Built and Spent by Words:
Forward Guidance under Inflation Targeting

Lukasz A. Drozd
Federal Reserve Bank of Philadelphia

August 2026

The views expressed herein are those of the author and do not necessarily represent those of the Federal Reserve System or the Federal Reserve Bank of Philadelphia.

Two Layers of Central Bank Communication

Inflation targeting

durable communication

- standing message: policy remains anchored to the target
- includes the state-contingent departures the regime permits

Forward guidance

operational communication

- announces a discretionary departure from the inherited rule
- issued shock by shock by the delegated authority

Two Layers of Central Bank Communication

Inflation targeting

durable communication

- standing message: policy remains anchored to the target
- includes the state-contingent departures the regime permits

Forward guidance

operational communication

- announces a discretionary departure from the inherited rule
- issued shock by shock by the delegated authority

Every new statement is filtered through the standing one.

The Layers Clash I: “Run the Economy Hot”

- Lower-bound recession ⇒ promise above-target inflation after liftoff
- The promise conflicts with the inherited target
  - → validating it later can require a costly real contraction
- Public anticipates the waiver temptation ⇒ discounts the promise today

The Layers Clash I: “Run the Economy Hot”

- Lower-bound recession ⇒ promise above-target inflation after liftoff
- The promise conflicts with the inherited target
  - → validating it later can require a costly real contraction
- Public anticipates the waiver temptation ⇒ discounts the promise today
- Aggressive guidance can become **self-defeating**

The Layers Clash II: “Looking Through” a Shock

- Large cost-push shock; central bank privately expects low persistence
- It promises a rapid return to target

  → restrain current expectations and reduce later disinflation costs

- If the shock fades: the central bank is vindicated
- If it persists: honor is contractionary; waiver damages credibility

The Layers Clash II: “Looking Through” a Shock

- Large cost-push shock; central bank privately expects low persistence
- It promises a rapid return to target

  → restrain current expectations and reduce later disinflation costs

- If the shock fades: the central bank is vindicated
- If it persists: honor is contractionary; waiver damages credibility

How much off-rule guidance can the institution sell?

Post-COVID Fact 1: Expectations Stayed Anchored

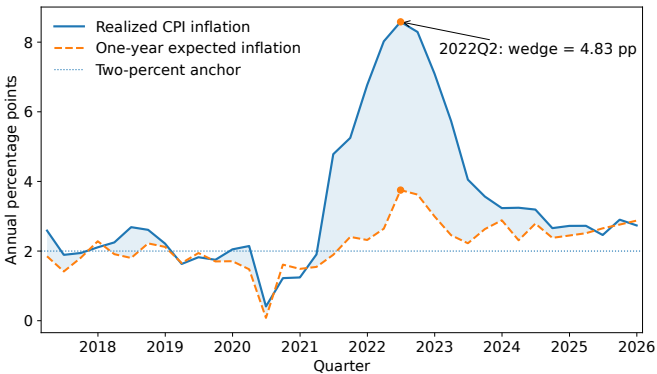

- CPI peaked near 9%; one-year expectations barely breached 4%
- Revision regression:

$$\widehat{w}_0 = 0.756$$

- 95% interval: $[0.696, 0.817]$

Expectations remained remarkably anchored.

Post-COVID Fact 2: Inflation News on FOMC Days Went Quiet

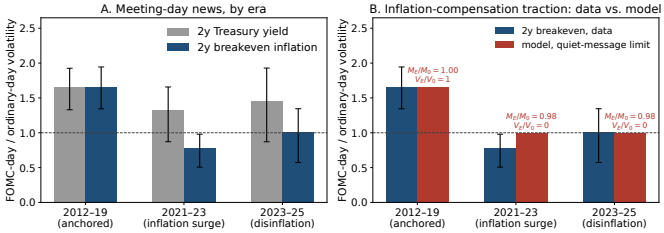

- 2-year inflation-compensation ratio:
- 2012–19: **1.66**
- 2021–23: 0.78

Post-COVID Fact 2: Inflation News on FOMC Days Went Quiet

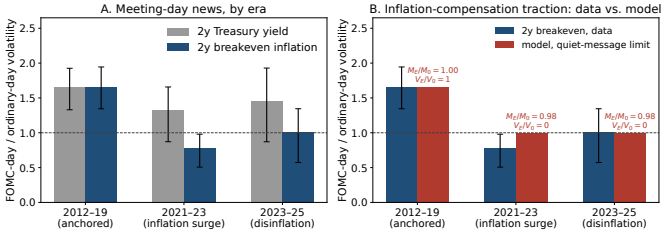

- 2-year inflation-compensation ratio:
- 2012–19: **1.66**
- 2021–23: 0.78

Strong average anchoring need not mean strong marginal traction.

This Paper

- Inflation target: **fixed inherited institution**
  - → public state-contingent default rule around a numerical target
- Forward guidance: **a discretionary announcement with a penalty**
  - → shifts the rule now; creates an accountable future promise
- Hidden institutional modes differ in resistance to waiving the promise

This Paper

- Inflation target: **fixed inherited institution**
  - → public state-contingent default rule around a numerical target
- Forward guidance: **a discretionary announcement with a penalty**
  - → shifts the rule now; creates an accountable future promise
- Hidden institutional modes differ in resistance to waiving the promise

What limits does the target impose on guidance?

What are the policy risks of “looking through the shock” guidance?

