conditions. Separate shock contingency therefore cannot improve welfare in this class. The proof below gives the comparison; Remark 4 establishes the stronger pure-forecast conclusion for any positive delivery stake.

A.9 Delegated Selective Opacity

Proof of Lemma 6. Part (i): separability. Fix a true type (v_1, s), an announcement g, and the expectation $e_1(g)$ that the public attaches to it through an arbitrary map. The central bank's expected loss at this announcement is $\mathcal{J}(v_1, s, g)$ in (27).

In period 1 the central bank chooses x_1 knowing v_1 and $e_1(g)$. The period-1 Phillips curve (7) contains the true disturbance v_1 and the expectation $e_1(g)$ but not s, and the period-1 loss in (3) is $\frac{1}{2}(\theta_1^2 + \lambda x_1^2)$. Minimizing over x_1 gives, as in Lemma 1(ii), the minimized current loss

$$\frac{C}{2}\big[v_1 + \beta e_1(g)\big]^2, \tag{A.48}$$

which is a function of (v_1, g) only.

In period 2 the central bank observes the disturbance $v_2 = \eta v_1 + s + \varepsilon$ and chooses x_2, equivalently the delivered center θ_2 through (8). By the identity (11), the period-2 loss is

$$\frac{A}{2}\big(\theta_2 - C v_2\big)^2 + \frac{C}{2} v_2^2. \tag{A.49}$$

Under the ideal rule the announced center is $Cv_2 + g$, so writing $a \equiv \theta_2 - Cv_2 - g$ for the shortfall, the delivered deviation from the ex-post ideal is $\theta_2 - Cv_2 = g + a$ and the first term of (A.49) is $\frac{A}{2}(g + a)^2$. The second term depends on v_2 only: no period-2 choice and no announcement affects it.

The audit adds the expected sanction (25). Its probability $\rho(g, z)$ depends on the announcement and on the miss $z = \pi_2 - Cv_2 - g$. Substituting $\pi_2 = \theta_2 + \varepsilon_2$ and the definition of a,

$$z = a + \varepsilon_2, \tag{A.50}$$

so the residual is the shortfall plus implementation noise. Because ε_2 is independent of (v_1, s, v_2), the expected sanction conditional on the central bank's choices is a function of (g, a) alone, which we write as $\mathcal{A}(g, a) \equiv \Psi\,\mathbb{E}[\rho(g, a + \varepsilon_2)]$. It does not depend on v_2 or on s.

The central bank's period-2 problem in (27) is therefore, after dropping the term $\frac{C}{2} v_2^2$

that no choice affects,

$$\mathcal{V}(g) \equiv \min_{a} \left\{ \frac{A}{2}(g+a)^2 + \mathcal{A}(g,a) \right\}, \tag{A.51}$$

whose objective involves only (g, a). Its value, which exists by the assumption that optimal delivery choices exist, is a function of g alone.

Collecting (A.48), (A.51), and the dropped term, and taking the expectation over v_2 conditional on the type,

$$\mathcal{J}(v_1, s, g) = \underbrace{\frac{C}{2}\big[v_1 + \beta e_1(g)\big]^2 + \beta\, \mathcal{V}(g)}_{\mathcal{I}(v_1, g)} + \underbrace{\frac{\beta C}{2}\mathbb{E}\big[v_2^2 \mid v_1, s\big]}_{Z(v_1, s)}, \tag{A.52}$$

which is (A.46). The first bracket does not involve s and the second does not involve g, as claimed; the implementation-noise constant from integrating π_t has been suppressed throughout.

Parts (ii) and (iii): common best responses and opacity. Fix the equilibrium expectation map $e_1(\cdot)$ and the audit. By part (i), for any two announcements g and g',

$$\mathcal{J}(v_1, s, g) - \mathcal{J}(v_1, s, g') = \mathcal{I}(v_1, g) - \mathcal{I}(v_1, g'). \tag{A.53}$$

The right-hand side does not depend on s. Thus any two types (v_1, s) and (v_1, s') with the same nowcast rank every pair of announcements in exactly the same way. In particular, their sets of minimizing announcements coincide.

If this common best-response set is a singleton, every type sharing v_1 must choose the same announcement. The equilibrium announcement is then a function of v_1 alone, so conditional on v_1 it conveys no additional information about s.

Part (ii) therefore holds unconditionally, and part (iii) under the stated hypothesis that announcement best responses are unique: separability does not by itself exclude several announcements attaining the same minimum of $\mathcal{I}(v_1, \cdot)$, and forecast types breaking such a tie differently would reveal s.

A.10 Shock-Independent Audit

Proof of Lemma 7. We first derive a lower bound on loss within the stated class of mandates, then show that an audit depending only on (g, z) attains it.

Step 1: a lower bound on any mandate. Fix an admissible mandate, with audit $\rho(g, z, v_2)$, and an equilibrium of the delegated game: an announcement strategy $g(\mathcal{T})$, delivery policies, and beliefs. Write q for the public information carried by the announcement. The

planner of Section 3, given the same communication strategy q, may choose any square-integrable current policy and delivery schedule consistent with the expectation condition (9); the equilibrium's policies are such a choice, and the equilibrium expectation equals the expected delivered center by the definition of equilibrium. Since sanctions carry no social cost, the social loss in the equilibrium is at least the minimized Ramsey loss of q given by Lemma 1(ii),

$$\mathcal{L}(\text{mandate}) \geq \mathcal{L}^R(q). \tag{A.54}$$

This is a relaxation argument; it does not require preserving reporting incentives across mechanisms and remains valid when the audit depends on v_2.

Step 2: separate information and the competitor's own nowcast message. Within the maintained class, $q = (q_1, q_f)$ with $q_1 = q_1(v_1)$ and $q_f = q_f(s)$. Corollary 2, applied to this q_1, gives

$$\mathcal{L}^R(q) = \mathcal{L}^R(q_1) + \tfrac{1}{2}ABC(\beta C)^2\mathbb{E}\big[\mathbb{E}[s \mid q]^2\big], \tag{A.55}$$

and optimality of $q_1^\star$ among nowcast messages gives $\mathcal{L}^R(q_1) \geq \mathcal{L}^R(q_1^\star)$. Hence

$$\mathcal{L}^R(q) \geq \mathcal{L}^R(q_1^\star) + \tfrac{1}{2}ABC(\beta C)^2\mathbb{E}\big[\mathbb{E}[s \mid q]^2\big]. \tag{A.56}$$

This covers every admissible nowcast message, including partial disclosure and the threshold case, without presuming what an optimal competitor reveals.

Step 3: a (g, z) audit attains the benchmark. If revelation of the nowcast is strictly preferred, $\Gamma(\eta) < 0$, take the ideal-rule case of the coarse-rule audit of Appendix A.17 with the candidate $g_C(r) = \Delta^a(g^R(r))$. Under the implementation hypothesis of the lemma, the ideal-rule argument in Appendix A.18 gives the reporting gap $\frac{\beta BC^2(1-A\eta)}{2}(r - v_1)^2$, strictly positive because $\eta < \eta^\star < 1/A$ (Corollary 6), so truthful nowcast reporting is the unique optimum. By Lemma 6, whose hypothesis this audit satisfies, the pure forecast adds the same constant to every report and does not affect the announcement; and the delivered center is the Ramsey center. The induced information is $q_1^\star$ and the loss is $\mathcal{L}^R(q_1^\star)$. If opacity is optimal, $\Gamma(\eta) \geq 0$, a constant announcement, say $g = 0$, with delivery Cv_2 attains the pooling benchmark: $e_1 = 0$, $\theta_1 = Cv_1$, $\theta_2 = Cv_2$. A constant announcement format with zero audit probability suffices, and under the graduated parameterization a zero promised departure has zero delivered deviation as its unique delivery optimum. This pooling equilibrium attains $\mathcal{L}^R(q_1^\star)$. In both cases the audit depends only on (g, z).

Step 4: contradiction. Combining (A.54) and (A.56) gives (A.47) for every admissible equilibrium. Suppose an optimal equilibrium had $\mathbb{E}[\mathbb{E}[s \mid q]^2] > 0$. Its loss would exceed $\mathcal{L}^R(q_1^\star)$, which Step 3 shows an admissible (g, z) audit attains; it would not be optimal.

Hence every optimal equilibrium has $\mathbb{E}[s \mid q] = 0$ almost surely, and the same bound shows that no admissible audit, shock-contingent or not, attains less than $\mathcal{L}^R(q_1^\star)$. This proves sufficiency of a shock-independent audit. The optimal audit's functional form may be non-unique, and the argument covers the maintained class of communication strategies.

The pure-forecast case of Lemma 7 takes the following form.

Remark 4 (Pure-forecast benchmark under delegation). Let $v_1 = 0$ almost surely. Under the ideal rule and the graduated audit $\rho_C(z)$ of Section 4.1.2 with any positive stake, the constant announcement $g = 0$ is the unique optimal announcement of every forecast type in the equilibrium in which the public expects delivery, whatever expectation is attached to announcements off the path.

To see this, write $\mathcal{V}_C(g) = \min_a\{\frac{A}{2}(g+a)^2 + \mathcal{A}_C(a)\}$, the delivery value (31) at $\Delta^a = g$. The expected audit burden $\mathcal{A}_C$ has its unique minimum at $a = 0$, so $\mathcal{V}_C(g) \geq \mathcal{A}_C(0)$ with equality only when $a = 0$ and $g + a = 0$, that is, at $g = 0$. Under $g = 0$ delivery is Cv_2 and the on-path expectation is zero; at any other announcement, with whatever finite expectation $e_1(g)$ the public attaches to it,

$$\mathcal{J}(0, s, g) - \mathcal{J}(0, s, 0) = \frac{C\beta^2}{2} e_1(g)^2 + \beta\big[\mathcal{V}_C(g) - \mathcal{V}_C(0)\big] > 0, \tag{A.57}$$

an inequality independent of s and of the off-path belief. This is opacity without any restriction on the announcement and without any curvature assumption on the announcement problem.

Two limits of Lemma 7 should be noted. First, it establishes that shock contingency is unnecessary; an optimal audit may still contain it. Audits are identified only up to the incentives they induce: two audits that induce the same relative valuation of announcements and deliveries are isomorphic for the allocation, and a term that depends only on the realized shock can be added to any audit without changing those valuations. Second, the bound is an inequality: an arbitrary mandate may also lose through inefficient current policy or delivery, and an announcement may vary with s without moving its posterior mean. What the lemma determines is the object that matters for the design: in every optimal design the announcement carries no posterior-mean information about the pure forecast.

A.11 Delivery and Optimal Audit Tolerance

From here on we restrict ρ of (24) to one parametric family, the *graduated parameterization*

$$\rho_C(z) = 1 - \exp\left(-\frac{z^2}{2\tau^2}\right), \tag{A.58}$$

where the subscript marks the component that disciplines *delivery*: a second component, ρ_I, will discipline the choice of announcement in Section 4.2.6, and the two are combined there. The probability is smooth and symmetric, and the expected sanction is bounded by Ψ. These properties provide a differentiable delivery problem without favoring overshooting or undershooting. The formulas below use the graduated specification. Choosing τ trades off the width of its responsive range against the curvature of the expected sanction at zero shortfall.

Integrating over the implementation noise gives the expected delivery sanction as a function of the hidden shortfall,

$$\mathcal{A}_C(a) = \Psi\left[1 - \frac{\tau}{\sqrt{\tau^2 + \sigma_\pi^2}} \exp\left(-\frac{a^2}{2(\tau^2 + \sigma_\pi^2)}\right)\right], \tag{A.59}$$

which is smooth, symmetric, and increasing in $|a|$. The magnitude of the marginal disciplinary force peaks at $|a| = \sqrt{\tau^2 + \sigma_\pi^2}$ and declines beyond, which makes the interval of half-width $\sqrt{\tau^2 + \sigma_\pi^2}$ around zero the *responsive range* of the audit. The tolerance therefore sets the scale on which the audit disciplines at all.

Proof of Lemma 3. Step 1: the first-order condition on the responsive range. Differentiating (A.59) gives

$$\mathcal{A}_C'(a) = \frac{\Psi\tau}{(\tau^2 + \sigma_\pi^2)^{3/2}} a \exp\left(-\frac{a^2}{2(\tau^2 + \sigma_\pi^2)}\right). \tag{A.60}$$

On $[-\sqrt{\tau^2 + \sigma_\pi^2}, \sqrt{\tau^2 + \sigma_\pi^2}]$,

$$\mathcal{A}_C''(a) = \frac{\Psi\tau}{(\tau^2 + \sigma_\pi^2)^{3/2}} \exp\left(-\frac{a^2}{2(\tau^2 + \sigma_\pi^2)}\right)\left(1 - \frac{a^2}{(\tau^2 + \sigma_\pi^2)}\right) \geq 0. \tag{A.61}$$

The maximum absolute derivative on this responsive range is attained at $|a| = \sqrt{\tau^2 + \sigma_\pi^2}$ and equals $\Psi\tau/(\exp(1/2)(\tau^2 + \sigma_\pi^2))$. Because $\Psi > 0$, (34) has a unique solution on this range, equal to zero at $\Delta^\star = 0$ and otherwise of the opposite sign, if and only if $A|\Delta^\star| \leq \Psi\tau/(\exp(1/2)(\tau^2 + \sigma_\pi^2))$, which is $|\Delta^\star| \leq \bar{\Delta}_C$.

Step 2: the delivery optimum. Set $\Delta^a = \Delta^\star - a^*$, where a^* is the shortfall solving (34). Then

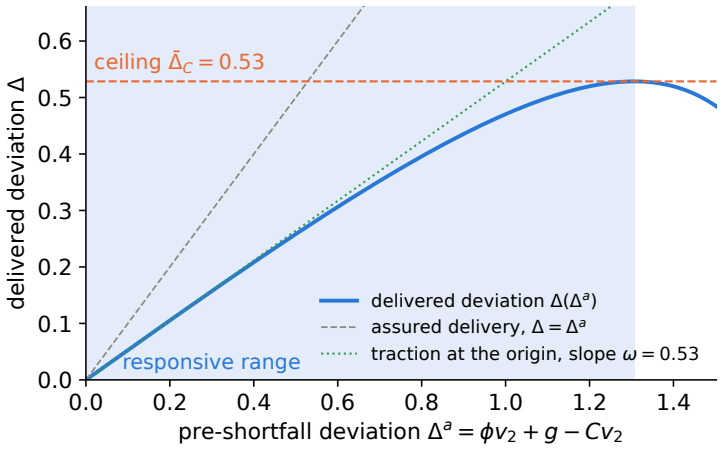

Figure A.2: Traction function at the baseline audit

Notes: Graduated audit (A.58), $\Psi = 18$ and $\tau = \sigma_\pi = 0.55$; baseline calibration. The curve maps the pre-shortfall deviation Δ^a into delivery $\Delta(\Delta^a) = \Delta^a + a^*(\Delta^a)$. Its slope at zero is $\omega = 0.528$, and its maximum is $\bar{\Delta}_C = 0.53$ at the edge of the responsive range (shaded). The horizontal distance to the 45-degree line measures optimal hyperbole. Delivery falls beyond the responsive range as the marginal sanction declines.

$\Delta^a + a^* = \Delta^\star$, and the delivery first-order condition for (31) is exactly (34). The objective is strictly convex on the responsive range, so a^* is the unique delivery optimum there; Step 4 shows that under the lemma's hypothesis it is the unique optimum outright.

Step 3: monotonicity. The map $\Delta^a(\Delta^\star)$ of (35), the inverse of the traction function on the responsive range, which we call the inverse traction function, is one-to-one. On the positive branch, write $y = -a \in [0, \sqrt{\tau^2 + \sigma_\pi^2}]$ and $\kappa_C \equiv \mathcal{A}_C''(0) = \Psi\tau/(\tau^2 + \sigma_\pi^2)^{3/2}$, the curvature of the expected sanction at zero shortfall. The first-order condition can be parameterized as

$$\Delta^\star(y) = \frac{\kappa_C}{A} y \exp\left(-\frac{y^2}{2(\tau^2 + \sigma_\pi^2)}\right), \qquad \Delta^a(y) = y + \Delta^\star(y). \tag{A.62}$$

For $y \in [0, \sqrt{\tau^2 + \sigma_\pi^2})$,

$$\frac{\mathrm{d}\Delta^\star}{\mathrm{d}y} = \frac{\kappa_C}{A} \exp\left(-\frac{y^2}{2(\tau^2 + \sigma_\pi^2)}\right)\left(1 - \frac{y^2}{(\tau^2 + \sigma_\pi^2)}\right) > 0, \tag{A.63}$$

so $\mathrm{d}\Delta^a/\mathrm{d}y = 1 + \mathrm{d}\Delta^\star/\mathrm{d}y > 0$. Writing $\Delta^a(y)$ for the value of (35) at $\Delta^\star(y)$, the traction function has slope $(\mathrm{d}\Delta^\star/\mathrm{d}y)/(1 + \mathrm{d}\Delta^\star/\mathrm{d}y) \in (0, 1)$. Continuity is immediate, and symmetry of the objective makes the map odd. Thus Δ^a is continuous, odd, and strictly increasing on the responsive range. Figure A.2 plots the traction function and its ceiling at the baseline audit.

Step 4: global optimality. Write $w = a^2/(\tau^2 + \sigma_\pi^2)$. Since $\mathcal{A}_C''(a) = \kappa_C \exp(-w/2)(1-w)$ and $\exp(-w/2)(1-w)$ reaches its minimum $-2\exp(-3/2)$ at $w = 3$, the negative curvature of the expected sanction is bounded in magnitude by $2\exp(-3/2)\kappa_C$. Hence the delivery objective satisfies

$$\frac{\mathrm{d}^2}{\mathrm{d}a^2}\left\{\frac{A}{2}(\Delta^a + a)^2 + \mathcal{A}_C(a)\right\} \geq A - 2\exp(-3/2)\kappa_C. \tag{A.64}$$

The hypothesis (32) of Lemma 3 therefore makes the objective globally strictly convex. Its quadratic term is coercive and the audit burden is bounded, so a global minimizer exists and is unique. The target-associated shortfall is then $a^*(\Delta^a(\Delta^\star))$, and $\Delta(\Delta^a(\Delta^\star)) = \Delta^\star$.

The same ceiling calculation gives the tolerance that maximizes the responsive domain.

Corollary 7 (Optimal audit tolerance)**.** *For fixed $\Psi > 0$, $A > 0$, and $\sigma_\pi > 0$, the ceiling (33) is uniquely maximized at*

$$\tau^* = \sigma_\pi. \tag{A.65}$$

At that tolerance,

$$\bar{\Delta}_C^{\max} = \frac{\Psi}{2\exp(1/2)\,A\sigma_\pi}, \qquad \Psi^{\min}(\bar{\Delta}) = 2\exp(1/2)\,A\sigma_\pi\bar{\Delta}. \tag{A.66}$$

The second expression is the smallest stake meeting the responsive-branch ceiling for targets up to $\bar{\Delta} > 0$.

Proof. Maximizing the ceiling amounts to maximizing $\tau/(\tau^2+\sigma_\pi^2)$. Its derivative has the sign of $\sigma_\pi^2 - \tau^2$, giving the unique maximizer and the two expressions above. Locally, $\mathcal{A}_C(a) = \mathcal{A}_C(0) + \kappa_C a^2/2 + o(a^2)$, so the delivery first-order condition gives $\Delta(\Delta^a) = \kappa_C\Delta^a/(A+\kappa_C) + o(\Delta^a)$, the traction ω of Section 4.1.2.

A.12 Point-Forecast and Mean-Delivery Calibration

This calculation uses the forecast-bundling nowcast case ($s = 0$ almost surely) and truthful reporting $r = v_1$, so $\varepsilon = v_2 - \eta v_1 = v_2 - \eta r$ is unknown at announcement and has conditional mean zero. A literal announcement cannot react to its realization. All uses of the inverse traction function assume the target lies on a responsive branch verified to be globally optimal, for example under (32). A convenient point-forecast calibration is

$$g_\phi^{\text{point}}(r) = \Delta^a(g^R(r)) - (\phi - C)\eta r. \tag{A.67}$$

It produces pre-shortfall deviation

$$\Delta^a = \Delta^a(g^R(r)) + (\phi - C)\varepsilon, \tag{A.68}$$

and hence delivered center

$$\begin{aligned}\theta_2 = Cv_2 + g^R(r) \\ + \left[\Delta(\Delta^a(g^R(r)) + (\phi - C)\varepsilon) - \Delta(\Delta^a(g^R(r)))\right].\end{aligned} \tag{A.69}$$

At $\varepsilon = 0$, the Ramsey center $Cv_2 + g^R(r)$ is delivered exactly. If both arguments lie on the responsive branch, the bracket is between zero and $(\phi - C)\varepsilon$, since the traction function has slope between zero and one. Under the ideal rule the bracket vanishes. Under the strict target or a coarse rule it records the partially repaired surprise.

The bracket need not have zero conditional mean: a nonlinear traction function cannot in general be evaluated at the mean shock in place of taking its expectation. To target the conditional Ramsey mean, choose $g_\phi(r)$ from

$$\mathbb{E}\left[\Delta((\phi - C)(\eta r + \varepsilon) + g_\phi(r)) \mid v_1 = r\right] = g^R(r). \tag{A.70}$$

Existence and a one-to-one announcement schedule require that the target lie in the range of the map $g \longmapsto \mathbb{E}[\Delta((\phi - C)(\eta r + \varepsilon) + g) \mid v_1 = r]$ on the branch used. When this equation holds, expectations formed from actual delivery, $e_\phi(r) = \mathbb{E}[Cv_2 + \Delta((\phi - C)v_2 + g_\phi(r)) \mid v_1 = r]$ (Appendix A.17), reduce to $e_\phi(r) = C\eta r + g^R(r)$. The calibration reduces to the point-forecast calibration with no residual uncertainty or with a linear traction function. Matching this mean leaves the state-by-state surprise in place, and the optimal revelation problem under nonlinear auditing is a separate question.

A.13 Revelation under Coarse Rules

Proof of Proposition 1 and Lemma 16. The nowcast (Lemma 16). For the current-state comparison, take $q = q(v_1)$, so (10) applies. Fix a message q with posterior mean $m \equiv \widehat{v}_1(q)$ and posterior variance $w \equiv \mathrm{Var}(v_1 \mid q)$. Under (23) the deviation from the ex-post ideal is $(\phi - C)v_2 + g$, the public expectation is $e_1 = \phi\eta m + g$, and the current allocation is $\theta_1 = C[v_1 + \beta e_1]$, $x_1 = -\kappa[v_1 + \beta e_1]/(\kappa^2 + \lambda)$ as in Lemma 1(ii). Averaging the minimized current loss $\frac{C}{2}[v_1 + \beta e_1]^2$ within the posterior cell and adding discounted period-2 loss

gives, up to terms independent of the communication strategy,

$$\frac{C}{2}w + \frac{C}{2}\big[m(1+\beta\phi\eta)+\beta g\big]^2 + \frac{\beta A}{2}\big[(\phi-C)\eta m + g\big]^2 + \frac{\beta A}{2}(\phi-C)^2\,\mathrm{Var}(v_2 \mid q), \tag{A.71}$$

where we used $\mathbb{E}[v_2 \mid q] = \eta m$. Write $p \equiv m(1+\beta\phi\eta)$ and $k \equiv (\phi - C)\eta m$, so that the middle two terms are $\frac{C}{2}(p+\beta g)^2 + \frac{\beta A}{2}(k+g)^2$. The first-order condition in g gives $g = -B(Cp + Ak)$, which is (B.19); direct substitution verifies $g^R_\phi - g^R = (C-\phi)\eta m$. At the optimum, $p + \beta g = AB(p - \beta k)$ and $k + g = -BC(p - \beta k)$, so the minimized middle terms equal

$$\frac{1}{2}ABC(p-\beta k)^2 = \frac{1}{2}ABCc_\eta^2 m^2, \qquad p - \beta k = m\big[1 + \beta\phi\eta - \beta(\phi - C)\eta\big] = c_\eta m. \tag{A.72}$$

The minimized mean component of message loss therefore does not depend on ϕ: the promise absorbs whatever the rule leaves out of the predictable part of the shock, and $v_1 + \beta e_1$ and the deviation from the ex-post ideal, $k + g = -BCc_\eta m = g^R(m)$, are the same as under the ideal rule. The rule slope enters only through the surprise term. By the law of total variance, $\mathbb{E}[w] = \mathrm{Var}(v_1) - \mathrm{Var}(m)$ and $\mathbb{E}[\mathrm{Var}(v_2 \mid q)] = \mathrm{Var}(v_2) - \eta^2\,\mathrm{Var}(m)$. Collecting the coefficient on $\mathrm{Var}(m)$,

$$-\frac{C}{2} + \frac{1}{2}ABCc_\eta^2 - \frac{\beta A}{2}(C-\phi)^2\eta^2 = \frac{C}{2}\left[\Gamma(\eta) - \frac{\beta A}{C}(C-\phi)^2\eta^2\right], \tag{A.73}$$

which is (B.18). Because the loss is affine in $\mathrm{Var}(m)$, the optimum is at an endpoint, which proves the lemma. Setting $\phi = C$ recovers (B.14).

The pure forecast (Proposition 1). First set $v_1 = 0$ to isolate the forecast-only calculation. The posterior-mean statistic is $\widehat{s}$, the public expectation is $\phi\widehat{s} + g$, and there is no current-loss term in s. The per-message loss is

$$\frac{C}{2}\big[\beta(\phi\widehat{s} + g)\big]^2 + \frac{\beta A}{2}\big[(\phi - C)\widehat{s} + g\big]^2 + \frac{\beta A}{2}(\phi - C)^2\,\mathrm{Var}(s + \varepsilon \mid q_f). \tag{A.74}$$

With $p \equiv \beta\phi\widehat{s}$ and $k \equiv (\phi - C)\widehat{s}$ the same algebra gives $g = -B(Cp + Ak)$, which is (37), and $p - \beta k = \beta C\widehat{s}$; hence $\phi\widehat{s} + g = ABC\widehat{s}$ and $(\phi - C)\widehat{s} + g = -B\beta C^2\widehat{s}$, and the minimized first two terms equal $\frac{1}{2}ABC(\beta C)^2\widehat{s}^2$. Under opacity $\widehat{s} = 0$. By the law of total variance, revealing the forecast reduces the expected posterior variance in the last term by $\mathrm{Var}(\widehat{s})$. The loss difference is therefore (36). For positive posterior-mean variance it is negative if and only if $(C-\phi)^2 > \beta BC^3$, or $|\phi - C| > \bar{c}$. No step requires $\phi \le C$. At equality, or with zero posterior-mean variance, the planner is indifferent.

For a nonzero current state already revealed to the public, the same calculation applies

to the incremental forecast component. Write $\widehat{s} = \mathbb{E}[s \mid q_f, v_1]$. The total optimal promise is the current-state promise plus the forecast increment, $g^R_\phi(v_1) + g^R_\phi(\widehat{s})$, where $g^R_\phi(v_1) = \gamma^R_\eta v_1$ is the nowcast promise (B.19) and $g^R_\phi(\widehat{s}) = \gamma^R \widehat{s}$ the forecast promise (37). The optimized mean-loss term is $\frac{1}{2}ABC\,\mathbb{E}[(c_\eta v_1 + \beta C\widehat{s})^2]$, and the expected posterior variance of v_2 falls by $\mathbb{E}[\widehat{s}^2]$. Because $\mathbb{E}[\widehat{s} \mid v_1] = 0$, the cross term with v_1 vanishes. Subtracting the no-forecast-message allocation therefore gives exactly (36), with the promise and expectation interpreted as forecast increments. Under full forecast revelation, $\widehat{s} = s$.

The strict target, $\phi = 0$, is the following corollary, which Appendix B.3 uses.

Corollary 8 (Information revelation under the strict target)**.** *Under the strict target, $\phi = 0$:*

(i) the revelation coefficient of the current state is

$$\Gamma_0(\eta) = -B\beta C(A\eta - 1)^2 \leq 0, \tag{A.75}$$

so full revelation of v_1 is optimal for every news content η; if $\mathrm{Var}(v_1) > 0$, revelation and pooling are indifferent only at $\eta = 1/A$;

(ii) conditional on revelation of v_1, full revelation of the pure forecast is optimal, and strictly improves on opacity when $\mathrm{Var}(s) > 0$, because $C > \bar{c}$;

(iii) under revelation of v_1 alone the promise is $g^R_0(v_1) = C\eta v_1 + g^R(v_1)$, which equals the expectation (A.42) that the ideal rule would induce, and vanishes at $\eta = 1/A$; the incremental promise attached to the forecast is $g^R_0(s) = ABCs$, while its contribution to the conditional-mean deviation from the ex-post ideal is $-B\beta C^2 s$.

Proof of Corollary 8. Setting $\phi = 0$ in (B.18) and using $B = 1/(A + \beta C)$ gives (A.75). A direct argument gives the sign without the algebra: under $\phi = 0$ the expectation is $e_1(q) = g(q)$, so a planner who has learned v_1 can set $g \equiv 0$ and reproduce the allocation under the uninformative message q^0 exactly, whence $\mathcal{L}(q) \leq \mathcal{L}(q^0)$ for every communication strategy and $\Gamma_0 \leq 0$. Part (ii) follows from $\bar{c} = C\sqrt{B\beta C} < C$, since $B\beta C < 1$. For part (iii), (B.19) at $\phi = 0$ gives $g^R_0(v_1) = C\eta v_1 + g^R(v_1)$, which is the bracket in (A.42); it vanishes when $\eta = B(1 + \beta C\eta)$, that is, when $\eta A = 1$, or $\eta = 1/A$. The forecast promise and deviation are (37) and the deviation of Proposition 1 at $\phi = 0$.

A.14 Reporting Distortion and Global Optimality

Proof of Proposition 2 and Lemma 17. The pure forecast (Proposition 2). Hold the current state public, $v_1 = 0$, and let the public read the announcement indexed by r as the forecast r,

attaching to it the promise $\gamma^R r$ and the expectation $\phi r + \gamma^R r = ABCr$ of Proposition 1. The reporting loss (39) is quadratic in r. Its derivative is

$$\partial_r J_\phi(s,r) = \beta^2 C(ABC)^2\, r + \beta A\gamma^R\big[(\phi - C)s + \gamma^R r\big], \tag{A.76}$$

and at $r = s$ the second bracket is $-B\beta C^2 s$ by Proposition 1, so that

$$\partial_r J_\phi(s,s) = AB(\beta C)^2\big[ABC - \gamma^R\big]s = \phi\, AB(\beta C)^2\, s, \tag{A.77}$$

using $ABC - \gamma^R = \phi$ from (37); this is (40): a unit reduction in the report lowers the loss by $\phi\, AB(\beta C)^2 s$, so at the truthful report the central bank prefers a smaller report. The correction (41) has derivative $-\phi\, AB(\beta C)^2 r$, which at $r = s$ cancels the distortion, so $r = s$ is a stationary point of the corrected loss for every forecast. The corrected loss is quadratic in r with second derivative

$$\beta^2 C(ABC)^2 + \beta A(\gamma^R)^2 - \phi\, AB(\beta C)^2 = AB(\beta C)^2\gamma^R + \beta A(\gamma^R)^2 = \beta A\gamma^R\big(\beta BC^2 + \gamma^R\big), \tag{A.78}$$

which, using $ABC = C - \beta BC^2$, equals

$$\beta A(ABC - \phi)(C - \phi). \tag{A.79}$$

Since $\bar{c} = C\sqrt{B\beta C} > \beta BC^2$, both factors are positive when $\phi < C - \bar{c}$. Both are negative when $\phi > C$, so the curvature is also positive throughout the upper revelation region $\phi > C + \bar{c}$. Thus $r = s$ is the unique global minimizer in both revelation regions. On (ABC, C) the curvature is negative, but that interval lies strictly inside the opacity region and is excluded from the proposition. At $\phi = 0$ the distortion is zero for every s and the uncorrected loss is itself strictly convex with minimizer $r = s$, so the central bank announces the planner's promise with no correction.

The nowcast (Lemma 17). Set $s = 0$ almost surely, so the central bank's private information is v_1 alone. Fix a rule slope ϕ and the candidate policy that assigns the rule-constrained Ramsey promise (B.20) to each nowcast. We first derive the reporting distortion and then establish global optimality under the correction.

By (B.20) the candidate promise has slope $\gamma^R_\eta \equiv -B\big(C + [A(\phi - C) + \beta C\phi]\eta\big)$, so that $g^R_\phi(r) = \gamma^R_\eta r$ and $e^R_\phi(r) = (\phi\eta + \gamma^R_\eta)r$. Substituting into (B.21),

$$J_\phi(v_1, r) = \frac{C}{2}\big[v_1 + \beta(\phi\eta + \gamma^R_\eta)r\big]^2 + \frac{\beta A}{2}\big[(\phi - C)\eta v_1 + \gamma^R_\eta r\big]^2. \tag{A.80}$$

The report enters the first bracket through the expectation, with slope $\phi\eta + \gamma_\eta^R$, and the second bracket through the promise alone, with slope γ_η^R; the rule's unrepaired response $(\phi - C)\eta v_1$ depends on the true nowcast, not on the report.

We first evaluate the two brackets at the truthful report. The proof of Lemma 16 in Appendix A.13 shows that at the optimal promise, for every ϕ,

$$v_1 + \beta e_\phi^R(v_1) = ABc_\eta v_1, \qquad (\phi - C)\eta v_1 + g_\phi^R(v_1) = g^R(v_1) = -BCc_\eta v_1, \tag{A.81}$$

and that the planner's first-order condition for the promise is

$$C\left[v_1 + \beta e_\phi^R(v_1)\right] + A\, g^R(v_1) = 0. \tag{A.82}$$

Differentiating (A.80) with respect to the report and evaluating at $r = v_1$ using (A.81),

$$\partial_r J_\phi(v_1, v_1) = \beta C\left[v_1 + \beta e_\phi^R(v_1)\right]\left(\phi\eta + \gamma_\eta^R\right) + \beta A\, g^R(v_1)\, \gamma_\eta^R. \tag{A.83}$$

Collecting the terms in γ_η^R gives $\beta\gamma_\eta^R\{C[v_1 + \beta e_\phi^R(v_1)] + Ag^R(v_1)\}$, which is zero by (A.82): every effect of the report that operates through the promise cancels, because the planner has already balanced the promise's cost tomorrow against its benefit today. What remains is the effect through the rule's anticipated response,

$$\partial_r J_\phi(v_1, v_1) = \beta C\left[v_1 + \beta e_\phi^R(v_1)\right]\phi\eta = ABC\beta\eta\phi\, c_\eta\, v_1, \tag{A.84}$$

where the second equality substitutes (A.81). This is (B.22).

The correction (B.23) has derivative $\partial_r \mathcal{A}_R^\phi(r) = -ABC\beta\eta\phi c_\eta r$, which at $r = v_1$ is the negative of the distortion, so $r = v_1$ is a stationary point of $J_\phi(v_1, r) + \mathcal{A}_R^\phi(r)$ for every type. The corrected objective is quadratic in r with second derivative

$$C\beta^2(\phi\eta + \gamma_\eta^R)^2 + \beta A(\gamma_\eta^R)^2 - ABC\beta\eta\phi c_\eta = \beta B\, \mathcal{K}_\phi(\eta), \tag{A.85}$$

where

$$\mathcal{K}_\phi(\eta) \equiv A^2\eta^2(C - \phi)^2 + AC\beta\eta^2\phi(\phi - C) + AC\eta\phi - 2AC^2\eta + C^2; \tag{A.86}$$

expanding the squares and using $B = 1/(A + \beta C)$ verifies the identity. Since the corrected objective is quadratic and its derivative vanishes at $r = v_1$, we can write the global loss

difference directly as

$$\begin{aligned}\left[J_\phi(v_1, r) + \mathcal{A}_R^\phi(r)\right] - \left[J_\phi(v_1, v_1) + \mathcal{A}_R^\phi(v_1)\right] \\ = \frac{\beta B \mathcal{K}_\phi(\eta)}{2}(r - v_1)^2.\end{aligned} \tag{A.87}$$

Thus $\mathcal{K}_\phi(\eta) > 0$ is sufficient for truth to be the unique global minimizer.

In the strict revelation region this condition is automatic. Using the definitions of Γ_ϕ, $\mathcal{K}_\phi$, and c_η and collecting terms gives

$$\beta B \mathcal{K}_\phi(\eta) = -C\Gamma_\phi(\eta) + ABC\beta\eta\phi c_\eta. \tag{A.88}$$

The second term on the right-hand side is nonnegative because $\eta \geq 0$ and $\phi \geq 0$. Therefore $\Gamma_\phi(\eta) < 0$ implies $\mathcal{K}_\phi(\eta) > 0$. Whenever society strictly prefers to reveal the nowcast under the coarse rule, the correction that removes the local reporting distortion also makes truth the unique global minimizer. Outside the strict revelation region the sign of $\mathcal{K}_\phi$ must still be checked.

For part (i), at $\phi = 0$ the slope is $\gamma_\eta^R = BC(A\eta - 1)$, the distortion (B.22) vanishes for every η, and (A.86) reduces to $\mathcal{K}_0(\eta) = C^2(A\eta - 1)^2 \geq 0$; it is zero only at $\eta = 1/A$, where $\gamma_\eta^R = 0$ and the candidate promise is the same for every type. For part (ii), at $\phi = C$ the slope is $\gamma_\eta^R = -BCc_\eta$, the distortion is $AC^2\beta\eta Bc_\eta v_1$, and $\mathcal{K}_C(\eta) = C^2(1 - A\eta)$, so uniqueness requires $\eta < 1/A$. This holds throughout the revelation region, because $\beta C\eta^\star = \sqrt{1 + \beta C/A} - 1 < \beta C/A$ and hence $\eta^\star < 1/A$.

Both components private. Let the revealed component be the pure forecast, with v_1 separately revealed and carrying its own promise $g_\phi^R(v_1) = \gamma_\eta^R v_1$ from (B.19). By Proposition 1, the public attaches to a forecast report r the additional promise $g_\phi^R(r) = \gamma^R r = -[\beta C\phi + A(\phi - C)]Br$ and the additional expectation $\phi r + g_\phi^R(r) = ABCr$. Under assured delivery the report-dependent loss of a central bank of type (v_1, s) is

$$J_\phi(v_1, s, r) \equiv \frac{C}{2}\Big[v_1 + \beta\big(e_\phi^R(v_1) + ABCr\big)\Big]^2 + \frac{\beta A}{2}\Big[(\phi - C)(\eta v_1 + s) + g_\phi^R(v_1) + g_\phi^R(r)\Big]^2. \tag{A.89}$$

At the truthful report $r = s$, the composite in the first bracket is $AB(c_\eta v_1 + \beta C s)$ and the deviation in the second is $-BC(c_\eta v_1 + \beta C s)$, by the proof of Lemma 16 and Proposition 1. Differentiating with respect to r therefore gives

$$\partial_r J_\phi(v_1, s, s) = ABC\beta\Big[ABC + B\big(\beta C\phi + A(\phi - C)\big)\Big]\big(c_\eta v_1 + \beta C s\big) = ABC\beta\phi\big(c_\eta v_1 + \beta C s\big), \tag{A.90}$$

since the bracket equals ϕ. The distortion is again the rule's loading on the future content of the report, here one for one, multiplied by the posterior-mean statistic $c_\eta v_1 + \beta C s$ (the composite $M(q)$ under full revelation of both components), and it is zero at $\phi = 0$. Holding $v_1 = 0$ to isolate the forecast, the same computation without the correction gives the uncorrected report $r^0 = \{(C-\phi)[A(C-\phi) - \beta C\phi]/[A(C-\phi)^2 + \beta C\phi^2]\}s$, whose slope equals one at $\phi = 0$, one half at $\phi = C - \bar{c}$, and zero at $\phi = C$.

The uncorrected report for the nowcast. Returning to the forecast-bundling nowcast case ($s = 0$ almost surely), without the correction the first-order condition of the uncorrected report problem is linear in (v_1, r) and gives $r^0 = \varsigma_\phi(\eta) v_1$ with

$$\varsigma_\phi(\eta) = 1 - \frac{ABC\beta\eta\phi c_\eta}{C\beta^2(\phi\eta + \gamma_\eta^R)^2 + \beta A(\gamma_\eta^R)^2} \tag{A.91}$$

$$= \frac{A^2\eta^2(C-\phi)^2 + AC\beta\eta^2\phi(\phi - C) - 2AC^2\eta + AC\eta\phi + C^2}{A^2\eta^2(C-\phi)^2 + AC\beta\eta^2\phi^2 - 2AC^2\eta + 2AC\eta\phi + C^2}. \tag{A.92}$$

At $\phi = C$ this gives

$$r^0(v_1) = \varsigma_C(\eta) v_1, \qquad \varsigma_C(\eta) = \frac{1 - A\eta}{1 + AC\beta\eta^2}. \tag{A.93}$$

At $\phi = 0$, $\varsigma_0 = 1$ when $\eta \neq 1/A$; at $\eta = 1/A$ the candidate promise is constant and all reports are payoff equivalent. These best responses are one-shot deviations against the truthful reading; the equilibrium in which the public anticipates them is the subject of Appendix A.15.

A.15 Guidance without a Correction

Preliminaries

Throughout this subsection Assumption 1 is in force, and we treat the pure-forecast case: the nowcast v_1 is public and equal to zero, the central bank's private information is the pure forecast $s \in [-S, S]$, drawn from a distribution symmetric about zero with a density that is positive on $[-S, S]$ (a regularity condition used only in Lemma 9, where it places the mean of a pool strictly inside it), and $v_2 = s + \varepsilon$ with $\mathbb{E}[\varepsilon \mid s] = 0$. Write $\Lambda \equiv A + \beta C = 1/B$. The nowcast case (Appendix B.3) is treated in Steps 5–8 and the nowcast parts of Steps 9–10 of the proof of Lemma 4, and at the end of the proof of Proposition 3.

The game. The central bank observes s and announces $g \in \mathbb{R}$. The public observes g, forms the posterior mean $\widehat{s}(g) = \mathbb{E}[s \mid g]$, and expects the delivered center, which under

assured delivery is $\phi v_2 + g$, so that

$$e_1(g) = \mathbb{E}[\phi v_2 + g \mid g] = \phi\, \widehat{s}(g) + g. \tag{A.94}$$

Given e_1, the central bank sets current policy optimally; by Lemma 1(ii) and (39) its current loss is $\frac{C}{2}[v_1 + \beta e_1]^2 = \frac{C}{2}[\beta e_1]^2$. In period 2 it delivers $\theta_2 = \phi v_2 + g$, so by (11) its period-2 loss is $\frac{A}{2}[(\phi - C)v_2 + g]^2 + \frac{C}{2}v_2^2$, whose conditional expectation given s is $\frac{A}{2}[(\phi - C)s + g]^2$ plus terms that do not depend on g. Dropping those terms, the central bank of type s that announces g and is read as $\widehat{s}$ incurs

$$J(s, g; \widehat{s}) \equiv \frac{C}{2}\big[\beta(\phi\widehat{s} + g)\big]^2 + \frac{\beta A}{2}\big[(\phi - C)s + g\big]^2. \tag{A.95}$$

A strategy is a map G from forecasts to announcements. A perfect Bayesian equilibrium is a strategy together with a belief map $\widehat{s}(\cdot)$ such that (a) $\widehat{s}(g) = \mathbb{E}[s \mid G(s) = g]$ at every announcement made by some type, (b) $\widehat{s}(g) \in [-S, S]$ at every other announcement, and (c) for every type s, $G(s)$ minimizes $J(s, g; \widehat{s}(g))$ over $g \in \mathbb{R}$. The equilibrium is *separating* if G is one-to-one, in which case $\widehat{s}(G(s)) = s$; a set of types sharing an announcement is a *pool*. An equilibrium is *monotone* if G is nondecreasing and *symmetric* if $G(-s) = -G(s)$ and $\widehat{s}(-g) = -\widehat{s}(g)$, which is the natural class given the symmetry of the forecast distribution. A reading in the sense of the text is the belief map restricted to the image of a separating strategy, $\widehat{s} = G^{-1}$.

Three properties of (A.95) carry every argument below.

(P1) *Common belief term.* The belief enters only through $\frac{C}{2}[\beta(\phi\widehat{s} + g)]^2$, which is the same for every type. Hence all types rank beliefs alike: at a given announcement g, every type prefers the belief that brings $\phi\widehat{s} + g$ closest to zero, and the set of beliefs at which a type is better off than at some benchmark is an interval centered at $-g/\phi$.

(P2) *Single crossing.* $\partial_g J(s, g; \widehat{s}) = C\beta^2(\phi\widehat{s} + g) + \beta A[(\phi - C)s + g]$ is affine in s with slope $\partial_{sg} J = \beta A(\phi - C) < 0$: at any announcement and any belief, a larger forecast has a smaller marginal cost of a larger promise.

(P3) *Complete-information promise.* For a fixed belief, J is a strictly convex quadratic in g with $\partial_{gg} J = C\beta^2 + \beta A$. Under the belief $\widehat{s} = s$ its minimizer is

$$\arg\min_g J(s, g; s) = \frac{\beta A(C - \phi) - C\beta^2\phi}{C\beta^2 + \beta A}\, s = \frac{A(C - \phi) - \beta C\phi}{A + \beta C}\, s = \gamma^R s, \tag{A.96}$$

which is the rule-constrained Ramsey promise of Proposition 1. In the revelation

region $\gamma^R > 0$: $\gamma^R > 0$ if and only if $\phi < AC/(A + \beta C) = ABC$, and $C - \bar{c} = C(1 - \sqrt{\beta BC}) < C(1 - \beta BC) = ABC$ because $\beta BC < 1$.

Two consequences of (P1) and (P2) are used repeatedly and are recorded here. First, for any $s > 0$, any belief, and any $g \leq 0$,

$$J(s, g; \widehat{s}) \geq \frac{\beta A}{2}\big[(C - \phi)s - g\big]^2 \geq \frac{\beta A}{2}(C - \phi)^2 s^2 \equiv \Pi(s), \tag{A.97}$$

because $(C - \phi)s \geq 0$ and $-g \geq 0$ have the same sign; $\Pi(s)$ is the loss of a forecast that shares the uninformative announcement $g = 0$ with belief zero, for which $e_1 = 0$ and the rule's omission goes unrepaired. Second, for a differentiable separating strategy G on an interval of positive forecasts, write $V(s, s') \equiv J(s, G(s'); s')$ for the loss of type s that announces what type s' announces and is read as s'. Then

$$\partial_{s'} V(s, s') = \partial_g J(s, G(s'); s')\, G'(s') + \partial_{\widehat{s}} J(s, G(s'); s'). \tag{A.98}$$

Criterion D1. For an announcement g made by no type, criterion D1 (Cho and Kreps, 1987; Banks and Sobel, 1987) restricts the public's belief as follows. Under assured delivery the public's best response is a singleton, $\mathrm{BR}(\mu, g) = \{\phi\, \mathbb{E}_\mu[s \mid g] + g\}$, so the receiver's action is the reading and the gain sets below are sets of readings. For each type s let $\mathcal{B}(s; g)$ be the set of beliefs $\widehat{s} \in [-S, S]$ under which type s is at least as well off announcing g as in equilibrium, $J(s, g; \widehat{s}) \leq J^{\mathrm{eq}}(s)$, where $J^{\mathrm{eq}}(s)$ is the equilibrium loss of type s, and $\mathcal{B}^\circ(s; g)$ the set under which it is strictly better off. If $\mathcal{B}(s; g) \subseteq \mathcal{B}^\circ(s'; g)$ for some s' with $\mathcal{B}^\circ(s'; g) \neq \varnothing$, the belief at g puts no weight on s. By (P1),

$$\mathcal{B}(s; g) = \Big\{\widehat{s} \in [-S, S] : (\phi\widehat{s} + g)^2 \leq c(s; g)\Big\}, \qquad c(s; g) \equiv \frac{2}{C\beta^2}\Big[J^{\mathrm{eq}}(s) - \frac{\beta A}{2}\big((\phi - C)s + g\big)^2\Big], \tag{A.99}$$

and $\mathcal{B}^\circ(s; g)$ is the same set with a strict inequality: intervals centered at $-g/\phi$, intersected with $[-S, S]$ (empty when $c < 0$). Hence $c(s; g) < c(s'; g)$ implies $\mathcal{B}(s; g) \subseteq \mathcal{B}^\circ(s'; g)$, and, whenever some strict-gain set is nonempty, every nonmaximizer of $c(\cdot; g)$ is deleted. A maximizer is not deleted: its weak-gain set contains a boundary point of any nonempty strict-gain set. Such a boundary exists because the equilibrium belief makes no deviation strictly profitable. D1 then puts all weight on maximizers of $c(\cdot; g)$. If all strict-gain sets are empty, D1 imposes no restriction; support on these maximizers remains admissible. We use this supporting choice below, call the maximizers the D1 types, and take $\widehat{s}(g)$ in their convex hull. The following two lemmas are used in the proofs of Propositions 3, 4, and 6.

Lemma 8. *Let $0 < \phi < C - \bar{c}$ and define*

$$\varphi(u) \equiv \frac{\beta A}{2}\left[2(C-\phi)u - u^2\right] - \frac{C\beta^2}{2}(\phi + u)^2.$$

For any schedule G on an interval of positive forecasts, with $u(s) \equiv G(s)/s$, the gain of type s from announcing $G(s)$ and being read as itself rather than sharing the uninformative announcement, $\Phi(s) \equiv \Pi(s) - J(s, G(s); s)$, equals $s^2\varphi(u(s))$. The function φ is a concave quadratic with $\varphi(\gamma^R) > 0$; if $\phi \leq \hat{\phi}$ it is positive on $[\gamma_1, \gamma^R]$, with $\varphi(\gamma_i) = \frac{1}{2}\left[\beta A\gamma_i^2 + \beta^2 C(\phi + \gamma_i)^2\right]$ at the two self-confirming slopes.

Proof. With $G(s) = us$, $J(s, us; s) = \frac{C\beta^2}{2}(\phi + u)^2 s^2 + \frac{\beta A}{2}(\phi - C + u)^2 s^2$ and $\Pi(s) = \frac{\beta A}{2}(C - \phi)^2 s^2$; subtracting and expanding $(C-\phi)^2 - (C - \phi - u)^2 = 2(C-\phi)u - u^2$ gives φ. Its leading coefficient is $-\frac{1}{2}(\beta A + C\beta^2) < 0$. At $u = \gamma^R$, $s^2\varphi(\gamma^R)$ is the loss under opacity minus the rule-constrained Ramsey loss, the value of the information, which is positive for $\phi < C - \bar{c}$ by Proposition 1. At a self-confirming slope γ_i the fixed-point relation (44), $A(C-\phi)\gamma_i = A\gamma_i^2 + \beta C(\gamma_i + \phi)^2$, gives $\frac{\beta A}{2} \cdot 2(C-\phi)\gamma_i = \beta A\gamma_i^2 + \beta^2 C(\gamma_i + \phi)^2$, hence $\varphi(\gamma_i) = \frac{1}{2}[\beta A\gamma_i^2 + \beta^2 C(\phi + \gamma_i)^2] > 0$. A concave function that is positive at γ_1, at γ_2, and at γ^R is positive on $[\gamma_1, \gamma^R]$.

Lemma 9. *In a symmetric monotone equilibrium satisfying D1, no pool of forecasts has a nonzero mean.*

Proof. In a monotone equilibrium a pool is an interval. A pool with a nonzero mean lies on one side of zero: if a pool contained types of both signs, its mirror image would be a pool announcing the negative of its announcement and would overlap it, so the two would be the same pool, with announcement zero and mean zero. By symmetry it suffices to treat a pool $[s_a, s_c]$ with $0 \leq s_a < s_c$, all of whose members announce g_p and are read as its mean $m > s_a$. Its announcement is positive, $g_p > 0$: a symmetric monotone strategy assigns nonnegative announcements to nonnegative forecasts, and $g_p = 0$ would make the pool the mirror of itself, with mean zero. For $\epsilon \in (0, g_p)$ let $g = g_p - \epsilon > 0$. If some type announces g in equilibrium, monotonicity places every such type below the pool, and since nonpositive forecasts announce nonpositive values, in $(0, s_a)$; hence $0 \leq \widehat{s}(g) < s_a$. At belief s_a, type s_a strictly gains for sufficiently small ϵ, since the expectations term falls strictly while the delivery term converges. Thus, if g is off path, its strict-gain set is nonempty and D1 applies. For a pooled type s, $J^{\text{eq}}(s) = J(s, g_p; m)$, and expanding the squares in $c(s; g)$,

$$c(s; g) = (\phi m + g_p)^2 + \frac{A}{C\beta}\left[\left((\phi - C)s + g_p\right)^2 - \left((\phi - C)s + g_p - \epsilon\right)^2\right]$$

$$= (\phi m + g_p)^2 + \frac{A}{C\beta}\Big[2\epsilon\big((\phi - C)s + g_p\big) - \epsilon^2\Big], \tag{A.100}$$

which is strictly decreasing in s because $\phi < C$; among pooled types, s_a has the largest c. For a type $s > s_c$ above the pool, equilibrium requires $J^{\text{eq}}(s) \le J(s, g_p; m)$ (it could join the pool), so $c(s; g)$ is at most the right side of (A.100) evaluated at s, which is below its value at s_c and hence below $c(s_a; g)$. For a type $-s < 0$, symmetry gives $J^{\text{eq}}(-s) = J^{\text{eq}}(s)$ while $((C - \phi)s + g)^2 \ge ((C - \phi)s - g)^2$ for $s, g \ge 0$, so $c(-s; g) \le c(s; g)$: a negative type is never a D1 type unless its mirror is one too. Therefore every D1 type at g lies in $[-s_a, s_a]$, and so does the belief $\widehat{s}(g)$. In both cases, then, $\widehat{s}(g) \le s_a < m$ and $\widehat{s}(g) \ge -s_a$. Now compare type s_a's losses. Its equilibrium loss is $J(s_a, g_p; m)$; at the deviation its loss is $J(s_a, g_p - \epsilon; \widehat{s}(g))$. As $\epsilon \to 0$ the delivery terms converge, while in the expectations terms, since $|\widehat{s}(g)| \le s_a$ and $g > 0$ give $|\phi\widehat{s}(g) + g| \le \phi s_a + g$,

$$(\phi\widehat{s}(g) + g)^2 \le (\phi s_a + g)^2 \to (\phi s_a + g_p)^2 < (\phi m + g_p)^2,$$

uniformly in the belief. For ϵ small enough the deviation is strictly profitable, contradicting equilibrium.

Proofs of Lemmas 4 and 18

Proof. Step 1: the best response to a reading of slope γ. Under the reading $\gamma \neq 0$ the public infers $\widehat{s} = g/\gamma$, so by (A.94) $e_1(g) = (1 + \phi/\gamma)g$ and the central bank's loss is (42). Its derivative in g is

$$\partial_g J = C\beta^2\big(1 + \tfrac{\phi}{\gamma}\big)^2 g + \beta A\big[(\phi - C)s + g\big],$$

and setting it to zero gives $g\big[C\beta^2(1 + \phi/\gamma)^2 + \beta A\big] = \beta A(C - \phi)s$; dividing by β,

$$g = \frac{A(C - \phi)}{A + \beta C(1 + \phi/\gamma)^2}\, s = b(\gamma)s,$$

which is (43). The second derivative $C\beta^2(1 + \phi/\gamma)^2 + \beta A$ is positive, so this is the minimizer.

Step 2: the fixed-point equation. The reading is self-confirming when $b(\gamma) = \gamma$, that is, when $\gamma[A + \beta C(1 + \phi/\gamma)^2] = A(C - \phi)$. Expanding $(1 + \phi/\gamma)^2 = (\gamma + \phi)^2/\gamma^2$ and multiplying by γ,

$$A\gamma^2 + \beta C(\gamma + \phi)^2 = A(C - \phi)\gamma \quad \Longleftrightarrow \quad (A + \beta C)\gamma^2 - \big[A(C - \phi) - 2\beta C\phi\big]\gamma + \beta C\phi^2 = 0,$$

which is (44). Write $Q(\gamma)$ for its left side. Since we multiplied by $\gamma \neq 0$, the roots of Q and

the fixed points of b coincide.

Step 3: existence and sign of the roots. The product of the roots of Q is $\beta C\phi^2/\Lambda > 0$, so real roots have a common sign, and their sum is $[A(C-\phi) - 2\beta C\phi]/\Lambda$. The discriminant is

$$\mathcal{C} \equiv \left[A(C-\phi) - 2\beta C\phi\right]^2 - 4\Lambda\beta C\phi^2.$$

If the sum is zero, $\mathcal{C} = -4\Lambda\beta C\phi^2 < 0$ and there are no real roots. Suppose $\mathcal{C} \geq 0$ with a negative sum, $A(C-\phi) < 2\beta C\phi$. Then $\mathcal{C} \geq 0$ requires $2\beta C\phi - A(C-\phi) \geq 2\phi\sqrt{\Lambda\beta C}$, that is, $A(C-\phi) \leq 2\phi(\beta C - \sqrt{\Lambda\beta C}) < 0$ because $\Lambda > \beta C$; this contradicts $\phi < C$. Hence real roots have a positive sum and are both positive, and they exist if and only if

$$A(C-\phi) - 2\beta C\phi \geq 2\phi\sqrt{\beta C\Lambda} \quad \Longleftrightarrow \quad AC \geq \phi\left[A + 2\beta C + 2\sqrt{\beta C\Lambda}\right].$$

The bracket is a perfect square, $A + 2\beta C + 2\sqrt{\beta C\Lambda} = \Lambda + \beta C + 2\sqrt{\Lambda\beta C} = (\sqrt{\Lambda} + \sqrt{\beta C})^2$, so the condition is $\phi \leq AC/(\sqrt{\Lambda} + \sqrt{\beta C})^2 = \hat{\phi}$, which is (45).

Step 4: both roots lie below γ^R. By (A.96), $\Lambda\gamma^R = A(C-\phi) - \beta C\phi$. Evaluating Q at γ^R and using this identity to eliminate $\Lambda(\gamma^R)^2$,

$$Q(\gamma^R) = \gamma^R\left[\Lambda\gamma^R - A(C-\phi) + 2\beta C\phi\right] + \beta C\phi^2 = \gamma^R\beta C\phi + \beta C\phi^2 = \beta C\phi(\phi + \gamma^R) > 0,$$

and $Q'(\gamma^R) = 2\Lambda\gamma^R - A(C-\phi) + 2\beta C\phi = A(C-\phi) > 0$. A convex quadratic that is positive and increasing at γ^R has both roots to the left of γ^R. With Step 3 this gives $0 < \gamma_1 \leq \gamma_2 < \gamma^R$, which completes the existence and location statements of Lemma 4.

Step 5: the nowcast best response and fixed-point equation. With $s = 0$ the private information is v_1, the public infers $\widehat{v}_1 = g/\gamma$ and expects $e_1(g) = \phi\eta\widehat{v}_1 + g = kg$ with $k \equiv 1 + \phi\eta/\gamma$, and the central bank of type v_1 minimizes $\frac{C}{2}[v_1 + \beta kg]^2 + \frac{\beta A}{2}[(\phi - C)\eta v_1 + g]^2$ (the current loss now contains the shock directly, by (B.21)). The first-order condition $C\beta k[v_1 + \beta kg] + \beta A[(\phi - C)\eta v_1 + g] = 0$ gives, after division by β,

$$g = \frac{A(C-\phi)\eta - Ck}{A + \beta Ck^2}\, v_1.$$

Setting the slope equal to γ, multiplying by $A + \beta Ck^2$, substituting $k = (\gamma + \phi\eta)/\gamma$, and multiplying by γ,

$$A\gamma^2 + \beta C(\gamma + \phi\eta)^2 = A(C-\phi)\eta\gamma - C\gamma - C\phi\eta,$$

which rearranges to (B.24):

$$\Lambda\gamma^2 + \big[C + 2\beta C\phi\eta - A(C-\phi)\eta\big]\gamma + C\phi\eta(1+\beta\phi\eta) = 0.$$

For the positive-news, positive-rule case, $\eta > 0$ and $\phi > 0$, call its left side $N(\gamma)$, its linear coefficient n_1, and its constant term $n_0 \equiv C\phi\eta(1+\beta\phi\eta) > 0$. Steps 6–8 treat this case; Step 9 treats the strict target. At $\eta = 0$, the quadratic factors as $\gamma(\Lambda\gamma + C) = 0$, so the only nonzero reading has slope $-C/\Lambda = -BC$.

Step 6: the discriminant as a quadratic in $\xi = \phi\eta$. Write $n_1 = (A + 2\beta C)\xi + C(1 - A\eta)$ and $n_0 = C\xi + \beta C\xi^2$. Then

$$\begin{aligned} n_1^2 - 4\Lambda n_0 &= (A+2\beta C)^2\xi^2 + 2(A+2\beta C)C(1-A\eta)\xi + C^2(1-A\eta)^2 - 4\Lambda C\xi - 4\Lambda\beta C\xi^2 \\ &= A^2\xi^2 - 2AC(1 + A\eta + 2\beta C\eta)\xi + C^2(1-A\eta)^2, \end{aligned}$$

using $(A + 2\beta C)^2 - 4\beta C\Lambda = A^2$ and $2(A + 2\beta C)C(1 - A\eta) - 4\Lambda C = -2AC(1 + A\eta + 2\beta C\eta)$. This is a convex quadratic in ξ with roots

$$\begin{aligned} \xi_\pm &= \frac{C}{A}\Big[(1 + A\eta + 2\beta C\eta) \pm \sqrt{(1 + A\eta + 2\beta C\eta)^2 - (1 - A\eta)^2}\Big] \\ &= \frac{C}{A}\Big[(1 + A\eta + 2\beta C\eta) \pm 2\sqrt{\eta\Lambda(1 + \beta C\eta)}\Big], \end{aligned}$$

since $(1 + A\eta + 2\beta C\eta)^2 - (1 - A\eta)^2 = (2A\eta + 2\beta C\eta)(2 + 2\beta C\eta) = 4\eta\Lambda(1 + \beta C\eta)$. The discriminant is nonnegative for $\xi \le \xi_-$ and for $\xi \ge \xi_+$; the second branch is irrelevant because $\xi_+ > (C/A)A\eta = C\eta \ge \phi\eta$ for every $\phi \le C$. Hence real roots exist if and only if $\phi\eta \le \xi_-$, that is, $\phi \le \xi_-/\eta \equiv \hat{\phi}(\eta)$. Finally, because $(\sqrt{\eta\Lambda})^2 + (\sqrt{1 + \beta C\eta})^2 = 1 + A\eta + 2\beta C\eta$, the bracket in ξ_- is a perfect square, $\xi_- = (C/A)\big(\sqrt{\eta\Lambda} - \sqrt{1 + \beta C\eta}\big)^2$, and dividing by η gives (B.25).

Step 7: sign and location of the nowcast roots. The product of the roots is $n_0/\Lambda > 0$, so they share a sign; their sum is $-n_1/\Lambda$. The rule-constrained Ramsey slope satisfies $\Lambda\gamma_\eta^R = A(C-\phi)\eta - C - \beta C\phi\eta$, so

$$-n_1 = \Lambda\gamma_\eta^R - \beta C\phi\eta. \tag{A.101}$$

Suppose $\gamma_\eta^R > 0$ and real roots exist, but the sum is not positive, so that $0 < \Lambda\gamma_\eta^R \le \beta C\phi\eta$ and $0 \le n_1 < \beta C\phi\eta$. Then $n_1^2 - 4\Lambda n_0 < \beta^2C^2\phi^2\eta^2 - 4\Lambda C\phi\eta(1+\beta\phi\eta) = C\phi\eta[\beta^2C\phi\eta - 4\Lambda(1+\beta\phi\eta)] < 0$, a contradiction. Hence with $\gamma_\eta^R > 0$ both roots are positive and, by (A.101), $\gamma_2 \le \gamma_1 + \gamma_2 = \gamma_\eta^R - \beta C\phi\eta/\Lambda < \gamma_\eta^R$. Suppose instead $\gamma_\eta^R < 0$ and real roots exist. By (A.101)

$n_1 > 0$, so both roots are negative. Real roots require $n_1 \geq 2\sqrt{\Lambda n_0} = 2\sqrt{\Lambda C\phi\eta(1+\beta\phi\eta)} > 2\phi\eta\sqrt{\Lambda\beta C} > 2\beta C\phi\eta$ (as $\Lambda > \beta C$), and $n_1 > 2\beta C\phi\eta$ is the statement $C > A(C-\phi)\eta$. The vertex of N is at $-n_1/(2\Lambda)$, and by (A.101) $\gamma_\eta^R = (-n_1 + \beta C\phi\eta)/\Lambda$ lies to the left of the vertex exactly when $n_1 > 2\beta C\phi\eta$. Moreover, eliminating $\Lambda(\gamma_\eta^R)^2$ with (A.101),

$$N(\gamma_\eta^R) = \beta C\phi\eta\,\gamma_\eta^R + C\phi\eta(1+\beta\phi\eta) = C\phi\eta\big[1+\beta(\phi\eta+\gamma_\eta^R)\big] = C\phi\eta\big[1+\beta BC(A\eta-1)\big] > 0,$$

where $\phi\eta + \gamma_\eta^R = BC(A\eta-1)$ is the expectation slope of Lemma 16, the same for every ϕ, and $1+\beta BC(A\eta-1) > 0$ because $\beta BC < 1$. A convex quadratic that is positive at a point left of its vertex has both roots to the right of that point: $\gamma_\eta^R < \gamma_1 \leq \gamma_2 < 0$. If $\gamma_\eta^R = 0$, then $n_1 = \beta C\phi\eta$ and $n_1^2 - 4\Lambda n_0 = C\phi\eta[\beta^2 C\phi\eta - 4\Lambda(1+\beta\phi\eta)] < 0$, so no real roots exist. In the two remaining cases the roots have the sign of γ_η^R and lie between 0 and γ_η^R.

Step 8: the shape of $\hat\phi(\eta)$. From (B.25), $\hat\phi(\eta) = \frac{C}{A}\big[\sqrt{\Lambda} - \sqrt{\beta C + 1/\eta}\big]^2$ vanishes exactly when $\Lambda = \beta C + 1/\eta$, that is, at $\eta = 1/A$. The term $\sqrt{\beta C + 1/\eta}$ decreases in η; it exceeds $\sqrt{\Lambda}$ for $\eta < 1/A$ and is below it for $\eta > 1/A$, so $\hat\phi$ decreases on $(0, 1/A)$ and increases on $(1/A, \infty)$, with limit $\frac{C}{A}(\sqrt{\Lambda} - \sqrt{\beta C})^2 = \frac{C}{A}\cdot\frac{A^2}{(\sqrt{\Lambda}+\sqrt{\beta C})^2} = \hat\phi$, using $(\sqrt{\Lambda} - \sqrt{\beta C})(\sqrt{\Lambda} + \sqrt{\beta C}) = \Lambda - \beta C = A$. Finally $\hat\phi(\eta) = C$ on the decreasing branch requires $\sqrt{\beta C + 1/\eta} - \sqrt{\Lambda} = \sqrt{A}$, that is, $1/\eta = (\sqrt{\Lambda} + \sqrt{A})^2 - \beta C = 2A + 2\sqrt{A\Lambda}$, so $\eta = 1/[2\sqrt{A}(\sqrt{A} + \sqrt{\Lambda})]$. The revelation threshold of Lemma 15 is $\eta^\star = (\sqrt{\Lambda/A} - 1)/(\beta C) = (\sqrt{\Lambda} - \sqrt{A})/(\sqrt{A}\beta C)$, and $(\sqrt{\Lambda} - \sqrt{A})(\sqrt{\Lambda} + \sqrt{A}) = \beta C$ gives $\eta^\star/2 = 1/[2\sqrt{A}(\sqrt{\Lambda} + \sqrt{A})]$: the two coincide. This completes the existence, location, and threshold statements of Lemma 18.

Step 9: the strict target. At $\phi = 0$, (A.94) gives $e_1(g) = g$ for every belief. For the pure forecast the loss is $\frac{C}{2}\beta^2 g^2 + \frac{\beta A}{2}(g - Cs)^2$, with first-order condition $C\beta^2 g + \beta A(g - Cs) = 0$ and minimizer $g = ACs/(A + \beta C) = ABC\,s = \gamma^R s$. For the nowcast the loss is $\frac{C}{2}(v_1 + \beta g)^2 + \frac{\beta A}{2}(g - C\eta v_1)^2$, with first-order condition $C\beta(v_1 + \beta g) + \beta A(g - C\eta v_1) = 0$ and minimizer $g = C(A\eta - 1)v_1/(A + \beta C) = BC(A\eta - 1)v_1 = g_0^R(v_1)$, the promise of Corollary 8(iii). Neither depends on γ, so the only self-confirming reading is the truthful one; for the nowcast this requires $\eta \neq 1/A$, since at $\eta = 1/A$ the promise is zero for every type and no strictly monotone reading exists. This is the reading statement of Proposition 4, and the strict-target statement of Lemma 18.

Step 10: the social loss. Under a self-confirming reading of slope γ the public reads correctly, so the allocation is that of the planner who promises γs. Per unit of $\mathrm{Var}(s)$ the planner's loss is $\mathcal{L}(\gamma) = \frac{C}{2}[\beta(\phi + \gamma)]^2 + \frac{\beta A}{2}[(\phi - C) + \gamma]^2$, a quadratic in γ with $\mathcal{L}'(\gamma) = C\beta^2(\phi + \gamma) + \beta A(\phi - C + \gamma)$, which vanishes at γ^R by (A.96), and $\mathcal{L}'' = C\beta^2 + \beta A$. The exact second-order expansion around the minimizer is $\mathcal{L}(\gamma) - \mathcal{L}(\gamma^R) = \frac{1}{2}(\beta A + \beta^2 C)(\gamma - \gamma^R)^2$. For the nowcast, $\mathcal{L}_\eta(\gamma) = \frac{C}{2}[1 + \beta(\phi\eta + \gamma)]^2 + \frac{\beta A}{2}[(\phi - C)\eta + \gamma]^2$ has the same second derivative

and is minimized at γ_η^R (its first-order condition rearranges to $\Lambda\gamma = -C - \beta C\phi\eta - A(\phi - C)\eta$). This is the loss statement of Lemma 4 and that of Lemma 18.

Proof of Proposition 3

The eight steps split at $\hat{\phi}$. Steps 1 to 3 pin the separating schedule and show there is at most one, Steps 4 to 6 establish that it exists and is incentive compatible when $\phi \le \hat{\phi}$ and that it fails above, and Steps 7 and 8 turn to pools: none below the threshold, and exactly one collapse equilibrium above it.

Proof. Step 1: the boundary condition, part (i). Let $(G, \widehat{s})$ be an equilibrium in which type S is read as itself, $\widehat{s}(G(S)) = S$. Suppose $G(S) \neq \gamma^R S$, and let type S announce $g' = \gamma^R S$ instead. Since $\gamma^R \ge 0$ (P3), $g' \ge 0$, and for every belief $\widehat{s}' \in [-S, S]$,

$$-\phi S + g' \le \phi\widehat{s}' + g' \le \phi S + g', \qquad \text{with} \qquad -(\phi S + g') \le -\phi S + g',$$

so $|\phi\widehat{s}' + g'| \le \phi S + g'$ and, by (P1), $J(S, g'; \widehat{s}') \le J(S, g'; S)$. By (P3), $J(S, g'; S) = \min_g J(S, g; S) < J(S, G(S); S)$, the last being type S's equilibrium loss. Whatever belief the public attaches to g', on or off the path, the deviation is profitable, contradicting equilibrium. Hence $G(S) = \gamma^R S$. The same argument for type $-S$ with $g' = -\gamma^R S$, using $|\phi\widehat{s}' + g'| \le \phi S - g' = |\phi(-S) + g'|$, gives $G(-S) = -\gamma^R S$.

Step 2: monotonicity and the first-order condition of a separating equilibrium. Let $(G, \widehat{s})$ be a separating equilibrium, so that every type is read as itself and, by Step 1 applied at both ends, $G(S) = \gamma^R S$ and $G(-S) = -\gamma^R S$. *(a) G is strictly increasing.* For two types $s \neq s'$, add the incentive constraints $J(s, G(s); s) \le J(s, G(s'); s')$ and $J(s', G(s'); s') \le J(s', G(s); s)$. By (P1) the belief terms on the two sides are the same and cancel, leaving

$$\big[(\phi - C)s + G(s)\big]^2 + \big[(\phi - C)s' + G(s')\big]^2 \le \big[(\phi - C)s + G(s')\big]^2 + \big[(\phi - C)s' + G(s)\big]^2,$$

which expands to $(\phi - C)(s - s')\big(G(s) - G(s')\big) \le 0$. Since $\phi < C$, $(s - s')(G(s) - G(s')) \ge 0$: G is nondecreasing, and being one-to-one it is strictly increasing. *(b) The first-order condition.* The function $s \mapsto \phi s + G(s)$ is strictly increasing, negative at $-S$ and positive at S; let I^+ be the set of types at which it is positive, an interval with right endpoint S, and I^- the set at which it is negative, an interval with left endpoint $-S$. The regularity argument associated with Mailath (1987) can be verified directly here. At continuity points in I^+, a nearby downward on-path deviation implies $G(s) \le \gamma^R s$; in I^-, an upward deviation implies the reverse inequality. These bounds extend to one-sided limits by monotonicity. At an interior jump, the two limiting incentive constraints give $J(s, G_-; s) = J(s, G_+; s)$,

hence $(G_- + G_+)/2 = \gamma^R s$, contradicting the corresponding bound on both limits. Thus G is continuous on each interval; continuity at the outer endpoints follows from Step 1 and the same limiting constraints. For the continuous inverse $s(g)$, the two incentive constraints for neighboring announcements squeeze its difference quotient to $\Lambda(\gamma^R s - g)/[\beta C\phi(\phi s + g)]$. This is continuous where $\phi s + g \neq 0$, so the inverse is continuously differentiable. Its derivative cannot vanish at an interior point: at such a point its second derivative would be $-\Lambda/[\beta C\phi(\phi s + g)] \neq 0$, forcing the first derivative to change sign, contrary to monotonicity. Thus G is differentiable on both interiors and optimality gives

$$C\beta^2(\phi s + G) + \beta A\big[(\phi - C)s + G\big] + \frac{C\beta^2\phi(\phi s + G)}{G'(s)} = 0.$$

Multiplying by G'/β and solving for G',

$$G'(s) = \frac{\beta C\phi(\phi s + G)}{A(C - \phi)s - AG - \beta C(\phi s + G)} = \frac{\beta C\phi(\phi s + G)}{\Lambda(\gamma^R s - G)}, \tag{A.102}$$

where the second form uses $A(C-\phi)s - \beta C\phi s - (A+\beta C)G = \Lambda(\gamma^R s - G)$ by (A.96). Equation (A.102) is invariant under $(s, G) \mapsto (-s, -G)$, so the mirror image of a solution is a solution.

Step 3: uniqueness of the separating equilibrium. Write the first-order condition for the inverse schedule $s(g)$:

$$\frac{ds}{dg} = \frac{\Lambda(\gamma^R s - g)}{\beta C\phi(\phi s + g)}. \tag{A.103}$$

The right side is a rational function of (g, s), continuously differentiable on the open set $\mathcal{R} \equiv \{\phi s + g > 0\}$, hence locally Lipschitz in s there. The point $(g, s) = (\gamma^R S, S)$ belongs to $\mathcal{R}$. By the Picard–Lindelöf theorem the initial-value problem for (A.103) from this point has exactly one maximal solution in $\mathcal{R}$; call it the top trajectory, and call its mirror image through $(-\gamma^R S, -S)$ the bottom trajectory. Let $(G, \widehat{s})$ be a separating equilibrium. By Steps 1 and 2 its inverse schedule passes through $(\gamma^R S, S)$ and solves (A.103) on the image of I^+, which lies in $\mathcal{R}$; by uniqueness it coincides with the top trajectory as long as both remain in $\mathcal{R}$, that is, on all of I^+ (if the two parted at some type s_1 of I^+, the initial-value problem from the common point at s_1, which lies in $\mathcal{R}$, would have two solutions). Likewise G coincides with the bottom trajectory on I^-. Step 4 shows that for $\phi \leq \hat{\phi}$ the top trajectory stays in $\mathcal{R}$ for every $s > 0$ and tends to the origin. Then $\inf I^+ = 0$: otherwise, with $s_1 \equiv \inf I^+ > 0$, every $s \in (0, s_1)$ has $\phi s + G(s) \leq 0$, hence $G(s) \leq -\phi s < 0$, while the mirror argument shows that G is positive on $(\sup I^-, 0)$ if $\sup I^- < 0$ and tends to zero along the bottom trajectory as $s \uparrow 0$ if $\sup I^- = 0$; either way some $s' < s$ has $G(s') > G(s)$, contradicting Step 2(a). Likewise $\sup I^- = 0$. Hence G equals the top trajectory on $(0, S]$ and the bottom trajectory

on $[-S, 0)$, and $G(0) = 0$ because G is monotone and both trajectories tend to zero at the origin. There is at most one separating equilibrium, and it is odd.

Step 4: the phase portrait, existence below $\hat{\phi}$ and nonexistence above. Let $u(s) = G(s)/s$ on $(0, S]$. Then $G' = u + su'$, and (A.102) reads $su' = F(u) - u$ with $F(u) = \beta C\phi(\phi + u)/[\Lambda(\gamma^R - u)]$. Bringing to a common denominator and using $\Lambda\gamma^R = A(C - \phi) - \beta C\phi$,

$$F(u) - u = \frac{\beta C\phi(\phi + u) - \Lambda u(\gamma^R - u)}{\Lambda(\gamma^R - u)} = \frac{\Lambda u^2 - [A(C - \phi) - 2\beta C\phi]u + \beta C\phi^2}{\Lambda(\gamma^R - u)} = \frac{Q(u)}{\Lambda(\gamma^R - u)},$$

so that, in the variable $t = \ln s$,

$$\frac{du}{dt} = \frac{Q(u)}{\Lambda(\gamma^R - u)}, \tag{A.104}$$

where $Q(u) = \Lambda(u - \gamma_1)(u - \gamma_2)$ for $\phi \leq \hat{\phi}$, and $Q(u) > 0$ for every u when $\phi > \hat{\phi}$. The top trajectory starts at $u = \gamma^R$ when $t = \ln S$, the pole of the right side, and is followed backward in t (downward in s). The singular start causes no difficulty: the inverse equation (A.103) is regular at the top, with $ds/dg = 0$ there (the schedule has an infinite slope at S), and $u < \gamma^R$ immediately below the top because $d(\gamma^R s - g)/dg = -1$ at $(\gamma^R S, S)$, so the phase equation applies on $(0, S)$.

Below the threshold. For $u \in (\gamma_2, \gamma^R)$ both $Q(u)$ and $\gamma^R - u$ are positive, so $du/dt > 0$: as s falls, u falls. The point γ_2 is a rest point of the autonomous equation (A.104), whose right side is continuously differentiable on (γ_1, γ^R), so trajectories cannot cross it; u decreases monotonically and stays above γ_2, hence converges as $t \to -\infty$, and its limit is a rest point, so it is γ_2. Therefore $\gamma_2 s \leq G(s) \leq \gamma^R s$ on $(0, S]$, $G(s) \to 0$ as $s \to 0$, and $G' = F(u) > 0$ because $\phi + u > 0$ and $\gamma^R - u > 0$: the solution is an increasing schedule from the origin to $(S, \gamma^R S)$, odd by construction on $[-S, 0]$. With Step 5 it is an equilibrium, and with Step 3 the only separating one.

Above the threshold. Now Q has no real root and $Q(u) \geq Q_{\min} > 0$ for all u, so for $u \in [0, \gamma^R)$, $du/dt \geq Q_{\min}/(\Lambda\gamma^R) \equiv \mu > 0$. Following the trajectory backward from $t = \ln S$, u falls at rate at least μ and reaches 0 after a t-interval of length at most γ^R/μ, that is, at a forecast $s_0 \geq S\exp(-\gamma^R/\mu) > 0$. At s_0 the top trajectory satisfies $G(s_0) = 0$ and $\phi s_0 + G(s_0) = \phi s_0 > 0$, so it is still in $\mathcal{R}$, and since $G' = F(u) > 0$ for $u > -\phi$ it is increasing wherever it is defined, with $G < 0$ below s_0. Now let $(G, \widehat{s})$ be a separating equilibrium, and take $s \in (0, s_0)$. If $s \in I^+$, then by Step 3 $G(s)$ lies on the top trajectory, so $G(s) < 0$; if $s \notin I^+$, then $\phi s + G(s) \leq 0$ and again $G(s) \leq -\phi s < 0$. Hence $G < 0$ on $(0, s_0)$, and by the mirror argument $G > 0$ on $(-s_0, 0)$, contradicting the monotonicity of Step 2(a). No separating equilibrium exists, for any S. (At $\phi = \hat{\phi}$ the two roots coincide at a double root γ_*; for $u \in (\gamma_*, \gamma^R)$ the phase equation still has $du/dt > 0$ with γ_* a rest point, so the trajectory from

γ^R descends to γ_* and the conclusions of the case $\phi < \hat{\phi}$ hold with $\gamma_1 = \gamma_2 = \gamma_*$; part (ii) therefore covers $\phi \leq \hat{\phi}$.) Figure A.3 shows the two trajectories at the baseline calibration.

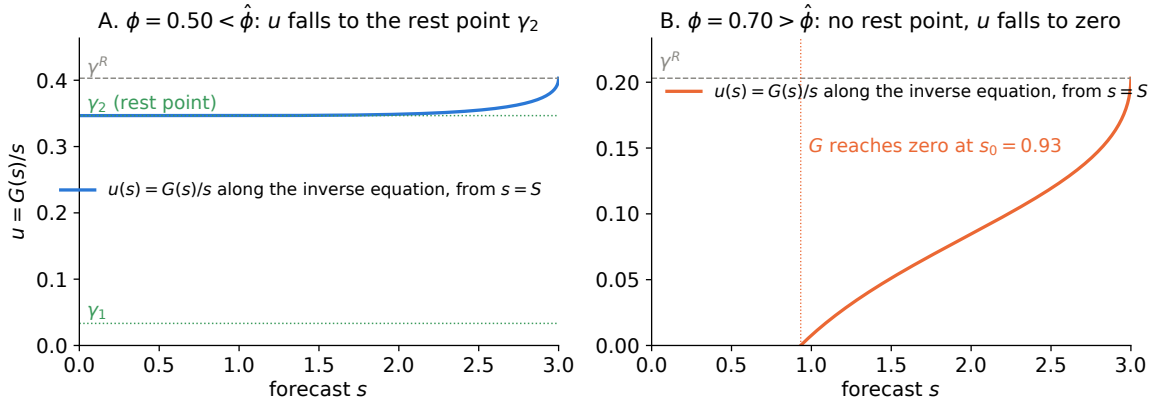

Figure A.3: Phase portrait of the inverse equation: the ratio $u = G(s)/s$ from the largest forecast

Notes: Top trajectory of (A.104), starting at $u(S) = \gamma^R$ and followed downward in s; baseline calibration, $S = 3\sigma_s$. Panel A, $\phi = 0.5 < \hat{\phi}$: u approaches γ_2 from above, and $G(s)$ reaches the origin. Panel B, $\phi = 0.70 > \hat{\phi}$: Q has no real root, and u reaches zero at $s_0 = 0.93 > 0$. This trajectory cannot extend to an increasing schedule through the origin.

Step 5: global incentive compatibility of the separating schedule below the threshold. Fix $s \in (0, S]$; by symmetry this covers negative forecasts. The equilibrium loss of type s is $J^{\text{eq}}(s) \equiv V(s, s) = J(s, G(s); s)$. Three classes of deviations exhaust the announcements. (a) *On-path announcements of positive forecasts*, $g = G(s')$ with $s' \in (0, S]$, read as s'. By (A.98) and the first-order condition of type s' from Step 2, $\partial_g J(s', G(s'); s')G'(s') + \partial_{\widehat{s}} J(s', G(s'); s') = 0$; subtracting it from (A.98) and using that $\partial_{\widehat{s}} J$ does not depend on the type (P1) and that $\partial_g J$ is affine in the type with slope $\beta A(\phi - C)$ (P2),

$$\partial_{s'} V(s, s') = \left[\partial_g J(s, G(s'); s') - \partial_g J(s', G(s'); s')\right] G'(s') = \beta A(\phi - C)(s - s')G'(s').$$

Since $G' > 0$ and $\phi < C$, this is negative for $s' < s$ and positive for $s' > s$ on $(0, S)$, and $V(s, \cdot)$ is continuous at S: $V(s, \cdot)$ decreases up to $s' = s$ and increases beyond, so $V(s, s') \geq V(s, s)$ for every $s' \in (0, S]$. (b) *Nonpositive announcements*, $g \leq 0$, on the path (read as $G^{-1}(g) \leq 0$) or off it. By (A.97), $J(s, g; \widehat{s}) \geq \Pi(s)$ for every belief, and $\Pi(s) \geq J^{\text{eq}}(s)$ by Lemma 8, since $u(s) \in [\gamma_2, \gamma^R]$ by Step 4. (c) *Announcements above the image*, $g > \gamma^R S$, to which the equilibrium attaches the belief $\widehat{s} = S$. Under that belief $J(s, \cdot; S)$ is a convex quadratic whose minimizer is

$$\frac{A(C - \phi)s - \beta C\phi S}{\Lambda} \leq \frac{A(C - \phi)S - \beta C\phi S}{\Lambda} = \gamma^R S < g,$$

so $J(s, g; S) > J(s, \gamma^R S; S) = J(s, G(S); S) = V(s, S) \geq V(s, s)$ by (a). No deviation is profitable.

Step 6: the large-support limit. Equation (A.102) is homogeneous of degree one in (s, G): if G solves it, so does $G_c(s) \equiv cG(s/c)$ for every $c > 0$, with $G_c(cS) = c\gamma^R S$. By the uniqueness of Step 3 the schedule for support S' is $G_{S'}(s) = (S'/S)G_S(sS/S')$, so its ratio satisfies $u_{S'}(s) = u_S(sS/S')$. For fixed s and $S' \to \infty$ the argument tends to zero, and by Step 4 $u_S(\sigma) \to \gamma_2$ as $\sigma \to 0$. Hence $G_{S'}(s) \to \gamma_2 s$ pointwise: the separating equilibrium converges to the near-Ramsey reading.

Unbounded support. With $S = \infty$ nothing pins a largest forecast, and the phase portrait shows what replaces uniqueness (this is the family mentioned in the footnote of Section 3.3). The constant solutions $u \equiv \gamma_1$ and $u \equiv \gamma_2$ are the two readings of a single slope. For $u \in (\gamma_1, \gamma_2)$ one has $Q(u) < 0$ and $\gamma^R - u > 0$, so $du/dt < 0$: a trajectory started at any such u rises to γ_2 as $s \to 0$ and falls to γ_1 as $s \to \infty$, with $G' = F(u) > 0$ throughout and every positive announcement on its path. Each such trajectory is a separating equilibrium of the unbounded game: on-path deviations are ruled out by the single-crossing argument of Step 5(a), nonpositive announcements by (A.97) and the positivity of φ on $[\gamma_1, \gamma^R]$ (Lemma 8), and there are no announcements above the image. Trajectories with $u < \gamma_1$ reach a zero announcement at a positive forecast, and trajectories with $u \in (\gamma_2, \gamma^R)$ reach the pole γ^R at a finite forecast beyond which no increasing schedule continues them; neither is an equilibrium. The self-confirming readings of the unbounded game are therefore a one-parameter family indexed by $u(1) \in [\gamma_1, \gamma_2]$, from the near-Ramsey reading to the reading that attaches the smaller promise $\gamma_1 s$ to large forecasts.

Step 7: criterion D1 and pools; the last sentence of (ii). When a strict gain is feasible, criterion D1 restricts the belief at an announcement made by no type to the types that maximize $c(\cdot; g)$ in (A.99), and Lemma 9 shows that in a symmetric monotone equilibrium satisfying D1 every pool has mean zero. Hence every pool is symmetric about zero, of the form $[-s_p, s_p]$, and by symmetry its common announcement is zero, read as zero, so its members incur $\Pi(s)$. Suppose such a pool exists with $s_p > 0$ and $\phi \leq \hat{\phi}$. Forecasts above s_p separate (there is no other pool), and by Steps 1–3 their schedule is the solution of Step 4. Consider a forecast s just below s_p announcing $G(s')$ for s' just above s_p, read as s'; as both tend to s_p its loss tends to $J(s_p, G(s_p); s_p) = \Pi(s_p) - \Phi(s_p) < \Pi(s_p)$ by Lemma 8, while its equilibrium loss tends to $\Pi(s_p)$. The deviation is profitable for s close enough to s_p: no such pool exists. Together with Steps 3–6 this proves (ii).

Step 8: the collapse band, part (iii). Let $\hat{\phi} < \phi < C - \bar{c}$; by Step 4 no separating equilibrium exists. We construct the equilibrium and then show it is the only symmetric monotone one satisfying D1.

Construction. Let $s(g)$ be the solution of (A.103) from $(\gamma^R S, S)$, followed down to $g = 0$; by Step 4 it reaches $g = 0$ at $s_0 > 0$, and on $[s_0, S]$ it defines an increasing schedule G with $G(s_0) = 0$ and $G(S) = \gamma^R S$, along which $u = G/s$ rises from 0 to γ^R. On this range $\Phi(s) = s^2\varphi(u(s))$ with φ as in Lemma 8: $\varphi(0) = -\frac{C\beta^2}{2}\phi^2 < 0$ and $\varphi(\gamma^R) > 0$, and a concave quadratic that is negative at 0 and positive at γ^R has exactly one root $u_b \in (0, \gamma^R)$, negative below it and positive above. Since $u(\cdot)$ is strictly increasing, there is exactly one $s_b \in (s_0, S)$ with $u(s_b) = u_b$, that is, with $\Phi(s_b) = 0$: the forecast s_b is indifferent between the uninformative announcement and its informative one, forecasts above it strictly prefer the latter, forecasts in (s_0, s_b) the former. Define the strategy: $G^*(s) = 0$ for $|s| \le s_b$, $G^*(s) = \text{sign}(s)G(|s|)$ for $|s| > s_b$; beliefs: $\widehat{s}(0) = 0$; $\widehat{s}(g) = \text{sign}(g)G^{-1}(|g|)$ for $G(s_b) \le |g| \le \gamma^R S$; $\widehat{s}(g) = \text{sign}(g)s_b$ for $0 < |g| < G(s_b)$; $\widehat{s}(g) = \text{sign}(g)S$ for $|g| > \gamma^R S$. Write $G_b \equiv G(s_b)$.

Verification: pooled types, $0 \le s \le s_b$, whose equilibrium loss is $\Pi(s)$. (a) To an on-path informative announcement $G(s')$, $s' \ge s_b$: by the computation of Step 5(a), $\partial_{s'}V(s, s') = \beta A(\phi - C)(s - s')G'(s') > 0$ for $s' > s$, so $V(s, s') \ge V(s, s_b) = J(s, G_b; s_b)$. (b) To an off-path $g \in (0, G_b)$, read as s_b: for a fixed $g \in (0, G_b]$ define $\Xi(s) \equiv J(s, g; s_b) - \Pi(s)$; then

$$\Xi'(s) = \beta A(\phi - C)\big((\phi - C)s + g\big) - \beta A(C - \phi)^2 s = -\beta A(C - \phi)g < 0,$$

so $\Xi(s) \ge \Xi(s_b) = J(s_b, g; s_b) - \Pi(s_b)$. At $s = s_b$, $J(s_b, \cdot; s_b)$ has minimizer $\gamma^R s_b \ge G_b$, so it is decreasing on $(0, G_b]$, and $J(s_b, g; s_b) \ge J(s_b, G_b; s_b) = \Pi(s_b)$ by indifference; hence $\Xi(s) \ge 0$. Cases (a) and (b) together show that no pooled type gains from any positive announcement. (c) To $g \le 0$: (A.97) gives a loss of at least $\Pi(s)$. (d) To $g > \gamma^R S$, read as S: $J(s, g; S) \ge J(s, \gamma^R S; S)$ by the convexity argument of Step 5(c), and $J(s, \gamma^R S; S) = V(s, S) \ge V(s, s_b) \ge \Pi(s)$ by (a) and (b).

Verification: informative types, $s > s_b$, with equilibrium loss $J^{\text{eq}}(s) = V(s, s)$. (a) To on-path announcements of other informative types: Step 5(a). (b) To the uninformative announcement, read as zero: the loss is $\Pi(s) = J^{\text{eq}}(s) + \Phi(s) \ge J^{\text{eq}}(s)$ since $u(s) \ge u_b$. (c) To an off-path $g \in (0, G_b)$, read as s_b: the minimizer of $J(s, \cdot; s_b)$ is $[A(C - \phi)s - \beta C\phi s_b]/\Lambda \ge \gamma^R s_b \ge G_b$, so $J(s, g; s_b) \ge J(s, G_b; s_b) = V(s, s_b) \ge V(s, s)$. (d) To $g \le 0$ and to $g > \gamma^R S$: as in Step 5(b)–(c). Negative forecasts follow by symmetry. Hence $(G^*, \widehat{s})$ is a symmetric monotone equilibrium.

The equilibrium satisfies D1. Its off-path announcements are $0 < |g| < G_b$ and $|g| > \gamma^R S$; by symmetry consider $g > 0$. For $g \in (0, G_b)$: for a pooled type, $c(s; g) = \frac{A}{C\beta}\big[(C - \phi)^2 s^2 - ((C - \phi)s - g)^2\big] = \frac{A}{C\beta}\, g\big[2(C - \phi)s - g\big]$, increasing in s, so among pooled types s_b has the largest c. For an informative type $s \ge s_b$, the envelope theorem applied to $J^{\text{eq}}(s) = V(s, s)$ with the

first-order condition of Step 2 gives $dJ^{\mathrm{eq}}/ds = \partial_s J(s, G(s); s) = \beta A(\phi - C)\big((\phi - C)s + G(s)\big)$, so

$$\begin{aligned}\frac{\partial c(s;g)}{\partial s} &= \frac{2}{C\beta^2}\Big[\beta A(\phi - C)\big((\phi - C)s + G(s)\big) - \beta A(\phi - C)\big((\phi - C)s + g\big)\Big] \\ &= \frac{2A}{C\beta}(\phi - C)\big(G(s) - g\big) < 0,\end{aligned}$$

because $G(s) \geq G_b > g$. Since $c(\cdot; g)$ is continuous at s_b ($J^{\mathrm{eq}}(s_b) = \Pi(s_b)$ by indifference), increasing below s_b and decreasing above it, s_b is its unique maximizer among nonnegative types; negative types are dominated by their mirrors as in Lemma 9. Hence s_b is the unique D1 type at g, and the belief $\widehat{s}(g) = s_b$ of the construction is D1-admissible (and required whenever a strict gain is feasible). Under it no type gains: for a pooled type, $(\phi\widehat{s}(g) + g)^2 \geq (\phi s_b + g)^2 \geq c(s_b; g) \geq c(s; g)$, where the middle inequality is $J(s_b, g; s_b) \geq \Pi(s_b)$ shown in (b) above; for an informative type s, $J(s, g; s_b) \geq J^{\mathrm{eq}}(s)$ by (c) above. For $g > \gamma^R S$: for any type $s < S$, equilibrium requires $J^{\mathrm{eq}}(s) \leq V(s, S) = J(s, \gamma^R S; S)$, so

$$c(s;g) \leq (\phi S + \gamma^R S)^2 + \frac{A}{C\beta}\big(g - \gamma^R S\big)\big[2(C - \phi)s - \gamma^R S - g\big],$$

which is increasing in s, while for $s = S$ the bound holds with equality; hence S is the unique D1 type, $\widehat{s}(g) = S$, and under this belief no type gains, by the argument of Step 5(c). The equilibrium satisfies D1.

Uniqueness among symmetric monotone D1 equilibria. Let $(G, \widehat{s})$ be any such equilibrium. By Step 4 it is not separating, so it has a pool; by Lemma 9 every pool has mean zero, and in a monotone strategy a pool is an interval, so there is exactly one pool, $[-s_p, s_p]$, announcing zero by symmetry. The pool is not everything, $s_p < S$: if type S pooled, its loss would be $\Pi(S)$, whereas announcing $\gamma^R S$ costs it at most $J(S, \gamma^R S; S) = \Pi(S) - S^2\varphi(\gamma^R) < \Pi(S)$ under any belief, by the argument of Step 1, since $\varphi(\gamma^R) > 0$ for every $\phi < C - \bar{c}$ (Lemma 8). Forecasts above s_p are read as themselves, so by Steps 1–3 their schedule is the top trajectory on $[s_p, S]$, which requires $s_p \geq s_0$. The boundary forecast is pinned by indifference: if $\Phi(s_p) > 0$, a forecast just below s_p gains by announcing $G(s')$ for s' just above s_p (its loss tends to $\Pi(s_p) - \Phi(s_p)$, its equilibrium loss to $\Pi(s_p)$); if $\Phi(s_p) < 0$, a forecast just above s_p gains by announcing zero (its loss tends to $\Pi(s_p)$, its equilibrium loss to $\Pi(s_p) - \Phi(s_p)$). Hence $\Phi(s_p) = 0$, and by the construction $s_p = s_b$. The announcement strategy is therefore G^*; payoff-irrelevant off-path beliefs may vary.

Homogeneity of s_b/S. By Step 6 the schedule for support S is $G_S(s) = SG_1(s/S)$, so $u_S(s) = u_1(s/S)$ and $\Phi_S(s) = S^2\Phi_1(s/S)$; the root s_b therefore scales with S, and s_b/S is the

root of Φ_1, a function of ϕ (and the calibration) only. This completes (iii).

The nowcast. With $s = 0$ and private v_1, the loss is $\frac{C}{2}[v_1+\beta(\phi\eta\widehat{v}_1+g)]^2+\frac{\beta A}{2}[(\phi-C)\eta v_1+g]^2$. The boundary condition of Step 1 holds with $\gamma_\eta^R\bar{v}$ in place of $\gamma^R S$: for the deviation $g' = \gamma_\eta^R\bar{v}$ and any belief in $[-\bar{v},\bar{v}]$, $|\bar{v}+\beta(\phi\eta\widehat{v}'+g')| \le \bar{v}+\beta(\phi\eta\bar{v}+g')$ because $\gamma_\eta^R > -1/\beta$; this inequality reads $\beta[A(\phi-C)+\beta C\phi]\eta < A$, and since $A(\phi-C) \le 0$ it holds whenever $\beta^2 C\phi\eta < A$, that is, for every $\phi \le C$ and $\eta < A/(\beta^2 C)$, which at the baseline calibration is $\eta < 21.5$. The first-order condition along a separating schedule is again a first-order equation in the inverse schedule, homogeneous of degree one, whose ratio equation has the roots of (B.24) as rest points and γ_η^R as its pole, so Steps 2–4 and 6 go through with $\hat{\phi}(\eta)$ in place of $\hat{\phi}$, the trajectory converging to the root nearest γ_η^R (Step 7 of the proof of Lemma 4 places γ_η^R outside the interval between the roots). Property (P1) fails for the nowcast, because the shock enters the expectations term, so the single-crossing argument of Step 5(a) and Lemma 9 do not transfer verbatim; the replication package verifies global incentive compatibility of the bounded nowcast reading directly, on a grid of announcements with the equilibrium beliefs, at six (ϕ,η) pairs in the revelation region below the threshold, including one with $\gamma_\eta^R < 0$ (relative gaps below 10^{-6}).

Supplementary Messages: Lemma 10 and Its Proof

The announcement format of Section 3.2 lets the central bank communicate only through the enforceable promise. The lemma below, cited there, enlarges the format by a separate message and shows that the enlargement supports no fully revealing outcome the promise alone does not and, within the equilibrium class of Proposition 3, can only raise the loss.

Lemma 10 (Supplementary messages and equilibrium welfare)**.** *Maintain the pure-forecast, assured-delivery assumptions of Proposition 3, with $0 < \phi < C - \bar{c}$. Enlarge the announcement space to $(g, m) \in \mathbb{R}\times\mathcal{M}$, where g is the enforceable departure and m is an additional public message that enters neither the standing rule nor the sanction. Every pair (g, m) is available to every type.*

(i) The enlarged and scalar announcement spaces support the same fully revealing pure-strategy equilibrium allocations.

(ii) Within the symmetric monotone pure-strategy equilibrium class satisfying D1, restricting communication to the enforceable promise is without loss for welfare. The scalar-message equilibrium of Proposition 3 weakly lowers every type's conditional expected loss relative to any equilibrium in the enlarged announcement space. Expected loss is strictly lower if supplementary messages change posterior means within a positive-probability group making the same promise.

For the enlarged game, symmetry requires an odd promise schedule and equal equilibrium losses for opposite forecasts; monotonicity concerns the promise schedule, and D1 applies to deviations in the full announcement (g, m).

Proof. Write

$$e(g, m) = g + \phi\widehat{s}(g, m), \qquad \widehat{s}(g, m) = \mathbb{E}[s \mid g, m],$$

and, with a slight abuse of notation, let $J(s; g, m) \equiv J\big(s, g; \widehat{s}(g, m)\big)$. Up to terms independent of the announcement, the loss of type s is

$$J(s; g, m) = \frac{C\beta^2}{2} e(g, m)^2 + \frac{\beta A}{2}\big[g - (C - \phi)s\big]^2.$$

The omitted terms are unchanged by the communication format.

Step 1: monotonicity and supplementary messages. Let $(G(s), m(s))$ be a pure-strategy equilibrium. For $s < t$, add the incentive constraints under which each type prefers its own announcement to the other's. The expectations terms cancel, leaving

$$\beta A(C - \phi)(t - s)\big[G(t) - G(s)\big] \geq 0.$$

Thus G is nondecreasing.

At a fixed promise g, all types rank supplementary messages identically: only the first term of J depends on m. Consequently, all messages used with that promise must induce the same squared expected-inflation level. Otherwise, any type using a more costly message would prefer a cheaper one.

Step 2: fully revealing allocations. In a fully revealing equilibrium,

$$e(G(s), m(s)) = \phi s + G(s).$$

If $G(s) = G(t) = g_0$ for some $s < t$, monotonicity implies $G(u) = g_0$ throughout $[s, t]$. Step 1 would then require $(\phi u + g_0)^2$ to be constant on that interval, which is impossible because $\phi > 0$. Hence G is strictly increasing: the promise alone identifies the forecast.

Support this allocation in the scalar game by assigning posterior mean s to each on-path promise $G(s)$. For any off-path promise g, use the belief assigned to (g, m_0) in the enlarged equilibrium, for a fixed admissible message m_0. Every scalar deviation then reproduces a deviation available in the enlarged game, so none is profitable. Conversely, any scalar equilibrium extends to the enlarged game by having every type send a fixed supplementary message and letting beliefs ignore that message, on and off path. This proves the first statement.

Step 3: the form of a symmetric D1 equilibrium. Consider a symmetric equilibrium of the enlarged game. A positive promise can be chosen only by positive forecasts. All supplementary messages accompanying that promise therefore induce positive expected inflation. By Step 1, their expected-inflation levels, and hence their posterior means, must coincide. Iterated expectations then imply that this common posterior mean equals the mean of the entire group making that promise.

The argument of Lemma 9 therefore continues to exclude any positive-promise pool. To see why the additional message does not alter that argument, consider a pool $[s_a, s_c]$ making a promise $g_p > 0$, with posterior mean $\mu_p > s_a$. A deviation to $(g_p - \epsilon, m)$ has the same delivery consequences for every supplementary message m. If it is on path, its posterior is supported below the pool. If it is off path, the D1 comparison is governed by

$$c(s; g) = \frac{2}{C\beta^2}\left[J^{\text{eq}}(s) - \frac{\beta A}{2}\big[g - (C - \phi)s\big]^2\right],$$

exactly as in (A.99). Within the pool, $c(s; g_p - \epsilon)$ is strictly decreasing in s. Types above the pool are dominated in this comparison by its lower endpoint, using their option to imitate a pooled announcement. Symmetry excludes negative types whose positive mirrors have a larger strict-gain set. Thus D1 restricts the posterior mean to $[-s_a, s_a]$. For sufficiently small ϵ, the lower endpoint strictly benefits: its delivery loss changes continuously, whereas its squared expectations term falls strictly. This contradicts equilibrium.

Negative-promise pools are excluded by symmetry. Therefore, any nondegenerate group sharing a promise is a central interval $[-p, p]$ making the promise zero. All types outside that interval separate and are correctly identified.

The central pool cannot contain every type. Its members incur at least

$$\Pi(s) \equiv \frac{\beta A}{2}(C - \phi)^2 s^2.$$

As established in the proof of Proposition 3 (the argument of Step 1 with Lemma 8), type S can announce $\gamma^R S$ and obtain loss strictly below $\Pi(S)$ under every admissible belief when $\phi < C - \bar{c}$.

The separating tails have the same incentive constraints, top boundary condition, and differential equation as in the scalar game. They therefore follow the same unique top-pinned trajectory, denoted by G below.

Step 4: supplementary disclosure in the central pool. By Step 1, every supplementary message used with promise zero induces the same squared expected-inflation level. Let

its contribution to loss be $Z \geq 0$. Every type in the central pool then incurs

$$J^{\text{eq}}(s) = \Pi(s) + Z.$$

If $Z = 0$, every such message induces posterior mean zero. If supplementary messages induce different posterior means, then $Z > 0$.

Along the separating trajectory, define

$$\Phi(s) = \Pi(s) - J(s, G(s); s),$$

as in Lemma 8. Incentive compatibility at the boundary between pooling and separation requires

$$\Pi(p) + Z = J(p, G(p); p), \qquad \text{hence} \qquad \Phi(p) = -Z.$$

Indeed, if either side of the first equality were strictly smaller, types sufficiently close to the boundary on the other side would profitably imitate it.

For $\phi \leq \hat{\phi}$, the separating trajectory satisfies $\Phi(s) > 0$ for every $s > 0$: by Step 4 of the proof of Proposition 3 its ratio $u(s) = G(s)/s$ lies in $[\gamma_2, \gamma^R]$, where φ is positive by Lemma 8. Thus no positive-width central pool can satisfy the boundary condition. The equilibrium allocation is the scalar separating allocation.

Now let $\hat{\phi} < \phi < C - \bar{c}$. The scalar equilibrium has central pooling boundary s_b, characterized by $\Phi(s_b) = 0$. Step 8 of the proof of Proposition 3 establishes that, on the relevant separating trajectory, $\Phi(s) < 0$ below s_b and $\Phi(s) > 0$ above it. Consequently,

$$p \leq s_b,$$

with strict inequality whenever $Z > 0$.

Step 5: welfare comparison. Compare the enlarged equilibrium with the scalar equilibrium. For $|s| \leq p$, the scalar equilibrium lowers loss from $\Pi(s) + Z$ to $\Pi(s)$. For $p < |s| < s_b$, the enlarged equilibrium separates, but

$$J(|s|, G(|s|); |s|) = \Pi(s) - \Phi(|s|) > \Pi(s),$$

which is the loss in the scalar equilibrium. For $|s| > s_b$, both equilibria use the same separating trajectory and yield the same loss. Boundary types satisfy the corresponding weak inequalities by indifference.

Thus the scalar equilibrium weakly lowers every type's conditional expected loss. The improvement in expected loss is strict whenever supplementary disclosure generates $Z > 0$

in a positive-probability central pool.

Finally, the scalar equilibrium remains available under the enlarged format, including under D1. At unused pairs with a given promise, the D1 comparisons depend only on that promise and the equilibrium loss function. Scalar D1 beliefs can therefore be retained at off-path promises. At promises already used, the original posterior is supported on types making that promise; their incentive constraints make them maximizers of the same deviation-gain comparison. Retaining that posterior at unused supplementary messages supports the scalar allocation. This completes the welfare comparison.

Remark 5 (The endpoints). The lemma assumes $0 < \phi < C - \bar{c}$. At the two endpoints the same conclusion is immediate from results already stated. Under the strict target, $\phi = 0$, the truthful promise is dominant whatever the public reads into an accompanying message, because the reading drops out of the loss (Proposition 4); the enlarged format supports the same outcome. Under the ideal rule, $\phi = C$, no type gains from being read as holding a larger forecast: revealing the forecast raises expected inflation today without improving the period-2 response (Corollary 2), so every type prefers an uninformative supplementary message and the promise alone loses nothing.

Proof of Proposition 4

Proof. Under $\phi = 0$, (A.94) reads $e_1(g) = g$ for every belief $\widehat{s}$, so the loss of every type is a function of its type and its announcement alone: $\frac{C}{2}(\beta g)^2 + \frac{\beta A}{2}(g - Cs)^2$ for the pure forecast, and $\frac{C}{2}(v_1+\beta g)^2+\frac{\beta A}{2}(g-C\eta v_1)^2$ for the nowcast. Each is a strictly convex quadratic in g whose minimizer was computed in Step 9 of the proof of Lemma 4: $ABC\, s$ and $g_0^R(v_1)$ respectively. Since the minimizer does not depend on the belief, it is a dominant announcement for every type. In a perfect Bayesian equilibrium each type plays a best response to the belief map, hence its dominant announcement; the strategy is therefore pinned, and with it the on-path beliefs. The strategy is one-to-one (for the nowcast, unless $\eta = 1/A$, where the promise of every type is zero and the announcement is uninformative by the planner's own choice, Corollary 8), so the equilibrium is the truthful reading, whatever the support and whatever beliefs are attached off the path.

The strict target is the only rule in the revelation region with this property. For $0 < \phi < C - \bar{c}$ the best response to a reading of slope γ is (43), $b(\gamma) = A(C-\phi)/[A+\beta C(1+\phi/\gamma)^2]$, whose derivative in γ is $2\beta C\phi(1+\phi/\gamma)A(C-\phi)/\{\gamma^2[A+\beta C(1+\phi/\gamma)^2]^2\}$; because $\phi > 0$ and $C - \phi > \bar{c} > 0$ it vanishes only at $\gamma = -\phi$, so b is not constant: the central bank's best promise varies with the reading, and no announcement is dominant. (At the ideal rule, $\phi = C$, the best promise is zero under every reading, but there the forecast does not enter

the announcement problem and society prefers it kept private, Corollary 2.)

Proof of Proposition 6

Proof. (i) With the correction. By Proposition 2, under assured delivery the correction for a revealed pure forecast attaches to the report r the burden $\mathcal{A}_R^{\phi}(r) = -\frac{1}{2}\phi\, AB(\beta C)^2\, r^2$; in announcement units, $r = g/\gamma^R$, this is $-\frac{1}{2}h_c g^2$ with $h_c \equiv \phi\, AB(\beta C)^2/(\gamma^R)^2$. The corrected loss is $J_c(s, g; \widehat{s}) = J(s, g; \widehat{s}) - \frac{1}{2}h_c g^2$; (P1) and (P2) are unaffected. *The top is pinned at the cap.* Under the belief S, $\partial_g J_c(S, g; S) = C\beta^2(\phi S + g) + \beta A[(\phi - C)S + g] - h_c g$; at $g = \gamma^R S$ the first two terms sum to $\beta S[C\beta(\phi + \gamma^R) + A(\phi - C + \gamma^R)] = 0$ by (A.96), so $\partial_g J_c(S, \gamma^R S; S) = -h_c \gamma^R S < 0$: the corrected loss is decreasing at the cap. If $J_c(S, \cdot; S)$ is convex ($\partial_{gg} J_c = C\beta^2 + \beta A - h_c > 0$), its unconstrained minimizer lies above the cap and the constrained minimizer is $\bar{g}$. If it is concave, its minimum over $[-\bar{g}, \bar{g}]$ is at an endpoint, and

$$\begin{aligned} J_c(S, -\bar{g}; S) - J_c(S, \bar{g}; S) &= -2C\beta^2\phi S\bar{g} + 2\beta A(C - \phi)S\bar{g} \\ &= 2\beta S\bar{g}\left[A(C - \phi) - \beta C\phi\right] = 2\beta\Lambda\gamma^R S\bar{g} > 0, \end{aligned}$$

so again the constrained minimizer is $\bar{g}$. The argument of Step 1 (worst belief S) then gives $G(S) = \bar{g}$ in every separating equilibrium. *The truthful schedule is the unique solution.* The corrected first-order condition along a separating schedule is (A.102) with $h_c G/\beta$ added to the denominator,

$$\frac{ds}{dg} = \frac{A(C - \phi)s - Ag - \beta C(\phi s + g) + (h_c/\beta)g}{\beta C\phi(\phi s + g)},$$

regular at $(\bar{g}, S)$ as before. The line $s = g/\gamma^R$ solves it: substituting, the numerator becomes $g[(A(C - \phi) - \beta C\phi)/\gamma^R - \Lambda + h_c/\beta] = (h_c/\beta)g$ and the denominator $\beta C\phi(\phi + \gamma^R)g/\gamma^R$, so the right side equals $h_c\gamma^R/[\beta^2 C\phi(\phi + \gamma^R)] = \phi\, AB(\beta C)^2\gamma^R/[(\gamma^R)^2\beta^2 C\phi\, ABC] = 1/\gamma^R$, using $\phi + \gamma^R = ABC$. By uniqueness of the initial-value problem the truthful reading is the only separating equilibrium. *It is an equilibrium.* Every admissible announcement is on the path, so only on-path deviations need checking: the corrected loss of type s reporting r, $J(s, \gamma^R r; r) - \frac{1}{2}h_c(\gamma^R r)^2$, is a quadratic in r with second derivative

$$\begin{aligned} C\beta^2(\phi + \gamma^R)^2 + \beta A(\gamma^R)^2 - \phi\, AB(\beta C)^2 &= AB(\beta C)^2(ABC - \phi) + \beta A(\gamma^R)^2 \\ &= \beta A\gamma^R\left[\beta BC^2 + \gamma^R\right] > 0, \end{aligned}$$

and stationary at $r = s$ by Proposition 2; so truth is each type's global optimum. Nothing in this argument used $\phi < \hat{\phi}$: the corrected ratio equation $F_c(u) = u$, with $F_c(u) \equiv \beta C\phi(\phi + u)/[\Lambda(\gamma^R - u) + h_c u/\beta]$, has $u = \gamma^R$ as a root by construction for every ϕ in the revelation region, band included.

(ii) Without the correction. Under belief S the minimizer of $J(S, \cdot; S)$ is $\gamma^R S \geq \bar{g}$, so by convexity the constrained minimizer is $\bar{g}$, and Step 1 pins $G(S) = \bar{g}$. The solution of (A.103) from $(\bar{g}, S)$ is unique (Step 3). Let $\bar{u} = \bar{g}/S$. In the phase equation (A.104): for $\bar{u} \in (\gamma_2, \gamma^R)$ the trajectory descends to γ_2 as in Step 4; for $\bar{u} = \gamma_2$ or $\bar{u} = \gamma_1$ it is the constant solution, the reading of a single slope; for $\bar{u} \in (\gamma_1, \gamma_2)$ one has $Q(u) < 0$ and $\gamma^R - u > 0$, so $du/dt < 0$ and, followed downward in s, u rises to γ_2 (at $\phi = \hat{\phi}$ the roots coincide and this case is empty); in all these cases $G' = F(u) > 0$ and the schedule reaches the origin, with $u(s) \in [\gamma_1, \gamma^R]$ throughout. For $\bar{u} < \gamma_1$, $Q > 0$ and the trajectory descends through zero at a positive forecast, and no separating equilibrium exists, as in Step 4. Global incentive compatibility of the schedule follows from Step 5 with two changes: announcements above $\bar{g}$ are inadmissible, and Lemma 8 applies because $\varphi > 0$ on $[\gamma_1, \gamma^R]$.

Misjudged Readings

The following remark quantifies the cost of acting on a misjudged reading (Section 4.2.3).

Remark 6 (Robustness to the reading). Under a coarse rule the central bank's best promise depends on how the public reads it, and so does the cost of misjudging that reading. If the central bank acts on a reading of slope γ' while the public holds γ, its loss exceeds its best-response loss by $\frac{1}{2} J_{gg}\, b'(\gamma)^2(\gamma' - \gamma)^2 s^2$ to second order, where $J_{gg} = C\beta^2(1 + \phi/\gamma)^2 + \beta A$ is the curvature of (42) and b' the derivative of (43); if the reading is uncertain to the central bank, the expected excess loss is the same expression with the mean squared error of its estimate of γ in place of $(\gamma' - \gamma)^2$. At the near-Ramsey reading b' is 0.07 at $\phi = 0.3$, 0.26 at $\phi = 0.5$, 0.71 at $\phi = 0.6$, and reaches one at $\hat{\phi}$, where the two readings merge and a misjudged reading passes one for one into the promise. Under the strict target $b' \equiv 0$: the central bank never has to judge how it will be read. The announcement correction is exposed in the same way through the model-specific coefficient of Corollary 3(i).

A.16 Welfare across Standing Rules

Proof of Proposition 5. In parts (i) and (ii) the forecast is s, the innovation is zero, and a promise g that the public reads correctly moves the expectation to $e_1 = \phi s + g$; with

sanctions socially free, the social loss of type s is then the central bank's own,

$$\mathcal{L}(s,g) = \frac{C}{2}\big[\beta(\phi s + g)\big]^2 + \frac{\beta A}{2}\big[(\phi - C)s + g\big]^2,$$

the current stabilization loss after expectations have moved plus the discounted loss from the delivered deviation, as in (42) with $\widehat{s} = s$. For fixed s this is strictly convex in g with the unique minimizer $g = \gamma^R s$, $\gamma^R = B[A(C-\phi) - \beta C\phi]$ (Proposition 1).

Part (i). At $g = \gamma^R s$,

$$\phi + \gamma^R = B\big[(A+\beta C)\phi + A(C-\phi) - \beta C\phi\big] = ABC, \qquad (\phi - C) + \gamma^R = ABC - C = -\beta BC^2,$$

using $1 - AB = \beta CB$. Substituting,

$$\mathcal{L}(s,\gamma^R s) = \frac{s^2}{2}\Big[C\beta^2A^2B^2C^2 + \beta A\beta^2B^2C^4\Big] = \frac{s^2}{2}\beta^2AB^2C^3(A+\beta C) = \frac{s^2}{2}ABC(\beta C)^2 = L^T s^2,$$

which does not involve ϕ. Since $\gamma^R s$ minimizes $\mathcal{L}(s,\cdot)$ for every s, no reading of guidance attains less, so L^T is the lowest loss attainable with guidance in the revelation region. Under the strict target the truthful promise $\gamma^R s$ with $\gamma^R = ABC$ is a dominant strategy (Proposition 4), so the strict target attains L^T with no correction. Without a correction and $0 < \phi < C - \bar{c}$, the equilibrium of Proposition 3 reveals types on its separating portion, where the loss is $\mathcal{L}(s, G(s))$, which exceeds $L^T s^2$ wherever $G(s) \neq \gamma^R s$. Below the threshold, G solves the separating first-order condition, and $G(s) = \gamma^R s$ on an interval would make γ^R a self-confirming slope; but γ^R is not a root of (44) for $\phi > 0$, because the quadratic evaluated at γ^R equals $\beta C\phi(\phi + \gamma^R) = \beta C\phi ABC > 0$. Hence $G(s) \neq \gamma^R s$ on a set of forecasts of positive measure, and the expected loss is strictly above L^T. In the collapse band, forecasts with $|s| \leq s_b$ issue $g = 0$ and incur $\frac{1}{2}\beta A(C-\phi)^2 s^2 > L^T s^2$ by part (ii), and $s_b > 0$.

Part (ii). With no guidance the expectation does not move, $e_1 = 0$, and the rule delivers ϕs against the ideal Cs, so the loss is $\frac{1}{2}\beta A(C-\phi)^2 s^2$. At $\phi = C - \bar{c}$ this equals $\frac{1}{2}\beta A\bar{c}^2 = \frac{1}{2}\beta AC^2B\beta C = L^T$ by $\bar{c} = C\sqrt{B\beta C}$ (Proposition 1); it is strictly decreasing in ϕ on $[0, C]$, so it lies above L^T for $\phi < C - \bar{c}$ and below it on $(C - \bar{c}, C]$, and it is zero at $\phi = C$.

Part (iii). With assured delivery the delivered center is $\phi v_2 + g$ and the promise is chosen before the innovation $\varepsilon = v_2 - s$ is realized, so the innovation enters the delivered deviation from the ideal as $(\phi - C)\varepsilon$ under every rule and every promise; its contribution to the expected loss is $\frac{1}{2}\beta A(C-\phi)^2\sigma^2_{2|1}$, largest at $\phi = 0$ and zero at $\phi = C$.

The nowcast. With $s = 0$, private v_1, and $\Gamma_\phi(\eta) < 0$, the per-message loss is the one

computed in the proof of Lemma 16: for a message with posterior mean m and posterior variance w of the nowcast, its minimum over the promise is $\frac{C}{2}w + \frac{1}{2}ABCc_\eta^2 m^2 + \frac{\beta A}{2}(C - \phi)^2 \operatorname{Var}(v_2 \mid q)$. Under truthful guidance $w = 0$ and $m = v_1$, so the loss is $\frac{1}{2}ABCc_\eta^2$ per unit of $\operatorname{Var}(v_1)$, which does not involve ϕ, plus the innovation term with $\operatorname{Var}(v_2 \mid v_1) = \sigma_{2|1}^2$; because the expected loss is affine in $\operatorname{Var}(m)$ with the negative coefficient $\frac{C}{2}\Gamma_\phi(\eta)$, no message attains less. With no guidance $m = 0$, $w = \operatorname{Var}(v_1)$, and the promise is zero, so the loss is $\frac{C}{2}\operatorname{Var}(v_1) + \frac{\beta A}{2}(C - \phi)^2 \operatorname{Var}(v_2)$ with $\operatorname{Var}(v_2) = \eta^2 \operatorname{Var}(v_1) + \sigma_{2|1}^2$, which is $\frac{C}{2} + \frac{1}{2}\beta A(C - \phi)^2\eta^2$ per unit of $\operatorname{Var}(v_1)$ plus the innovation term.

A.17 Implementation of the Coarse-Rule Mandate

This subsection carries the construction summarized in Section 4.2.6.

The mandate that restores guidance under a coarse rule uses one bounded sanction with two components in its probability: an announcement component fixed before delivery, and an outcome component that responds to the inflation miss. We state it for the forecast-bundling nowcast case ($s = 0$ almost surely), so that all type-conditional expectations condition on the central bank's full private information v_1; the same construction applies to the pure forecast when the nowcast is fixed and public, with the type and report indexing that forecast. It elicits one private component at a time.

Let $g_\phi(r)$ be a continuously differentiable, one-to-one candidate announcement schedule on a compact report interval containing zero; the announcement format admits its image as the announcement range. The associated public expectation must equal expected *delivered* inflation when the true type is r:

$$e_\phi(r) = \mathbb{E}\left[Cv_2 + \Delta((\phi - C)v_2 + g_\phi(r)) \mid v_1 = r\right]. \tag{A.105}$$

With no unanticipated innovation, the inverse $\Delta^a(\cdot)$ of the traction function, (35), matches the desired Ramsey center exactly. With innovations, calibrating the announcement at the mean shock does not in general match the mean delivered center; Appendix A.12 gives the conditional-mean calibration.

The loss of a central bank with true type v_1 that selects the announcement indexed by r is

$$\mathcal{J}_\phi^{\mathrm{del}}(v_1, r) = \frac{C}{2}[v_1 + \beta e_\phi(r)]^2 + \beta\mathbb{E}\left[\mathcal{V}_C((\phi - C)v_2 + g_\phi(r)) \mid v_1\right], \tag{A.106}$$

up to terms independent of the report. This includes the expected delivery sanction: a central bank that cares about that sanction takes it into account when announcing. Define

the additional expected announcement burden by

$$\mathcal{A}_I(g_\phi(r)) = -\int_0^r \partial_r \mathcal{J}_\phi^{\mathrm{del}}(x, x)\, \mathrm{d}x. \tag{A.107}$$

Its slope offsets the reporting distortion of the type that should issue each announcement; ∂_r differentiates the report argument before both arguments are evaluated at x. Under assured delivery, no report-dependent delivery burden, and the rule-constrained Ramsey candidate schedule, it reduces to the correction $\mathcal{A}_R^\phi$ of Proposition 2; with a noisy audit, it also accounts for the announcement incentives created by the expected delivery sanction. The closed form for the ideal-rule case is given in Appendix A.18.

Let $\underline{\mathcal{A}_I}$ and $\overline{\mathcal{A}_I}$ be the minimum and maximum of $\mathcal{A}_I$ over the candidate announcement range. The stake allocated to announcement incentives, its normalized probability, and the coarse-rule audit are

$$\Psi_I = (\overline{\mathcal{A}_I} - \underline{\mathcal{A}_I})/\beta, \tag{A.108}$$

$$\rho_I(g) = [\mathcal{A}_I(g) - \underline{\mathcal{A}_I}]/(\beta\Psi_I), \tag{A.109}$$

$$\Psi_A = \Psi + \Psi_I, \tag{A.110}$$

$$\rho_A(g, z) = \frac{\Psi}{\Psi_A}\rho_C(z) + \frac{\Psi_I}{\Psi_A}\rho_I(g). \tag{A.111}$$

When $\Psi_I = 0$, set $\rho_I = 0$. The mandate imposes one sanction of size Ψ_A with probability ρ_A, a valid probability because it is a convex combination. Its expected discounted burden is

$$\beta\Psi_A \mathbb{E}[\rho_A(g, a + \varepsilon_2) \mid g, a] = \beta\mathcal{A}_C(a) + \mathcal{A}_I(g) - \underline{\mathcal{A}_I}. \tag{A.112}$$

Once the central bank has announced g, the second term is fixed: it changes the incentive to select the announcement but not the marginal incentive to deliver it. The following lemma states the implementation result cited in Section 4.2.6 with the construction in place. Maintain that $v_1 \in [-\bar{v}, \bar{v}]$, that the announcement format admits the image of a continuously differentiable, one-to-one schedule g_ϕ, that the reporting loss is continuously differentiable, that delivery is uniquely optimal, and that public expectations satisfy (A.105).

Lemma 11 (Coarse accountability delegation under the coarse-rule audit)**.** *Suppose that either (a), for every type v_1, the corrected reporting loss*

$$r \longmapsto \mathcal{J}_\phi^{\mathrm{del}}(v_1, r) + \mathcal{A}_I(g_\phi(r)) \tag{A.113}$$

is strictly convex, or (b) $\mathcal{J}_\phi^{\text{del}}$ *is twice continuously differentiable and* $\partial_{v_1 r}\mathcal{J}_\phi^{\text{del}}(v_1, r) < 0$ *throughout the admissible type–report rectangle. If society's sanctioning capacity is at least* Ψ_A*, the coarse-rule audit supports a separating equilibrium with truthful report* $r = v_1$ *and delivery* $\theta_2 = Cv_2 + \Delta((\phi - C)v_2 + g_\phi(v_1))$*. When the candidate schedule implements the Ramsey center with no unanticipated innovation, that center is delivered exactly. Under the strict target with assured delivery and no report-dependent delivery burden,* $\mathcal{A}_I$ *is constant and no announcement correction is required when the candidate schedule is the rule-constrained Ramsey promise.*

Proof. See Appendix A.18.

The construction identifies the information and capacity needed to restore guidance. The designer must know the candidate schedule g_ϕ, the traction function, the corrected loss over all admissible types and reports, and the total stake Ψ_A. Conditions (a) and (b) must hold globally, not merely at the truthful report. In the assured-delivery pure-forecast case, the format also makes the truthful reading the unique separating equilibrium (Proposition 6).

Under assured delivery, the reporting condition holds wherever revelation is strictly desirable: $\beta B\mathcal{K}_\phi = -C\Gamma_\phi + ABC\beta\eta\phi c_\eta > 0$ when $\Gamma_\phi < 0$ (Proposition 2). Under the ideal rule, the reporting gap is $\frac{\beta BC^2(1-A\eta)}{2}(r - v_1)^2$ throughout the revelation region, provided the noisy audit implements the candidate center. Away from the ideal rule, delivery optimality alone is insufficient. At $\phi = 0.2$, $\eta = 0.2$, and $\Psi = 18$, the corrected loss is locally concave at truth. A stake of $\Psi = 35$ restores positive local curvature; the proof below establishes global convexity for all sufficiently large stakes.

The knowledge requirement is explicit under assured delivery. The nowcast correction is $\mathcal{A}_R^\phi(r) = -\frac{1}{2}ABC\beta\eta\phi c_\eta r^2$, so the designer must know the composite $\phi\eta c_\eta ABC$ that combines the rule's response to the forecast with stabilization tradeoffs. With noisy delivery, (A.107) also depends on the traction function and expected delivery sanction. Both components use the available sanctioning capacity. The discussion of report-dependent delivery stakes in Appendix A.18 explains why the additive construction leaves delivery incentives unchanged; Corollary 3 states the cross-model requirement.

A.18 Joint Implementation

This subsection proves Lemma 11 in the form stated in Appendix A.17. Throughout, $a^*(\Delta^a)$ is the central bank's delivery best response to a pre-shortfall deviation Δ^a, and $\Delta^a(\Delta^\star)$ is the inverse of the traction function for a desired delivered deviation, (35); the shortfall associated with that target is $a^*(\Delta^a(\Delta^\star))$, on a responsive branch that is globally optimal.

Proof of Lemma 11. By (A.108)–(A.111), both component probabilities lie in $[0, 1]$ and their

weights sum to one. Conditional on announcement g and hidden shortfall a,

$$\beta\Psi_A\mathbb{E}[\rho_A(g, a+\varepsilon_2) \mid g, a] = \beta\Psi\mathbb{E}[\rho_C(a+\varepsilon_2) \mid a] + \beta\Psi_I\rho_I(g) \tag{A.114}$$

$$= \beta\mathcal{A}_C(a) + \mathcal{A}_I(g) - \underline{\mathcal{A}}_I. \tag{A.115}$$

The last two terms are fixed once g is chosen. Thus the central bank's delivery problem is unchanged and its unique delivery action remains the one used in (A.105).

Write $H(r) = \mathcal{A}_I(g_\phi(r))$. By (A.107),

$$H'(r) = -\partial_r \mathcal{J}_\phi^{\mathrm{del}}(r, r). \tag{A.116}$$

Therefore the derivative of the corrected reporting loss of a true type v_1 is

$$\partial_r\Big[\mathcal{J}_\phi^{\mathrm{del}}(v_1, r) + H(r)\Big] = \partial_r \mathcal{J}_\phi^{\mathrm{del}}(v_1, r) - \partial_r \mathcal{J}_\phi^{\mathrm{del}}(r, r). \tag{A.117}$$

At $r = v_1$ the two terms coincide, so truth is stationary for every interior type. This establishes what the announcement correction does locally.

There are two ways to turn this local condition into a global incentive condition. Under condition (a) of the lemma, strict convexity makes the stationary report the unique global minimizer. Under condition (b), fix a true type v_1. If $r < v_1$, the fundamental theorem of calculus gives

$$\partial_r \mathcal{J}_\phi^{\mathrm{del}}(v_1, r) - \partial_r \mathcal{J}_\phi^{\mathrm{del}}(r, r) = \int_r^{v_1} \partial_{v_1 r}\mathcal{J}_\phi^{\mathrm{del}}(x, r)\, \mathrm{d}x < 0. \tag{A.118}$$

If $r > v_1$, the same expression equals

$$-\int_{v_1}^{r} \partial_{v_1 r}\mathcal{J}_\phi^{\mathrm{del}}(x, r)\, \mathrm{d}x > 0. \tag{A.119}$$

Thus the corrected loss falls as the report approaches truth from below and rises after the report passes truth. Truth is again the unique global minimizer. Endpoint types follow from the corresponding one-sided inequalities.

The same calculation also identifies the local obstruction. Differentiating (A.117) once more and evaluating at truth yields

$$\partial_{rr}\Big[\mathcal{J}_\phi^{\mathrm{del}}(v_1, r) + H(r)\Big]\Big|_{r=v_1} = -\partial_{v_1 r}\mathcal{J}_\phi^{\mathrm{del}}(v_1, v_1). \tag{A.120}$$

A positive cross-derivative at truth therefore makes the corrected objective locally concave

there. Delivery convexity alone cannot rule this out. The obstruction arises at the baseline calibration. With $v_2 = \eta v_1$ and the candidate $g_\phi(r) = (C - \phi)\eta r + \Delta^a(g^R(r))$, the chain rule gives, at the type $v_1 = 0$,

$$\partial_{rr}\Big[\mathcal{J}_\phi^{\text{del}}(v_1, r) + H(r)\Big]\Big|_{r=v_1=0} = \beta\Big\{A(C-\phi)\eta\big[(C-\phi)\eta\,\omega - BCc_\eta\big] - BC^2(A\eta - 1)\Big\}, \quad \text{(A.121)}$$

where $\omega = \Delta'(0)$ is the traction of Section 4.1.2; at $\omega = 1$ this is $\beta B\mathcal{K}_\phi(\eta)$, and at $\phi = C$ it is $\beta BC^2(1 - A\eta)$ for every ω. At $\phi = 0.2$ and $\eta = 0.2$ revelation is strictly desirable, $\Gamma_\phi(\eta) = -0.093$, and delivery is uniquely optimal, yet with $\Psi = 18$ and $\tau = \sigma_\pi$ the traction is $\omega = 0.528$ and the curvature is -0.061: truth is a local maximum of the corrected reporting loss. The curvature is increasing in ω and crosses zero at $\bar{\omega} = 0.680$, which the baseline audit reaches at $\Psi = 34.2$; at $\Psi = 35$ the curvature is $+0.002$. The large-capacity result below shows that it stays positive for every sufficiently large stake.

Since g_ϕ is one-to-one, unique truthful reporting allows the public to infer the type from the announcement. Its expectation is therefore (A.105). Delivery, reporting, and beliefs are mutually consistent, proving the separating-equilibrium claim. Exact matching of the Ramsey center in the no-innovation case follows by substituting $\Delta^a(\cdot)$. Under the strict target with assured delivery and no report-dependent delivery burden, the derivative vanishes by Proposition 2 when the candidate schedule is the rule-constrained Ramsey promise; H is then constant and $\Psi_I = 0$.

Ideal rule. For the ideal-rule candidate $g_C(r) = \Delta^a(g^R(r))$, with $e_C(r) = C\eta r + g^R(r)$ and the targets on a globally optimal responsive branch, the delivery argument is independent of the true current type. The announcement adjustment can then be written

$$\mathcal{A}_I(g_C(r)) = \mathcal{A}_R^C(r) + \frac{\beta A}{2}[g^R(r)]^2 - \beta\mathcal{V}_C(\Delta^a(g^R(r))) + \text{const.} \quad \text{(A.122)}$$

Using the delivery value function, this is equivalently

$$\mathcal{A}_I(g_C(r)) = \mathcal{A}_R^C(r) - \beta\mathcal{A}_C(a^*(\Delta^a(g^R(r)))) + \text{const.} \quad \text{(A.123)}$$

For the ideal rule the global reporting condition can be verified without assuming affine traction. Let $g_C(r) = \Delta^a(g^R(r))$ implement the desired deviation on a branch that is globally optimal at delivery. Consistent public expectations are

$$e_C(r) = C\eta r + g^R(r) = BC(A\eta - 1)r. \quad \text{(A.124)}$$

The delivery value associated with report r is independent of the central bank's true nowcast. Consequently, after the announcement correction has removed the derivative evaluated at type r, the remaining reporting derivative is entirely the current-stabilization difference:

$$\partial_r \left[\mathcal{J}_C^{\text{del}}(v_1, r) + H(r)\right] = C\beta e_C'(r)(v_1 - r) \tag{A.125}$$

$$= \beta BC^2(1 - A\eta)(r - v_1). \tag{A.126}$$

Integrating from the truthful report v_1 to an arbitrary report r gives

$$\begin{aligned} &\left[\mathcal{J}_C^{\text{del}}(v_1, r) + H(r)\right] - \left[\mathcal{J}_C^{\text{del}}(v_1, v_1) + H(v_1)\right] \\ &\qquad = \frac{\beta BC^2(1 - A\eta)}{2}(r - v_1)^2. \end{aligned} \tag{A.127}$$

Throughout the ideal-rule revelation region, $\eta < \eta^\star < 1/A$, so the coefficient is strictly positive. Truth is therefore the unique global minimizer. This argument uses the actual nonlinear delivery value and requires no affine traction function; nor is it tied to the linear Ramsey schedule, since for any one-to-one candidate implementing a differentiable expectation schedule $e_C(r)$ the corrected derivative is $C\beta(v_1 - r)e_C'(r)$, which points toward truth whenever $e_C' < 0$, and conversely makes truth a local maximum where $e_C'(v_1) > 0$. It also survives the addition of the independent private forecast s: by Lemma 6(i), s adds the same constant to the loss at every report.

In the audit-rebate expression (A.123), the first term corrects the rule-induced reporting distortion and the second removes the additional report incentive created by the expected delivery sanction. With assured delivery, no report-dependent delivery burden, and the rule-constrained Ramsey candidate schedule, the correction is $\mathcal{A}_R^\phi$ for any rule. With a coarse rule and nonlinear delivery, the integral (A.107) must instead be used.

Sufficiently large sanctioning capacity. The reporting condition cannot be deduced from delivery convexity for an arbitrary given noisy audit. There is, however, an existence result in the bounded nowcast case. Set $s = 0$ and suppose $v_2 = \eta v_1$, with $v_1 \in [-\bar{v}, \bar{v}]$, and $\Gamma_\phi(\eta) < 0$. Hold $\tau > 0$ and $\sigma_\pi > 0$ fixed. Then, for all sufficiently large finite delivery stakes Ψ, the candidate

$$g_\phi(r) = (C - \phi)\eta r + \Delta^a(g^R(r)) \tag{A.128}$$

is one-to-one, delivery is uniquely optimal for every admissible type–report pair, and the corrected reporting loss is strictly convex. Hence the coarse-rule audit implements the revelation benchmark whenever total sanctioning capacity covers this delivery stake and the associated announcement correction.

To prove the claim, write the expected delivery audit in (A.59) as $\mathcal{A}_C(a) = \Psi h(a)$, where

$$h(a) = 1 - \frac{\tau}{\sqrt{\tau^2 + \sigma_\pi^2}} \exp\left(-\frac{a^2}{2(\tau^2 + \sigma_\pi^2)}\right). \tag{A.129}$$

The function h has a unique global minimum at zero. Fix $0 < a_0 < \sqrt{\tau^2 + \sigma_\pi^2}$ and a bounded set of pre-shortfall deviations u. If $a^*(u)$ is any global delivery optimum, comparison with the feasible choice $a = 0$ gives

$$\Psi\big[h(a^*(u)) - h(0)\big] \le \frac{A}{2}u^2. \tag{A.130}$$

The right-hand side is uniformly bounded. Since $h(a) - h(0)$ is strictly positive away from zero, every global optimum must lie in $(-a_0, a_0)$ once Ψ is sufficiently large. On this interval $h''(a) > 0$, so the delivery objective is strictly convex there. Because every global optimum has already been shown to lie there, the global delivery optimum is unique.

Two further steps complete the argument. First, all delivery arguments that the reporting problem evaluates stay in one bounded set. For a target $t = g^R(r)$ on the compact report interval, the equation $\Psi h'(a) = -At$ that defines $\Delta^a(t)$ has, for Ψ large, a solution in $(-a_0, a_0)$ that is uniformly $O(1/\Psi)$, so the inverse input $t - a$ is uniformly bounded; hence so is $(\phi - C)\eta v_1 + g_\phi(r)$ over the type–report rectangle, and the comparison with $a = 0$ above applies on that fixed set. Second, the delivery-value derivatives, not only the candidate schedule, converge. Differentiating the first-order condition $A(u + a^*(u)) + \Psi h'(a^*(u)) = 0$ gives

$$(a^*)'(u) = -\frac{A}{A + \Psi h''(a^*(u))}, \qquad \Delta'(u) = \frac{\Psi h''(a^*(u))}{A + \Psi h''(a^*(u))}, \tag{A.131}$$

$$\Delta''(u) = -\frac{\Psi A^2 h'''(a^*(u))}{[A + \Psi h''(a^*(u))]^3}. \tag{A.132}$$

Because every optimum lies uniformly near zero, h'' is bounded away from zero there and h''' is bounded, so uniformly on bounded sets

$$a^*(u) \longrightarrow 0, \qquad \Delta(u) \longrightarrow u, \qquad \Delta'(u) \longrightarrow 1, \qquad \Delta''(u) \longrightarrow 0 \tag{A.133}$$

as $\Psi \to \infty$. By the envelope theorem $\mathcal{V}_C'(u) = A\Delta(u)$ and $\mathcal{V}_C''(u) = A\Delta'(u)$, the derivatives that enter the reporting loss, and these converge uniformly to Au and A; the constant $\Psi h(0)$ in the delivery value may diverge but is report-independent and irrelevant to curvature. The inverse traction function consequently converges to the identity, with its first derivative

converging to one and its second derivative to zero on the bounded target range, so g_ϕ and its first two derivatives converge uniformly to the assured-delivery candidate. The sufficient global-curvature bound of Lemma 3 is unnecessary here: it places an upper bound on Ψ, and the direct comparison with $a = 0$ gives unique global optima on the relevant bounded set.

Under assured delivery the corrected reporting curvature is $\beta B \mathcal{K}_\phi(\eta)$, which is strictly positive because $\Gamma_\phi(\eta) < 0$ by (A.88). Uniform convergence therefore implies that the corrected noisy-delivery reporting curvature is also strictly positive for all sufficiently large finite Ψ. The assured-delivery candidate has nonzero slope in this region; otherwise its corrected reporting curvature would be nonpositive. Hence g_ϕ is one-to-one for sufficiently large Ψ as well. Finally, on the compact report interval the announcement correction has bounded range, so the additional sanction requirement is finite. This proves the claim. It is an existence result for sufficiently large capacity; at a given delivery stake the reporting condition must be verified directly, as the example above shows.

If opacity is desired, an announcement format restricted to a public formula produces an uninformative announcement. The resulting expectation must still be computed from actual delivery; a noisy implementation need not reproduce the assured-delivery allocation under the uninformative announcement state by state. Under the ideal rule, selective opacity about the pure forecast follows from Lemma 6 when announcement best responses are unique.

Report-dependent delivery stake. A report-dependent stake $\Psi(g)$ or tolerance $\tau(g)$ changes the marginal expected sanction at delivery and hence the traction function. The coarse-rule audit instead holds its announcement component fixed after the announcement, so it changes reporting incentives and leaves marginal delivery incentives unchanged. This decomposition uses the linearity of the central bank's loss in expected sanctions.

A.19 Model Uncertainty

Proof of Corollary 3. Write $H(r) = \mathcal{A}_I(g_\phi(r))$ for the additional expected discounted sanction. The correction computed under $\widehat{\Omega}$ has $H'(r) = -w_\phi(r; \widehat{\Omega})$, giving (48). A common H' cancels every model's derivative exactly when those derivatives coincide. Integrating the common derivative constructs H; strict convexity makes the stationary report the unique optimum, and the coarse-rule audit normalization implements its bounded range when capacity permits.

Under assured delivery and the strict target, the public expects the observed promise

itself: $e_1(g) = g$. Conditional on the central bank's information, both the central bank and the rule-constrained Ramsey planner minimize, up to constants,

$$\frac{C}{2}(v_1 + \beta g)^2 + \frac{\beta A}{2}\mathbb{E}[(g - Cv_2)^2 \mid v_1, s]. \tag{A.134}$$

Its derivative is $\beta\{Cv_1 + g/B - AC(\eta v_1 + s)\}$ and its curvature is $\beta/B > 0$. The unique optimum is (49). Because delivery makes $e_1(g) = g$ regardless of the hidden state, the public need not decode that state to understand the promised center.

Corollary 3 concerns one mandate used across several models; the central bank and the public know the true model. This subsection gives the additional implications of its compatibility condition and an explicit example showing why the strict-target result requires assured delivery.

What the designer must know. Let I be the compact report interval and retain the same announcement schedule $g_\phi : I \to \mathbb{R}$ and obedience audit in every model, with $J_\Omega(v, r) = \mathcal{J}_\phi^{\mathrm{del}}(v, r \mid \Omega)$ the reporting loss under the model-consistent expectation of actual delivery. Continuity on the compact interval makes the range of H finite, and the coarse-rule audit implements that range when its additional stake, $(\sup_I H - \inf_I H)/\beta$, is available. Two models are equivalent for this correction when they have the same function w_ϕ, even if their primitives differ: evaluated at a model's own rule-constrained Ramsey candidate in the assured-delivery nowcast case, the distortion depends on the primitives only through the composite $ABC\beta\eta\phi c_\eta$ (which has fewer free parameters than the model, since $A = 1/(1 - C)$). For the common schedule held fixed across models here, the relevant slope must instead be recomputed from $J_\Omega(v, r)$; the candidate schedule is not reoptimized. Errors that preserve the reporting slope are harmless on this margin.

Lower bound on remaining incentives. The compatibility condition gives a quantitative measure of the design conflict. At a fixed interior type v, suppose the reporting slopes are bounded over the designer's model class, with extrema $w_-(v)$ and $w_+(v)$ (or their infimum and supremum). For any differentiable common correction,

$$\sup_{\Omega \in \mathcal{M}} |w_\phi(v; \Omega) + H'(v)| \;\geq\; \frac{w_+(v) - w_-(v)}{2}. \tag{A.135}$$

To see this, put $c = H'(v)$. The triangle inequality gives $|w_+ - w_-| \leq |w_+ + c| + |w_- + c|$; at least one absolute value is at least half the difference. The same argument applies by limits when an extremum is not attained. Pointwise, the bound is attained at $c = -(w_+ + w_-)/2$.

A feasible global audit must additionally integrate these slopes, satisfy its stake constraint, and preserve global reporting incentives. Thus (A.135) is a lower bound on the local incentive error; the minimax welfare-design problem is outside the scope of this paper. Increasing the common stake alone cannot cancel two different required slopes at the same announcement.

Why noisy delivery changes the strict-target conclusion. Use the pure-forecast experiment in Supplementary Appendix 2.5: $v_1 = 0$, $v_2 = s$, $p = ABC$, and $\bar{p} = \beta BC^2$. Under the strict target, $\Delta^a(\cdot)$ supports the desired delivered center ps with the schedule

$$g_0(r) = Cr + \Delta^a(-\bar{p}r), \qquad e_0(r) = pr. \tag{A.136}$$

The report-dependent loss is

$$J_0(s, r) = \frac{C\beta^2 p^2}{2} r^2 + \beta \mathcal{V}_C(-Cs + g_0(r)). \tag{A.137}$$

At truth, $\mathcal{V}_C' = A\Delta = -A\bar{p}s$. Since $C\beta p = A\bar{p}$ and $p = C - \bar{p}$, differentiation gives

$$\partial_r J_0(s, s) = \beta A \bar{p}^2 s\{\Delta^{a\prime}(-\bar{p}s) - 1\}. \tag{A.138}$$

For the finite audit of Section 4.1.2 on the interior responsive branch, the traction function has slope strictly between zero and one. Its inverse therefore has $\Delta^{a\prime} > 1$, so (A.138) is positive for $s > 0$ and nonzero for $s \neq 0$. This provides an explicit counterexample to extending zero announcement correction to arbitrary noisy delivery. There is no rule-response margin here: the remaining incentive comes from how the smaller announcement changes actual delivery and the expected audit burden. Under a coarse rule the additional term is $\beta A \bar{p} \phi s$, as shown in Supplementary Appendix 2.5.

An assured promise of a strict-target departure moves expected and actual future inflation equally and aligns the central bank's conditional choice with the planner's. A promise enforced only through a finite noisy audit need not do so after a reporting deviation, because delivery depends on the true shock as well as the promise. Model uncertainty about that delivery response can matter under either rule. The advantage proved in the main text is the absence of the *rule-induced* reporting distortion under a strict target, together with parameter-independent conditional alignment in the assured-delivery benchmark; the revelation, delivery, and welfare claims are stated for the model in which they are proved.

B Extensions

Three extensions of the pure-forecast analysis: delivery that is noisy rather than assured, rules that prescribe more than the ideal response, and guidance about current conditions.

B.1 Noisy Delivery and the Collapse Threshold

This subsection relaxes Assumption 1(a). Delivery is governed by the graduated parameterization of Section 4.1.2 with a finite stake: for a pre-shortfall deviation $u \equiv \Delta^a = (\phi - C)s + g$ the central bank chooses the shortfall $a^*(u)$ of (31), delivers $\Delta(u) = u + a^*(u)$, and bears the delivery value $\mathcal{V}_C(u)$, whose envelope derivative is $\mathcal{V}_C'(u) = A\Delta(u)$; the traction of guidance is $\omega = \Delta'(0) = \mathcal{A}_C''(0)/[A + \mathcal{A}_C''(0)] \in (0, 1)$. The innovation is set to zero so that the delivered center is a function of the announcement and the forecast alone. A public that reads the announcement as $\widehat{s}$ expects the delivered center $e_1 = \phi\widehat{s} + g + a^*((\phi - C)\widehat{s} + g)$, and the central bank of type s incurs $\frac{C}{2}[\beta e_1]^2 + \beta\mathcal{V}_C((\phi - C)s + g)$.

The threshold in closed form. Whether a separating reading can reach the origin is decided by the first-order condition near the origin, and the threshold below is a statement about the loss to leading quadratic order near the origin, where delivery is locally linear: $a^*(u) = -(1 - \omega)u + O(u^3)$ because a^* is odd and $\Delta'(0) = \omega$, and $\mathcal{V}_C(u) = \mathcal{V}_C(0) + \frac{1}{2}A\omega u^2 + O(u^4)$ by the envelope condition. Substituting, the loss of type s read as $\widehat{s}$ is, up to higher-order terms,

$$\frac{C}{2}\Big[\beta\big(\phi_\omega\widehat{s} + \omega g\big)\Big]^2 + \frac{\beta A\omega}{2}\big[(\phi - C)s + g\big]^2, \qquad \phi_\omega \equiv \omega\phi + (1 - \omega)C.$$

Only a fraction ω of the promise is delivered and moves expectations, while the rule's anticipated response to the inferred forecast is ϕ_ω: the public expects the rule's response $\phi\widehat{s}$ plus the part $(1 - \omega)(C - \phi)\widehat{s}$ of the rule's omission that the central bank will repair on its own by moving delivery toward its ex-post ideal. Repeating Steps 1–3 of the proof of Proposition 3 with these coefficients, a reading of slope γ is self-confirming near the origin if and only if

$$(A + \beta C\omega)\gamma^2 - \big[A(C - \phi) - 2\beta C\phi_\omega\big]\gamma + \frac{\beta C\phi_\omega^2}{\omega} = 0,$$

and, writing $c_1(\omega) \equiv 2\beta C + 2\sqrt{\beta C(A + \beta C\omega)/\omega}$ for the separation cost per unit of promise under traction ω, positive real roots exist if and only if $A(C - \phi) \geq \phi_\omega\, c_1(\omega)$. Since ϕ_ω is affine in ϕ, this is

$$\phi \leq \hat{\phi}(\omega) \equiv C\,\frac{A - (1 - \omega)\,c_1(\omega)}{A + \omega\,c_1(\omega)}, \qquad \hat{\phi}(1) = \hat{\phi}. \tag{B.1}$$

The threshold falls as the traction falls, because the delivery term disciplines the promise at rate $A\omega$ while the anticipated response ϕ_ω exceeds ϕ. Two consequences follow. First, with noisy delivery the strict target acquires a leak of its own, $\phi_\omega = (1-\omega)C$ at $\phi = 0$, because the public's forecast of the shortfall depends on what it infers; the truthful promise is then no longer dominant, and a positive separating slope at the origin exists under the strict target if and only if $\hat{\phi}(\omega) \geq 0$, that is, $A \geq (1-\omega)c_1(\omega)$, which at the baseline calibration holds for every traction above $\omega^* = 0.16$ (a stake above 3.1). Second, the below-threshold statements of the main text are assured-delivery statements; the table gives the threshold and the excess loss of the equilibrium of Proposition 3 at two rule slopes as the stake varies, with the tolerance at $\tau = \sigma_\pi$, bounded forecasts $S = 3\sigma_s$, and both the reading and the rule-constrained benchmark computed with the same delivery technology on the responsive range of the audit.

stake Ψ	traction ω	$\hat{\phi}(\omega)$	excess loss at $\phi = 0.3$	at $\phi = 0.45$
10	0.38	0.36	0.010	–
18	0.53	0.46	0.007	0.029
25	0.61	0.50	0.006	0.025
35	0.69	0.53	0.005	0.023
50	0.76	0.55	0.005	0.021
100	0.86	0.58	0.004	0.019
assured	1.00	0.61	0.004	0.017

Notes: At $\Psi = 10$ the rule $\phi = 0.45$ lies above $\hat{\phi}(\omega) = 0.36$, so no positive separating slope exists at the origin ("–").

At the baseline stake $\Psi = 18$ the threshold is $\hat{\phi}(\omega) = 0.46$; the excess losses are per unit of forecast variance and remain second order below the threshold. Numerically integrated readings from the boundary condition of Proposition 3(i) confirm (B.1): the forecast at which the trajectory reaches a zero announcement becomes positive at the closed-form threshold, up to the resolution of the detection.

The rule-constrained optimum under a finite stake. The comparison line of Figure 4, the truthful-guidance loss L^T, is a property of the delivered deviation, attainable within the implementable range. Let delivery be governed by the graduated parameterization with any stake $\Psi > 0$, set the innovation to zero and sanctions socially free, and let the rule-constrained planner choose, for each forecast s, the announcement whose delivered deviation minimizes the social loss on the responsive range of the audit.

Lemma 12 (Optimal hyperbole under a coarse rule)**.** *Suppose the responsive delivery choice is globally optimal. For every rule $0 \le \phi < C - \bar{c}$ and every forecast whose target deviation $\beta BC^2|s|$ lies below the ceiling $\bar{\Delta}_C$ of Lemma 3, the planner's delivered deviation is $\Delta = -\beta BC^2 s$, its expected inflation is $ABCs$, and its loss per unit of forecast variance is $L^T = \frac{1}{2}ABC(\beta C)^2$, the same as under assured delivery.*

Proof. With the innovation set to zero the delivered center is $\phi s + g + a^*(u)$ with $u = (\phi - C)s + g$, so the delivered deviation from the ex-post ideal is $\Delta(u)$ and the expectation of a public that reads the announcement correctly is $e_1 = Cs + \Delta(u)$. Sanctions being socially free, the social loss of type s is $\frac{C}{2}[\beta(Cs + \Delta)]^2 + \frac{\beta A}{2}\Delta^2$, a strictly convex function of the delivered deviation Δ alone, minimized at $\Delta = -\beta BC^2 s$ (its first-order condition is $\beta C(Cs+\Delta)+A\Delta = 0$), with $e_1 = Cs - \beta BC^2 s = ABCs$ and minimized value $\frac{1}{2}ABC(\beta C)^2 s^2$, the objects of Proposition 1. By Lemma 3, every target deviation below the ceiling is delivered by exactly one announcement on the responsive range, $g = Cs - \phi s + \Delta^a(-\beta BC^2 s)$, so the planner attains the minimum by announcing more than it delivers, by the shortfall the central bank will take. The loss therefore equals L^T per unit of forecast variance under these implementation conditions, and only the size of the announcement depends on the stake.

Under a finite stake the loss without guidance also changes: a central bank that announces nothing still repairs part of the rule's omission at delivery, delivering the deviation $\Delta((\phi - C)s)$ instead of $(\phi - C)s$, so the no-guidance loss per unit of forecast variance is $\frac{1}{2}\beta A\,\mathbb{E}[\Delta((\phi - C)s)^2]/\mathrm{Var}(s)$, and it crosses L^T at a smaller rule slope than $C - \bar{c}$ (at the baseline stake $\Psi = 18$, near $\phi = 0.54$). The value of guidance under a coarse rule therefore shrinks as the stake falls, while the truthful-guidance loss does not: this is the sense in which the collapse is a property of institutions strong enough to make words binding (Section 4.2.3).

B.2 Rules above the Ideal Response

This extension maintains the pure-forecast environment, assured delivery, and the obedience audit alone. A rule with $\phi > C$ prescribes more accommodation than the ex-post ideal. Guidance can correct its predictable excess through a negative promise. Propositions 1 and 2 cover this case: revelation is desirable when $\phi > C + \bar{c}$, and the reporting correction restores the rule-constrained Ramsey promise there.

The promise $g^R_\phi(s) = (ABC - \phi)s$ changes sign at $\phi = ABC$, not at C. Since $C - \bar{c} < ABC < C$, the interval (ABC, C) lies entirely inside the opacity region. The corrected reporting loss has curvature $\beta A(ABC - \phi)(C - \phi)$: it is positive for $\phi > C$, as well as for

$\phi < C - \bar{c}$, but negative on (ABC, C). No truthful-reporting implementation is claimed on that interval, where society prefers opacity.

Lemma 13 (Single-slope readings above the ideal response)**.** *Let $\phi > C$ and, as in Lemma 4, apply a nonzero single-slope reading to every announcement before imposing a support bound. Such a reading is self-confirming if and only if*

$$\phi \geq \hat{\phi}_+ \equiv \frac{AC}{[\sqrt{A+\beta C} - \sqrt{\beta C}]^2}, \qquad \hat{\phi}\hat{\phi}_+ = C^2. \tag{B.2}$$

When it exists, its slopes satisfy $\gamma^R < \gamma_1 \leq \gamma_2 < 0$, with equality between the two roots only at $\hat{\phi}_+$. Thus the two reading thresholds are geometric reflections about C. In the intervals $(\hat{\phi}, C - \bar{c})$ and $(C + \bar{c}, \hat{\phi}_+)$, revelation is socially desirable but no single-slope reading is self-confirming.

Proof. The first-order condition for a linear reading gives the same quadratic as below the ideal rule:

$$Q(\gamma) = (A + \beta C)\gamma^2 - [A(C - \phi) - 2\beta C\phi]\gamma + \beta C\phi^2 = 0. \tag{B.3}$$

Its discriminant is $A[A(C-\phi)^2 - 4\beta C^2\phi]$. Its zeros as a function of ϕ are $\hat{\phi}$ and $\hat{\phi}_+$, whose product is C^2. For $\phi > C$, the sum of the roots is negative and their product is positive. Moreover, $Q(\gamma^R) > 0$ and $Q'(\gamma^R) = A(C - \phi) < 0$, so both roots exceed γ^R. The objective facing a type under any nonzero linear reading is strictly convex in its promise. A root therefore gives a global best response, not merely a stationary one. Finally, $C + \bar{c} < \hat{\phi}_+$: writing $t = \sqrt{B\beta C} \in (0, 1)$ gives $\bar{c} = Ct$ and $\hat{\phi}_+ = C(1+t)/(1-t) > C(1+t)$.

A bounded forecast support also permits an explicit equilibrium above C. The following result establishes existence and specifies the equilibrium used in the numerical exercise. It does not extend the uniqueness claim of Proposition 3 to all equilibria above C.

Lemma 14 (A bounded-support equilibrium above the ideal response)**.** *Let the prior have a symmetric, positive density on $[-S, S]$, and let $\phi > C$. There exists a symmetric perfect Bayesian equilibrium satisfying D1 with $G(s)$ nonincreasing. Define the promised restraint $H(s) \equiv -G(s)$ and $h_R \equiv \phi - ABC > 0$.*

(i) If $C < \phi \leq C + \bar{c}$, every type announces zero.

(ii) If $C + \bar{c} < \phi < \hat{\phi}_+$, a central interval $[-s_b, s_b]$ announces zero and the two tails separate. On the positive tail, H is the increasing trajectory pinned by $H(S) = h_R S$, with inverse equation

$$\frac{ds}{dh} = \frac{(A + \beta C)(h_R s - h)}{\beta C\phi(\phi s - h)}. \tag{B.4}$$

The boundary is the unique point on this trajectory at which

$$\frac{H(s_b)}{s_b} = u_b \equiv h_R - \sqrt{h_R^2 - \frac{\beta C}{A+\beta C}\phi^2}. \tag{B.5}$$

(iii) If $\phi \geq \hat{\phi}_+$, the same trajectory separates every type, with $H(0) = 0$ and $H(s)/s \longrightarrow -\gamma_1$ as $s \downarrow 0$.

The negative tail is the odd extension. The construction is homogeneous in S: $H_S(s) = SH_1(s/S)$ and s_b/S is independent of the support bound and of the shape of the symmetric prior.

Proof. The loss and the separating trajectory. Put $a = \phi - C > 0$. If an announcement of restraint $h = -g$ is read as forecast r, the central bank's loss, up to an announcement-independent term, is

$$\widetilde{J}(s, h; r) = \frac{C\beta^2}{2}(\phi r - h)^2 + \frac{\beta A}{2}(h - as)^2. \tag{B.6}$$

The zero announcement with posterior mean zero gives $\Pi(s) = \beta A a^2 s^2/2$. The first-order condition along an informative schedule is

$$\beta C\phi(\phi s - H(s)) + (A + \beta C)[H(s) - h_R s]H'(s) = 0, \tag{B.7}$$

which yields (B.4). The inverse equation is regular at $h = h_R S$, since $\phi S - h_R S = ABCS > 0$.

Write $u = H(s)/s$ along the trajectory. Then

$$\frac{du}{d\log s} = \frac{Q_h(u)}{(A+\beta C)(h_R - u)}, \qquad Q_h(u) = (A+\beta C)u^2 - [Aa + 2\beta C\phi]u + \beta C\phi^2. \tag{B.8}$$

Starting at $u = h_R$, the ratio falls as s falls. If $\phi \geq \hat{\phi}_+$, it approaches the larger positive root $u_+ = -\gamma_1 < h_R$ only as $s \downarrow 0$. If $C + \bar{c} < \phi < \hat{\phi}_+$, the polynomial is positive everywhere and the trajectory reaches $u = 0$ at a positive forecast. Throughout the relevant trajectory, $0 \leq u \leq h_R < \phi$ and H is increasing.

The pooling boundary. The gain from separation is

$$\begin{aligned} \Pi(s) - \widetilde{J}(s, H(s); s) &= s^2 F(u), \\ F(u) &= \frac{\beta}{2}\left[-(A+\beta C)u^2 + 2(A+\beta C)h_R u - \beta C\phi^2\right]. \end{aligned} \tag{B.9}$$

The function F is strictly increasing on $[0, h_R)$, with $F(0) < 0$. At the top, $F(h_R) > 0$ exactly when $\phi > C + \bar{c}$, by Proposition 1. In the band, it therefore has the unique zero (B.5), reached before the trajectory reaches zero. Outside the band, $F(u_+) = \beta A a u_+/2 > 0$, using

$Q_h(u_+) = 0$; hence separation is preferred to zero at every positive forecast. Homogeneity follows directly from (B.4) and the ratio boundary condition.

Beliefs. The posterior mean is zero at $h = 0$ and equals the inverse trajectory on each separating tail. If there is a central pool, assign mean s_b to $0 < h < H(s_b)$, and assign mean S to $h > h_R S$. Negative announcements receive the opposite means. In the fully pooling case assign mean S to every $h > 0$ and $-S$ to every $h < 0$. These beliefs are Bayes-consistent on path; symmetry of the prior gives mean zero in the pool.

Global incentive compatibility. Differentiating the loss of type s along an informative announcement indexed by r and using (B.7) gives

$$\frac{d}{dr}\widetilde{J}(s, H(r); r) = \beta A a (r - s) H'(r). \tag{B.10}$$

Thus a separating type prefers its own announcement to every other announcement on its tail. The gain calculation shows that it also prefers its own announcement to zero.

For a pooled type $0 \le s \le s_b$, the best positive tail announcement is the boundary one, by (B.10). For any fixed $h > 0$ read as s_b,

$$\frac{\partial}{\partial s}\left[\widetilde{J}(s, h; s_b) - \Pi(s)\right] = -\beta A a h < 0. \tag{B.11}$$

At $s = s_b$, the boundary announcement and zero give equal losses. With the belief fixed at s_b, the optimal restraint of that type is $h_R s_b \ge H(s_b)$, so every smaller positive announcement costs at least as much as the boundary announcement. Equation (B.11) then rules out these deviations for every pooled type. For a separating type $s \ge s_b$, the optimal restraint at belief s_b is $[Aas + \beta C \phi s_b]/(A + \beta C) \ge h_R s_b$. Its loss from an announcement in the gap is therefore at least its loss from the boundary announcement, which truthful separation already dominates.

At belief S, every type $s \le S$ has optimal restraint at most $h_R S$. Announcements beyond the top are consequently dominated by the top announcement. For $s \ge 0$, any $h < 0$ has delivery loss alone at least $\Pi(s)$; such a deviation cannot improve on either pooling or separation. Symmetry gives the negative-type conditions.

In the fully pooling case, apply (B.11) with $s_b = S$. For $h > 0$, the type most tempted to deviate is S. Its minimum loss at belief S is $L^T S^2$, whereas its pooling loss is $\Pi(S) \le L^T S^2$ when $\phi \le C + \bar{c}$. No type gains by deviating; negative deviations follow by symmetry.

D1. Let $J^{\text{eq}}(s)$ denote the equilibrium loss. For a fixed off-path h, the posterior means

that make the deviation weakly profitable satisfy

$$(\phi r - h)^2 \leq c(s;h), \qquad c(s;h) \equiv \frac{2}{C\beta^2}\left[J^{\text{eq}}(s) - \frac{\beta A}{2}(h - as)^2\right]. \tag{B.12}$$

These sets are nested intervals, intersected with the feasible means $[-S, S]$. The equilibrium loss is continuous at the pool boundary. The envelope condition, including the pooled region where $H = 0$, gives

$$\frac{\partial c(s;h)}{\partial s} = \frac{2Aa}{C\beta}[h - H(s)] \tag{B.13}$$

where the derivative exists, with the same one-sided signs at a boundary. For $0 < h < H(s_b)$, the gain set is largest at s_b; for $h > h_R S$, it is largest at S. Under full pooling it is largest at S for every $h > 0$. The specified off-path beliefs put weight only on these maximizing types and therefore satisfy D1. Negative announcements follow by symmetry. This completes the construction.

At the baseline, the lower band is approximately $(0.612, 0.743)$ and the upper band is $(1.150, 1.465)$. The reading thresholds satisfy (B.2); the bands themselves are not exact geometric reflections, since their revelation boundaries are the arithmetic reflections $C \pm \bar{c}$. In either band, the constructed equilibrium leaves ordinary forecasts uninformative even though society would benefit from disclosure. Above C, the missing policy adjustment is restraint rather than additional accommodation.

B.3 Guidance about Current Conditions

This appendix studies guidance about current conditions. The pure forecast is absent ($s = 0$ almost surely), and the private nowcast v_1 predicts tomorrow's shock through $\mathbb{E}[v_2 \mid v_1] = \eta v_1$. Revealing current conditions therefore also changes expectations of the standing rule's future response. Each result below points to its proof in Appendix A.

B.3.1 Overview: Revelation and Collapse for the Nowcast

The main-text propositions include the corresponding nowcast results. This section brings them together in Figure B.1 and states them as Lemmas 15–18.

For the planner, revealing the nowcast has a benefit the pure forecast lacks: it permits a promise of future restraint that improves stabilization today. Under a coarse rule, guidance also supplies the response the rule omits to the predictable part of tomorrow's shock. Against both stands the leak of the bundled forecast into current expectations. Under

the ideal rule only the first benefit exists, and it is small: revelation is optimal only for a news content below $\eta^\star = 0.026$ at the baseline calibration. As the rule coarsens, the value of repairing its omission grows with the square of the news content and the revelation region widens; under the strict target revelation is optimal for every η (Proposition 1). The revelation edge in Figure B.1 traces this region.

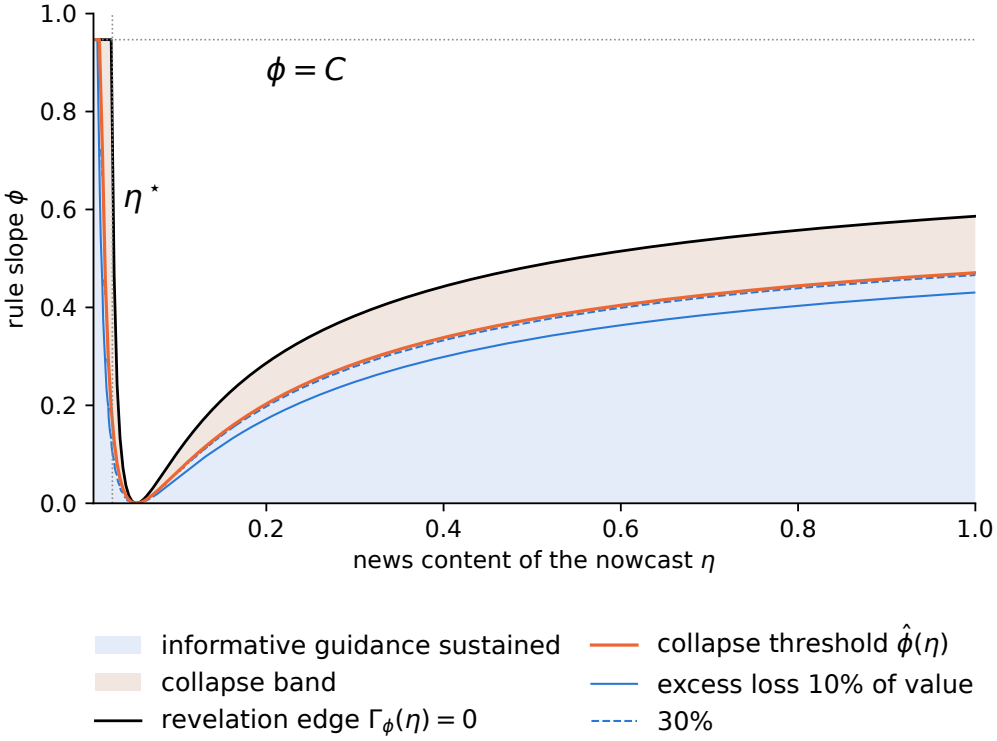

Figure B.1: Guidance about current conditions: revelation and collapse

Notes: Nowcast news content η and rule slope ϕ; assured delivery, obedience audit only, baseline calibration. The upper dotted line is the ideal rule, $\phi = C$. Black: revelation edge, $\Gamma_\phi(\eta) = 0$; revelation is preferred below it. Orange: single-slope threshold $\hat{\phi}(\eta)$; self-confirming informative readings exist below it (blue). In the shaded band, revelation is desirable but no such reading exists. Blue contours: excess loss of the near-Ramsey reading, at 10% and 30% of the information value. The figure does not characterize nowcast pooling.

As in the pure-forecast case, the existence of a self-confirming reading of a single slope depends on the rule. Here the threshold also depends on the news content: such readings require $\phi \leq \hat{\phi}(\eta)$, the V-shaped threshold of Lemma 4. Its zero at $\eta = 1/A$ is not a costly failure of communication under the strict target: there the planner's promise is zero for every nowcast, and revelation and opacity yield the same loss (Proposition 4); under positive rule slopes the planner prefers opacity at this news content. The collapse that matters occurs where society strictly benefits from revelation but the announcement game cannot sustain it. The threshold equals C at $\eta = \eta^\star/2$, so under the ideal rule the central

bank sustains informative guidance on its own over only half the range of news content in which society wants it. At the baseline calibration with $\eta = 0.2$ the threshold is 0.203, roughly a fifth of the optimal response C: the collapse can arise under rules far from the ideal one. With a bounded nowcast the separating schedule below the threshold is pinned as it is for the forecast (Proposition 3).

B.3.2 Revelation under the Ideal Rule

Consider the unrestricted Ramsey planner of Section 3 and a communication strategy $q = q(v_1)$ that conveys the nowcast and whatever it bundles about v_2. By Lemma 1 the planner attaches to message q the promise $g^R(q) = -BCM(q)$, where $M(q) = \widehat{v}_1(q) + \beta C \widehat{v}_2(q)$ is the inflationary pressure that the released information creates through expectations, and the value of the strategy is (18); Appendix A.6 records the intermediate objective. For messages generated from v_1, (10) gives $M(q) = (1 + \beta C\eta)\widehat{v}_1(q)$, and the law of total variance reduces (18) to

$$\mathcal{L}(q) = Z + C \operatorname{Var}(v_1)/2 + \frac{C}{2}\Gamma(\eta) \operatorname{Var}\big(\widehat{v}_1(q)\big), \qquad \Gamma(\eta) \equiv AB(1 + \beta C\eta)^2 - 1, \tag{B.14}$$

where the first two terms are independent of the communication strategy. Every admissible communication strategy produces a posterior-mean variance between zero and $\operatorname{Var}(v_1)$, and the loss is linear in it, so the planner either reveals v_1 or reveals nothing (policy-relevant opacity: no message moves the posterior mean of v_1); the sign of $\Gamma(\eta)$ decides which.

Lemma 15 (Ramsey revelation threshold)**.** *Suppose* (5) *holds,* $\operatorname{Var}(v_1) > 0$*, and the planner can choose any communication strategy* $q = q(v_1)$*.*

(i) If $\Gamma(\eta) < 0$*, full revelation is optimal, and every optimal communication strategy satisfies* $\widehat{v}_1(q) = v_1$ *almost surely.*

(ii) If $\Gamma(\eta) > 0$*, policy-relevant opacity is optimal, and every optimal communication strategy satisfies* $\widehat{v}_1(q) = 0$ *almost surely.*

(iii) If $\Gamma(\eta) = 0$*, every admissible communication strategy yields the same loss.*

For $\eta \geq 0$*, define*

$$\eta^\star \equiv \frac{1/\sqrt{AB} - 1}{\beta C}. \tag{B.15}$$

Then full revelation is optimal for $\eta < \eta^\star$ *and policy-relevant opacity for* $\eta > \eta^\star$*. Equivalently,*

revelation is strictly preferred to pooling whenever

$$\mathrm{Cov}(v_1, v_2) < \eta^\star \,\mathrm{Var}(v_1). \tag{B.16}$$

Proof. See Appendix A.6.

What the formula for $\eta^\star$ shows is the magnitude of the threshold. At the baseline calibration, $\kappa = 0.15$, $\lambda = 0.4$, and $\beta = 0.96$, the threshold is $\eta^\star = 0.026$; more generally, $\eta^\star \leq 1/(2A)$ (Appendix A.7 establishes this bound and its invariance to measurement noise). The region is small because revealing the nowcast buys the classic gain from commitment, a promise of restraint tomorrow that lowers inflation today, and that gain is small: carrying out the restraint contracts output tomorrow. Against it stands the bundled forecast, which the public expects the ideal rule to accommodate; revealing it raises expected inflation today and improves nothing tomorrow, because the rule was going to respond anyway. Writing $\mathcal{L}^{\text{reveal}}$ for the Ramsey loss when the message reveals v_1 and $\mathcal{L}^{\text{opacity}}$ when it reveals nothing, (B.14) gives

$$\mathcal{L}^{\text{opacity}} - \mathcal{L}^{\text{reveal}} = -\frac{C}{2}\Gamma(\eta)\,\mathrm{Var}(v_1) = \frac{C}{2}\left[1 - AB(1+\beta C\eta)^2\right]\mathrm{Var}(v_1), \tag{B.17}$$

which is positive exactly when $\eta < \eta^\star$. When revelation is optimal, the promise attached to the revealed state is the scalar Ramsey deviation $g^R(v_1) = -BCc_\eta v_1$, $c_\eta \equiv 1 + \beta C\eta$: the ideal rule handles the response to the shock in period 2, and $g^R(v_1)$ carries the commitment made in the stabilization period into the future.

B.3.3 Revelation under Coarse Rules

Under the rule-plus-promise form (23) with slope ϕ, the period-2 deviation from the ex-post ideal is $(\phi - C)v_2 + g(q)$ and the public's expectation is $e_1(q) = \phi\widehat{v}_2(q) + g(q)$. Relative to the ideal rule, the rule mishandles the predictable part of the shock by $(C - \phi)\eta\widehat{v}_1(q)$ unless the promise supplies it, and the surprise $v_2 - \eta\widehat{v}_1(q)$ passes through at slope ϕ rather than C.

Lemma 16 (Information revelation under coarse rules)**.** *For communication strategies $q = q(v_1)$ satisfying (10), the rule class (22) and restriction (23) imply that the rule-constrained Ramsey loss is affine in $\mathrm{Var}(\widehat{v}_1(q))$ with coefficient $\frac{C}{2}\Gamma_\phi(\eta)$, where*

$$\Gamma_\phi(\eta) = \Gamma(\eta) - \frac{\beta A}{C}(C - \phi)^2\eta^2. \tag{B.18}$$

Full revelation is optimal when $\Gamma_\phi(\eta) < 0$ and policy-relevant opacity when $\Gamma_\phi(\eta) > 0$; at equality,

all admissible communication strategies are payoff-equivalent. Given the message, the optimal promise is

$$g^R_\phi(q) = -B\big(C + \big[A(\phi - C) + \beta C\phi\big]\eta\big)\,\widehat{v}_1(q) = g^R(q) + (C - \phi)\eta\,\widehat{v}_1(q), \tag{B.19}$$

where g^R is the Ramsey promise of Lemma 1(ii).

Proof. See Appendix A.13.

The promise (B.19) is the unrestricted Ramsey promise plus a repair, $(C - \phi)\eta\widehat{v}_1(q)$, the response to the predictable shock that the rule fails to carry. The revelation coefficient (B.18) is the ideal-rule coefficient minus a nonnegative term, quadratic in η: the larger the bundled forecast the message carries, the larger the predictable error a coarse rule would commit, and the more valuable it is to reveal the nowcast so that the promise can repair it. The revelation region is therefore wider than under the ideal rule, $\Gamma_\phi \le \Gamma$ with equality only at $\phi = C$ (for $\eta > 0$); Figure B.1 maps it. At the optimal promise the composite $v_1 + \beta e_1$ and the conditional-mean deviation from the ex-post ideal are the same for every ϕ: the promise absorbs whatever the rule leaves out of the predictable part of the shock. Under the strict target the coefficient is $\Gamma_0(\eta) = -B\beta C(A\eta - 1)^2 \le 0$, so full revelation of v_1 is optimal for every η, and the optimal nowcast promise $g^R_0(v_1) = BC(A\eta - 1)v_1$ vanishes at $\eta = 1/A$ (Corollary 8 in Appendix A.13).

B.3.4 Reporting Distortion for the Nowcast

Fix a rule slope ϕ in the revelation region and the candidate policy that assigns the rule-constrained Ramsey promise to each nowcast. Write r for the nowcast the public infers from the observed promise. From Lemma 16, the promise and the expectation the public attaches to the report r are

$$g^R_\phi(r) = -B\big(C + [A(\phi - C) + \beta C\phi]\eta\big)\,r, \qquad e^R_\phi(r) = \phi\eta r + g^R_\phi(r). \tag{B.20}$$

The expectation has two parts: the rule's anticipated response $\phi\eta r$ to the predictable shock the report implies, and the promise. A central bank of true type v_1 that reports r chooses the current output gap optimally afterwards and, under assured delivery, its loss is

$$J_\phi(v_1, r) = \frac{C}{2}\Big[v_1 + \beta e^R_\phi(r)\Big]^2 + \frac{\beta A}{2}\Big[(\phi - C)\eta v_1 + g^R_\phi(r)\Big]^2 \tag{B.21}$$

up to terms that are independent of the report. The second bracket contains the unrepaired part of the rule's response to the predictable shock, which depends on the true v_1 rather

than on the report.

Lemma 17 (Reporting distortion for the nowcast)**.** *At the truthful report, the marginal value of reporting a smaller nowcast, the reduction in the central bank's loss per unit reduction in the report, is*

$$\frac{\partial J_\phi}{\partial r}(v_1, v_1) = ABC\beta\eta\phi\, c_\eta\, v_1, \tag{B.22}$$

positive for a positive nowcast when $\eta\phi > 0$, and zero when $\eta\phi = 0$. Let

$$\mathcal{A}_R^\phi(r) \equiv -\frac{1}{2} ABC\beta\eta\phi\, c_\eta r^2. \tag{B.23}$$

If the expected audit burden associated with report r varies according to $\mathcal{A}_R^\phi$, then $r = v_1$ is a stationary point of $J_\phi(v_1, r) + \mathcal{A}_R^\phi(r)$ for every type. Throughout the strict revelation region $\Gamma_\phi(\eta) < 0$, it is the unique global minimizer. More generally, uniqueness holds whenever $\mathcal{K}_\phi(\eta) > 0$, where the polynomial $\mathcal{K}_\phi$ is given in Appendix A.14. In particular:

(i) under the strict target, $\phi = 0$, the distortion is zero for every η, $\mathcal{K}_0(\eta) = C^2(A\eta - 1)^2 \geq 0$, and the central bank chooses the truthful promise with no correction; at $\eta = 1/A$, that promise is constant and does not reveal the nowcast;

(ii) under the ideal rule, $\phi = C$, the distortion is $AC^2\beta\eta Bc_\eta v_1$ and the condition reduces to $\eta < 1/A$, which holds throughout the revelation region $\eta < \eta^\star$.

Proof. See Appendix A.14.

The distortion is the product of the rule's loading on the future content of the report, $\phi\eta$, and the value to the central bank of a unit reduction in expected inflation, $\beta C \cdot ABc_\eta v_1$, where $ABc_\eta v_1$ is the stabilization composite $v_1 + \beta e_1$ at the truthful report. Nothing operating through the promise appears, for the reason given in the main text for the pure forecast: at the truthful report the planner's first-order condition makes the marginal cost and benefit of a different promise cancel. The distortion is proportional to ϕ: it is largest under the ideal rule, in the thin region where that rule reveals at all, it shrinks as the rule coarsens, and it vanishes under the strict target. What a coarse rule does is make the distortion matter: it widens the region in which the planner wants the nowcast revealed, so that truthful reporting must be induced where the ideal rule would have had the central bank announce nothing.

B.3.5 Readings of Nowcast Guidance and the Collapse Threshold

Under Assumption 1(a), (b), and (d), a reading of slope γ, $G(v_1) = \gamma v_1$, is decoded as $\widehat{v}_1 = g/\gamma$ and moves the expectation to $e_1(g) = \phi\eta\widehat{v}_1 + g$; the comparison point is the

rule-constrained Ramsey slope $\gamma_\eta^R = -B(C + [A(\phi - C) + \beta C\phi]\eta)$ of Lemma 16.

Lemma 18 (Self-confirming readings of the nowcast)**.** *Let $\Gamma_\phi(\eta) < 0$. A reading with nonzero slope γ, $G(v_1) = \gamma v_1$, is self-confirming if and only if*

$$(A + \beta C)\gamma^2 + \left[C + 2\beta C\phi\eta - A(C - \phi)\eta\right]\gamma + C\phi\eta(1 + \beta\phi\eta) = 0, \tag{B.24}$$

and, for $\eta > 0$, such readings exist if and only if $\phi \leq \hat{\phi}(\eta)$,

$$\hat{\phi}(\eta) \equiv \frac{C}{A}\left[\sqrt{A + \beta C} - \sqrt{\beta C + 1/\eta}\right]^2. \tag{B.25}$$

For $\phi > 0$ and $\eta > 0$, the two roots, when they exist, have the sign of γ_η^R and lie between 0 and γ_η^R; they coincide at the threshold. At $\eta = 0$, the only nonzero root is $-BC$, independently of the rule slope. The threshold falls from $\hat{\phi}(\eta) = C$ at $\eta = \eta^\star/2$ to $\hat{\phi}(1/A) = 0$ and then rises with η toward the pure-forecast threshold $\hat{\phi}$. Under the strict target the best response is $g = BC(A\eta - 1)v_1$ whatever γ, and the truthful reading is the only self-confirming reading provided $\eta \neq 1/A$. The social loss under a self-confirming reading of slope γ exceeds the rule-constrained Ramsey loss by $\frac{1}{2}(\beta A + \beta^2 C)(\gamma - \gamma_\eta^R)^2$ per unit of variance of the nowcast.

Proof. See Steps 5–10 of the proof of Lemma 4 in Appendix A.15.

The nowcast threshold depends on its news content. It is zero at $\eta = 1/A$. Under the strict target, the planner's promise is then zero for every type, and opacity entails no loss. Under a positive rule slope, the planner instead prefers opacity at this news content. At $\eta = \eta^\star/2$, the threshold equals C: under the ideal rule, informative guidance is self-confirming over only half the revelation region. At the baseline calibration with $\eta = 0.2$, $\hat{\phi}(\eta) = 0.203$, roughly a fifth of the ideal response. For news content between 0.1 and 1, the gap between the revelation edge and the single-slope threshold is a fifth to two fifths of the revelation range in ϕ (Figure B.1).

With a bounded nowcast, $v_1 \in [-\bar{v}, \bar{v}]$, the boundary condition and differential equation of Proposition 3 carry over under the conditions given in its proof. Below $\hat{\phi}(\eta)$, global incentive compatibility is checked numerically. The pure-forecast pooling theorem does not transfer: it uses the property that all types rank beliefs alike, which fails when the current shock enters the expectations term (Online Appendix A.15).

C Numerical Example

The numerical illustrations use $\beta = 0.96$, $\kappa = 0.15$, and $\lambda = 0.4$, unless a figure caption states otherwise; the forecast has $\sigma_s = 1$ and support $S = 3\sigma_s$, the innovation has $\sigma_{2|1} = 0.6$,

and the obedience audit has $\sigma_\pi = \tau = 0.55$ and stake $\Psi = 18$. Table C.1 reports the implied coefficients. The stationary illustration additionally sets $\rho_v = 0.70$.

Table C.1: Coefficients at the baseline calibration

Object	Definition	Value
Stabilization curvature	$A = 1 + \lambda/\kappa^2$	18.778
Discretionary accommodation per unit of shock	C	0.94675
Commitment accommodation per unit of forecast, truthful guidance	ABC	0.90304
Unrestricted Ramsey promise per unit of forecast	$-\beta BC^2$	−0.04371
Revelation threshold for the nowcast, ideal rule	$\eta^\star$	0.02631
Revelation boundary for the pure forecast, $C - \phi > \bar{c}$	$\bar{c}$	0.20342
Edge of the revelation region	$C - \bar{c}$	0.74332
Collapse threshold, pure forecast	$\hat{\phi}$	0.61186
Decay rate of the inherited promise	ρ_μ	0.80293
Infinite-horizon Ramsey rule slope	ϕ_R	1.74388
Continuation-adjusted temptation	Υ_R	5.07440

Introduction's example. With no current shock and a unit forecast of the future shock held privately, the discretionary center is C and the commitment center under truthful guidance is ABC; under the rule $\phi = 0.70$ the planner's promise is $\gamma^R = ABC - \phi = 0.203$ per unit of forecast. An announcement associated with a forecast of 0.8 induces the expectation $0.8ABC$ when the public takes the announcement at face value; under assured delivery the delivered accommodation is $0.8ABC$ at the strict target and $0.70 + 0.8\gamma^R$ under the coarse rule, approximately 0.72 and 0.86. The equilibrium numbers are those of Proposition 3 with $S = 3$ ($\sigma_s = 1$): the largest forecast announces $\gamma^R S$, the announcement schedule is the trajectory of (A.103) from the largest forecast's announcement $\gamma^R S$, and the boundary s_b of the uninformative announcement is the root of the gain function Φ constructed in Step 8 of the proof of Proposition 3 (with φ as in Lemma 8). The recovered share of the value of information and the mass of the uninformative announcement are computed under a normal forecast distribution truncated at $\pm S$. The replication package produces every number in the example.

Delivery and accountability. For the numerical checks of the graduated parameterization, $\sigma_\pi = \tau = 0.55$ and $\Psi = 18$ satisfy the global delivery condition (32). Integration over Gaussian implementation noise verifies (A.59). Root finding and direct minimization check the inverse traction function, including the target-associated shortfall $a^*(\Delta^a(\Delta^\star))$. Differentiating the delivery-valued loss verifies the ideal-rule audit rebate (A.123), the additional-capacity formula (10) of the supplementary appendix, and the sign of the noisy

strict-target counterexample (A.138). These checks complement the global reporting and delivery conditions in the corresponding results.

Independent checks. The replication package compares the revelation and promise formulas with direct optimization, checks the reporting curvatures, integrates the Gaussian audit and total-variation expressions, verifies the coefficients of Table C.1, and verifies the dynamic value matrix by a discounted Lyapunov solution. Algebraic and integration checks use tolerances of 10^{-8} or tighter; numerical minimization and finite-difference checks use tolerances no greater than 2×10^{-6}.

D Infinite Horizon

This appendix extends the results to an infinite-horizon model: the Ramsey decomposition into a stationary rule plus an inherited promise, the forecast interaction, and recursive delivery of the inherited promise. The communication strategy is taken as given. The operating implementation through interest rates is in the supplementary appendix (Supplementary Appendix 1.1).

D.1 Stationary Rule plus Inherited Promise

The conditional-mean Phillips curve in the fully dynamic setup is

$$\theta_t = v_t + \kappa x_t + \beta \mathbb{E}_t \theta_{t+1}, \tag{D.1}$$

where v_t is the cost-push disturbance, public at date t, and the Ramsey planner solves

$$\begin{aligned} \min_{\{\theta_t, x_t\}_{t \geq 0}} \quad & \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2}(\theta_t^2 + \lambda x_t^2) \\ \text{subject to} \quad & \theta_t = v_t + \kappa x_t + \beta \mathbb{E}_t \theta_{t+1} \qquad \text{for every } t. \end{aligned} \tag{D.2}$$

The nominal rate remains unrestricted and implements the chosen allocation through the IS equation, so the Phillips curve is the only independent forward-looking constraint.

The compression of Lemma 1 to a single number survives here. Let $e_t \equiv \mathbb{E}_t \theta_{t+1}$ be the expectation selected at date t and e_{t-1} the expectation inherited from the previous date. Define the composite

$$\Xi_t \equiv v_t + \beta e_t, \qquad \delta_t \equiv \theta_t - C \Xi_t. \tag{D.3}$$

Eliminating x_t from (D.1) and completing the square gives

$$\frac{1}{2}\theta_t^2 + \frac{\lambda}{2\kappa^2}(\theta_t - \Xi_t)^2 = \frac{A}{2}\delta_t^2 + \frac{C}{2}\Xi_t^2, \tag{D.4}$$

the dynamic counterpart of (11): $C\Xi_t$ is the ex-post ideal at date t, the no-promise center, and δ_t the departure from it. The inherited expectation fixes the conditional mean of δ_t, and, as in the episode, variation around that mean changes no inherited expectation and only raises loss.

The lemma below characterizes Ramsey policy. The rule slope ϕ_R is the response of the Ramsey center to the current shock when no promise is outstanding. The inherited promise G_{t-1}^R is the departure from that response that earlier commitments require today; it is the dynamic counterpart of the promise g of the episode. The decay rate ρ_μ governs how fast an inherited promise is worked off. The lemma states that the Ramsey center at every date is the rule plus the inherited promise, that the promise follows a first-order law driven by the current shock, and that the policy satisfies the familiar target criterion.

Lemma 19 (Infinite-horizon Ramsey rule and inherited promise)**.** *Suppose $v_{t+1} = \rho_v v_t + \varepsilon_{t+1}$ with $|\rho_v| < 1$ and innovations of finite, state-independent variance, and that the Ramsey allocation is bounded and satisfies the usual transversality condition. Let $\rho_\mu \in (0,1)$ be the smaller root of*

$$\beta\rho^2 - \left(1 + \beta + \frac{\kappa^2}{\lambda}\right)\rho + 1 = 0, \tag{D.5}$$

and define the rule slope and the inherited promise by

$$\phi_R \equiv \frac{\rho_\mu}{1 - \beta\rho_v\rho_\mu}, \qquad G_t^R \equiv \rho_\mu G_{t-1}^R - \phi_R(1-\rho_\mu)v_t, \qquad G_{-1}^R = 0. \tag{D.6}$$

Then Ramsey policy is exactly

$$\theta_t^R = \bar{\theta}^R(v_t) + G_{t-1}^R, \qquad \bar{\theta}^R(v_t) \equiv \phi_R v_t, \tag{D.7}$$

and the allocation satisfies the target criterion

$$\theta_t^R + \frac{\lambda}{\kappa}(x_t^R - x_{t-1}^R) = 0. \tag{D.8}$$

Proof. See Appendix D.3.

Equation (D.7) is the rule-plus-promise form (23) of the episode in a recursive setting. The rule $\bar{\theta}^R(v_t)$ responds to the current observable shock with the slope ϕ_R, which plays the

role of C; under the baseline calibration $\rho_\mu = 0.803$ and $\phi_R = 1.744$. The inherited promise carries every consequence of earlier commitments that remains relevant today, and (D.6) says how it evolves: a positive cost-push shock raises the current rule center and creates a negative promise for later policy, of size $\phi_R(1-\rho_\mu)$ per unit of shock, and the promise decays at the rate ρ_μ, which depends on the loss function and the Phillips curve alone.

The recursive model has an inflation center and an audit in every period. The gap component of the audit in Lemma 5 also applies in every period because the output-gap target differs from society's throughout. The two-period episode starts with no inherited promise, $G^R_{-1} = 0$, whereas the recursive model carries past commitments forward. The target criterion (D.8) expresses the tradeoff between inflation and the change in the output gap. Economies with the same current shock can therefore have different Ramsey allocations because their inherited promises differ.

A rule restricted to the current shock cannot encode this history. The inherited promise summarizes the policy obligations created by past shocks and remains necessary even without private information. The shock decays at rate ρ_v and the promise at rate ρ_μ, so the obligation may persist after the original disturbance has subsided. The remaining analysis characterizes the incentives required to deliver it.

Temptation to undo the inherited promise. To implement the recursive promise, the mandate must measure the future central bank's temptation in the same units as the public announcement. Conditional on the Ramsey continuation, a central bank that ignores its inherited obligation would choose the no-promise center $C\Xi_t$, and the proof of Lemma 19 shows that Ramsey policy exceeds it by $G^R_{t-1}/[A(1-\rho_\mu)]$. Since current loss has curvature A around the no-promise center, the marginal temptation to reduce the inherited promise is

$$A\left(\theta^R_t - C\Xi_t\right) = \Upsilon_R G^R_{t-1}, \qquad \Upsilon_R \equiv \frac{1}{1-\rho_\mu}. \tag{D.9}$$

The coefficient Υ_R is the continuation-adjusted temptation per unit of the public promise, the dynamic counterpart of the curvature A that priced the temptation in the episode. It is smaller than A, 5.074 against 18.778 under the baseline calibration, because G^R_{t-1} is a normalized recursive obligation rather than a one-period displacement from discretion: part of an inherited promise is worked off by future promises that the central bank itself wants to keep.

D.2 Stationary Rules and Private Forecasts

The interaction between the standing rule and private forecasts can also be studied in the stationary model. The relevant object is the cross term between the public shock and the operational promise in the Ramsey continuation value. In coordinates (v_t, G^R_{t-1}), discounted Ramsey loss has no cross term:

$$\mathcal{V}^R(v, G) = \frac{1}{2}P_v v^2 + \frac{1}{2}\Upsilon_R G^2 + \text{constant}. \tag{D.10}$$

The absence of a vG term is the dynamic analogue of the completion of the square in (11): under the ideal rule the public shock and the promise are separate sufficient statistics, and the curvature of the value function in the promise is the temptation Υ_R of (D.9). Both facts are proved in Appendix D.4 and verified numerically in Appendix C by solving the discounted Lyapunov equation of the stationary allocation.

Now suppose the standing rule has slope ϕ and define the operational promise H_{t-1} by $\theta_t = \phi v_t + H_{t-1}$: the departure from the standing rule that the central bank must announce. As an algebraic decomposition, the same allocation satisfies

$$G^R_{t-1} = H_{t-1} - (\phi_R - \phi)v_t, \tag{D.11}$$

so the operational departure equals the inherited Ramsey promise plus the rule's error. Under a coarse rule, the component of that error due to an unanticipated current shock cannot have been promised one period earlier. Substituting into (D.10) produces a cross term between the public state and the operational promise whose coefficient is the rule's error times the promise curvature.

The second part of the lemma concerns a local promise problem, which we now specify. At date t the central bank delivers $\theta_t = \phi v_t + H_{t-1}$ and chooses the new operational promise H_t, its announcement for date $t+1$. It privately observes a mean revision m_t to the next shock, $\mathbb{E}_t v_{t+1} = \rho_v v_t + m_t$, the dynamic counterpart of the pure forecast s of the episode; the public forms its expectation from the public shock and the announced promise, $e_t = \phi\rho_v v_t + H_t$. With θ_t pinned by delivery, a change in e_t is absorbed by the output gap, $\kappa x_t = \theta_t - v_t - \beta e_t$, and the continuation is the Ramsey value (D.10) written in the coordinates (v, H). Holding every other term fixed, the local promise problem is

$$\min_{H_t} \frac{\lambda}{2\kappa^2}\left[\theta_t - v_t - \beta(\phi\rho_v v_t + H_t)\right]^2 + \beta\,\mathbb{E}_t\,\mathcal{V}^R(v_{t+1}, H_t). \tag{D.12}$$

Its first-order condition is linear in (m_t, H_t); the lemma reports the slope of its solution

in m_t, the response of the new promise to the news, and the loss from forbidding that response. The calculation holds fixed the public's mapping from the announcement to expected inflation. It identifies the marginal response to private news; the announcement game with persistent private information remains open.

Lemma 20 (Rule error and the use of private forecasts)**.** *In the stationary linear–quadratic Ramsey problem, the cross term of the continuation value between the public shock and the operational promise is*

$$\frac{\partial^2 \mathcal{V}^R}{\partial v \, \partial H} = \Upsilon_R(\phi - \phi_R). \tag{D.13}$$

Suppose the central bank privately observes a mean revision m_t to v_{t+1}. In the local promise problem (D.12)*, the response of the new operational promise to the revision is*

$$\frac{\partial H_t}{\partial m_t} = -\frac{\beta \Upsilon_R(\phi - \phi_R)}{(\lambda/\kappa^2)\beta^2 + \beta \Upsilon_R}, \tag{D.14}$$

and the loss from suppressing this response, per unit of variance of the revision m_t, is

$$\frac{\beta^2 \Upsilon_R^2(\phi - \phi_R)^2}{2[(\lambda/\kappa^2)\beta^2 + \beta \Upsilon_R]}. \tag{D.15}$$

A one-period-ahead forecast is therefore orthogonal to the promise if and only if $\phi = \phi_R$; any rule error makes private forecasts marginally relevant for operational guidance.

Proof. See Appendix D.4.

Equation (D.13) is the dynamic version of Lemma S4 of the supplementary appendix: the cross term is the rule's error scaled by the curvature of the value function in the promise. Equation (D.14) turns it into a behavioral response, the change in the new promise per unit of private revision, whose denominator is the curvature of the promise problem: the cost of holding current inflation at its delivered center through the output gap when expectations move, $(\lambda/\kappa^2)\beta^2$, plus the discounted curvature of the continuation value in the promise, $\beta \Upsilon_R$. Equation (D.15) prices a mandate that forbids the response. Under the ideal dynamic rule, private news changes the level of expected continuation loss but not the marginal return to the promise. Under a coarse rule, the central bank anticipates that the rule will under- or over-react to the future state and uses guidance to repair the forecastable part of the error. The rule therefore also decides whether private forecasting and operational guidance remain orthogonal, in addition to prescribing policy and keeping the audit residual clean.

Delivery and operating implementation. The same audit principle can support an inherited promise, with continuation-adjusted temptation replacing the episode's static curvature. Appendix D.5 derives the required time-varying audit. Finally, a public and credible target path can be translated into nominal interest rates using the real-rate implementation of Holden (2024); Supplementary Appendix 1.1 sets out that mapping. The operating rule implements the target; the announcement and delivery incentives are settled by the mandate of Appendix D.5.

D.3 Proof of Lemma 19

We show that the Ramsey first-order conditions reduce to a one-dimensional state, the normalized multiplier μ_t, whose law of motion yields the rule slope ϕ_R, the inherited promise G_t^R, and the target criterion; the displacement from the ex-post ideal used in (D.9) follows at the end.

Proof. For the recursive representation, attach the multiplier $A\mu_t$ to the date-t Phillips curve and set $\mu_{-1} = 0$; the normalized multiplier μ_t is the one-dimensional state that records past commitments. The first-order conditions are

$$\theta_t + A(\mu_t - \mu_{t-1}) = 0, \tag{D.16}$$

$$\lambda x_t - \kappa A \mu_t = 0. \tag{D.17}$$

Eliminating the multiplier gives the target criterion (D.8). Combining the Phillips curve with the two first-order conditions gives

$$\beta \mathbb{E}_t \mu_{t+1} = \frac{1}{A} v_t + \left(1 + \beta + \frac{\kappa^2}{\lambda}\right) \mu_t - \mu_{t-1}. \tag{D.18}$$

Conjecture $\mu_t = \rho_\mu \mu_{t-1} + c_\mu v_t$ and use $\mathbb{E}_t v_{t+1} = \rho_v v_t$. Matching the coefficients on μ_{t-1} gives $\beta\rho_\mu^2 - (1 + \beta + \kappa^2/\lambda)\rho_\mu + 1 = 0$, whose stable root is the smaller root (D.5); matching the coefficients on v_t gives $c_\mu = -\rho_\mu/[A(1 - \beta\rho_v\rho_\mu)]$. Substituting the law of μ_t into $\theta_t = -A(\mu_t - \mu_{t-1})$ gives

$$\theta_t = -Ac_\mu v_t + A(1 - \rho_\mu)\mu_{t-1} = \phi_R v_t + G_{t-1}^R, \tag{D.19}$$

with $\phi_R = -Ac_\mu = \rho_\mu/(1 - \beta\rho_v\rho_\mu)$ and $G_t^R \equiv A(1 - \rho_\mu)\mu_t$. Multiplying the law of μ_t by $A(1 - \rho_\mu)$ yields (D.6). Finally, the displacement from the no-promise center used in (D.9) follows from the Phillips curve and (D.17), which give $\Xi_t = \theta_t - \kappa^2 A\mu_t/\lambda$. Since $C\kappa^2 A/\lambda = 1$

and $1 - C = 1/A$, $\theta_t - C\Xi_t = \theta_t/A + \mu_t = \mu_{t-1}$ by (D.16), so $\theta_t^R - C\Xi_t = G_{t-1}^R/[A(1-\rho_\mu)]$ and the marginal temptation is $A\mu_{t-1} = G_{t-1}^R/(1-\rho_\mu)$.

D.4 Proof of Lemma 20

The proof has two steps. We first solve the discounted Lyapunov equation of the Ramsey allocation and show that the continuation value is diagonal in (v, G), (D.10), with curvature Υ_R in the promise. We then change coordinates to the operational promise H, which produces the cross term (D.13), and solve the local promise problem (D.12).

Proof. For this calculation write $h = \kappa^2/\lambda$. The stable-root equation implies

$$h = \frac{(1-\rho_\mu)(1-\beta\rho_\mu)}{\rho_\mu}. \tag{D.20}$$

At state (v, G), Ramsey policy and the next inherited promise are

$$\theta = \phi_R v + G, \qquad G' = \rho_\mu G - \phi_R(1-\rho_\mu)v, \qquad x = \frac{\kappa}{\lambda(1-\rho_\mu)}G'. \tag{D.21}$$

Consequently the quadratic stage loss is $(v, G)Q(v, G)'/2$, where

$$Q = \begin{pmatrix} \phi_R^2(1+h) & \beta\rho_\mu\phi_R \\ \beta\rho_\mu\phi_R & (1-\beta\rho_\mu^2)/(1-\rho_\mu) \end{pmatrix}, \qquad T = \begin{pmatrix} \rho_v & 0 \\ -\phi_R(1-\rho_\mu) & \rho_\mu \end{pmatrix}, \tag{D.22}$$

and T is the conditional-mean transition matrix. With mean-zero innovations of state-independent variance, as in the stationary benchmark, innovations contribute only a constant to the quadratic value. Its matrix solves the discounted Lyapunov identity $P = Q + \beta T'PT$. Direct substitution gives

$$P = \begin{pmatrix} \dfrac{\phi_R^2}{\rho_\mu(1-\beta\rho_v^2)} & 0 \\ 0 & \dfrac{1}{1-\rho_\mu} \end{pmatrix}. \tag{D.23}$$

Indeed, the off-diagonal term is $\beta\rho_\mu\phi_R - \beta\phi_R\rho_\mu = 0$, the lower diagonal term is $(1 - \beta\rho_\mu^2)/(1-\rho_\mu) + \beta\rho_\mu^2/(1-\rho_\mu) = 1/(1-\rho_\mu)$, and the upper diagonal identity follows from $1 + h + \beta(1-\rho_\mu) = 1/\rho_\mu$. Stability gives the unique finite quadratic value. Thus $P_v = \phi_R^2/[\rho_\mu(1-\beta\rho_v^2)]$ and the promise curvature is $\Upsilon_R = 1/(1-\rho_\mu)$, as stated in (D.10).

In the notation of Appendix D.1, under the rule ϕv, $G = H - (\phi_R - \phi)v$, so

$$\mathcal{V}^R(v, H) = \frac{1}{2}\left[P_v + \Upsilon_R(\phi_R - \phi)^2\right] v^2 - \Upsilon_R(\phi_R - \phi)vH + \frac{1}{2}\Upsilon_R H^2 + Z. \tag{D.24}$$

The cross-derivative is $\Upsilon_R(\phi - \phi_R)$, proving (D.13). In the local promise problem (D.12), the revision m_t enters only through $\mathbb{E}_t v_{t+1}$ in the continuation value, whose cross term $-\Upsilon_R(\phi_R - \phi)vH$ contributes $\beta\Upsilon_R(\phi - \phi_R)m_t H_t$; the current-period term is free of m_t because the public forms e_t without it. The curvature of the objective in H_t is $Q_H = (\lambda/\kappa^2)\beta^2 + \beta\Upsilon_R$, the first term from the output gap that holds θ_t at its delivered center, the second from the continuation value. Differentiating the first-order condition gives

$$\frac{\partial H_t}{\partial m_t} = -\frac{\beta\Upsilon_R(\phi - \phi_R)}{Q_H}. \tag{D.25}$$

Completing the square gives the loss from suppressing this response, $[\beta\Upsilon_R(\phi - \phi_R)]^2/(2Q_H)$ per unit of variance of m_t. This proves (D.14)–(D.15).

D.5 Delivery of the Inherited Promise

At date t, write $\Delta_t = \theta_t - \bar{\theta}^R(v_t)$ and $a_t = \Delta_t - g_{t-1}$. The audit residual is $z_t = a_t + \varepsilon_t$, with $\varepsilon_t \sim N(0, \sigma^2_{\pi,t})$. Use the obedience audit of Section 4.1.2 with date-t parameters $\Psi_t > 0$, $\tau_t > 0$, $V_t = \tau_t^2 + \sigma^2_{\pi,t}$, and $\kappa_{C,t} = \Psi_t\tau_t/V_t^{3/2}$. The audit parameters used to set the literal promise are public at issuance.

Holding the Ramsey continuation from date $t+1$ fixed, the no-promise center relative to the rule is $(1 - \Upsilon_R/A)G^R_{t-1}$. The delivery problem is

$$\min_{\Delta_t}\left\{\frac{A}{2}\left[\Delta_t - \left(1 - \Upsilon_R/A\right)G^R_{t-1}\right]^2 + \mathcal{A}_{C,t}(\Delta_t - g_{t-1})\right\}. \tag{D.26}$$

At the Ramsey target $\Delta_t = G^R_{t-1}$, the marginal temptation is $\Upsilon_R G^R_{t-1}$, but the stabilization curvature remains A. This distinction is needed to apply the episode's global-optimality condition.

Lemma 21 (Conditional recursive delivery)**.** *Define*

$$\bar{G}_t = \frac{\Psi_t\tau_t}{\Upsilon_R \exp(1/2)\, V_t}. \tag{D.27}$$

For $|G^R_{t-1}| \le \bar{G}_t$*, the equation*

$$\mathcal{A}'_{C,t}(a_t^*) = -\Upsilon_R G^R_{t-1} \tag{D.28}$$

has a unique solution on $[-\sqrt{V_t}, \sqrt{V_t}]$. *The literal promise*

$$g^*_{t-1} = G^R_{t-1} - a^*_t \tag{D.29}$$

supports the target on the responsive branch. If $\kappa_{C,t} < \exp(3/2)A/2$, *this action is the unique global delivery optimum. Let* $\bar{G}$ *be a bound on the inherited promise,* $|G^R_{t-1}| \le \bar{G}$. *The stake at which the ceiling* (D.27) *covers every such promise,* $\bar{G}_t \ge \bar{G}$, *is*

$$\Psi_t \ge \Upsilon_R \exp(1/2)\, \frac{V_t}{\tau_t} \bar{G}. \tag{D.30}$$

For $\sigma_{\pi,t} > 0$, *the minimizing tolerance is* $\tau_t = \sigma_{\pi,t}$, *yielding*

$$\Psi^{\min}_t = 2 \exp(1/2)\, \Upsilon_R \sigma_{\pi,t} \bar{G}. \tag{D.31}$$

At this minimum stake, the global uniqueness condition $\kappa_{C,t} < \exp(3/2)A/2$ *holds if* $\bar{G} < \exp(1)A\sigma_{\pi,t}/(\sqrt{2}\Upsilon_R)$.

Proof. The audit derivative has maximum absolute value $\Psi_t \tau_t / (\exp(1/2) V_t)$ on the responsive range and is strictly increasing there. Comparing this maximum to $\Upsilon_R |G^R_{t-1}|$ gives (D.27). Substitution of (D.29) makes the target satisfy the delivery first-order condition. The global curvature bound from Appendix A.11 applies to (D.26) with A unchanged, establishing uniqueness. Solving the ceiling inequality for the stake gives (D.30); minimizing V_t/τ_t gives (D.31). At that minimum, $\kappa_{C,t}/A = \exp(1/2)\, \Upsilon_R \bar{G} / (\sqrt{2}\, A \sigma_{\pi,t})$, giving the stated sufficient condition.

This construction prices delivery of a specified inherited promise; the audit parameters that size the announcement are public when the promise is made.