Our Approach

- Standard New Keynesian allocation block
- Central bank inherits $\bar{\Theta}(s)$ and privately observes persistent signal v_t
  - $\rightarrow$ communication instrument: one announcement g_t on the rule
- Hidden modes $j_t \in \{\mathrm{G}, \mathrm{R}\}$ differ only in commitment technology
- Public learns twice:
  - $\rightarrow$ from the announcement; from noisy inflation delivery
- $\lambda_t = \Pr(j_t = \mathrm{R})$ is **credibility capital in inverse form**

Preview of Results

1. **Optimized waiver value**: the recursive Bellman gap is locally quadratic, with curvature Ψ_j
2. **Endogenous caution**: expected implementation curvature makes the ordinary mode speak less diffusely, conditionally
3. **Finite communication frontier**: inference and waiver temptation jointly create a Laffer curve of speech
4. **Credibility dynamics**: a message may build credibility in levels while spending it at the margin; delivery failures tether the reset

Related Literature – What We Add

- *Reputation and action learning* (Cukierman–Meltzer; King–Lu; Dovis–Kirpalani)
  - → beliefs move with delivery → **the message itself also moves beliefs**
- *Forward guidance with imperfect credibility*
  - → credibility limit imposed → **generated by the inherited target and waiver incentives**
- *Information effects* (Bassetto)
  - → private information gives speech an independent role

Model

Environment

- Public state s_t: cost-push $u(s_t)$ and natural rate $r^n(s_t)$
- Central bank privately observes persistent signal v_t; public learns it next cycle
- New Keynesian block:

$$\theta_t = \beta\widetilde{\mathbb{E}}_t\pi_{t+1} + \kappa x_t + u(s_t), \qquad x_t = \widetilde{\mathbb{E}}_t x_{t+1} - \frac{i_t - \widetilde{\mathbb{E}}_t\pi_{t+1} - r^n(s_t)}{\sigma}$$

- Realized inflation: $\pi_t = \theta_t + \varepsilon_t^\pi$

  → control noise makes delivery informative but not perfectly revealing

Layer 1: The Inherited Targeting Rule

Inflation targeting contains:

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

(ii) a public state-contingent default center $\bar{\Theta}(s)$

(iii) an ex-post institutional cost of not delivering a shifted center

- The rule is fixed outside the model
- A movement of $\bar{\Theta}(s)$ with the public state is *not* a breach

Layer 2: Forward Guidance Is a Shift on the Rule

- The central bank announces the shifted schedule

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

- Issuing it is costless; its size is disciplined by the credibility it spends
- It becomes next period's inherited promise $g'_{t-1} = g_t$
- Saying nothing means $g_t = 0$: **the central bank reaffirms the rule**
- Public state includes $(s_t, v_{t-1}, \lambda_t, g_{t-1})$

Credibility Capital: Hidden Institutional Modes

- Persistent hidden mode $j_t \in \{\mathrm{G}, \mathrm{R}\}$: ordinary vs. override

  $\rightarrow$ same target, objective, information, and instruments

- Waiver cost:

$$F_{j,t} = \bar{F}_j(\Omega_t) + \sigma_F \nu_t, \qquad \bar{F}_{\mathrm{G}}(\Omega) > \bar{F}_{\mathrm{R}}(\Omega)$$

- Ordinary-mode costs first-order stochastically dominate
- Public belief $\lambda_t = \Pr(j_t = \mathrm{R})$ summarizes institutional doubt

One Communication Cycle

1. Central bank observes $(v_t, j_t, F_{j,t})$
2. It resolves inherited g_{t-1} – honor or waive – and issues new g_t
3. Public observes g_t and updates from **words**
4. Expectations form; a waiver reset is selected conditional on those expectations
5. Inflation realizes with noise; public updates from **delivery**
6. g_t becomes the next inherited promise; the mode then transitions

Expectations: Traction Times the Promise

General local forecast:

$$e = \bar{d} + \tau\,[h - \bar{d}], \qquad \tau = \frac{p}{1 - \delta_W(1 - p)} \in (0, 1]$$

- h: honored center; $\bar{d}$: waiver fallback; p: prospective plan survival

Expectations: Traction Times the Promise

General local forecast:

$$e = \bar{d} + \tau\,[h - \bar{d}], \qquad \tau = \frac{p}{1 - \delta_W(1-p)} \in (0, 1]$$

- h: honored center; $\bar{d}$: waiver fallback; p: prospective plan survival

In the analytical benchmark: $\bar{d} = 0$, $h = g$, $\delta_W = 0$

$$\boxed{m(g) = g\,p(g)}$$

The theory is about how the institution makes $p(g)$ endogenous.

Theory: From Bellman Values to a Frontier

How the Theory Is Closed

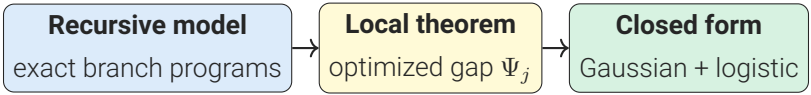

- Frontier results are exact in the approximate benchmark
- The Gaussian/logistic restrictions are *not* imposed on the recursive model

The Waiver Value Comes from Two Optimized Programs

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

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

Both branches choose g **Waiver also chooses θ**

$\mathcal{J}_j$ includes beliefs, λ', and the full continuation value

The Waiver Value Comes from Two Optimized Programs

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

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

Both branches choose g **Waiver also chooses θ**

$\mathcal{J}_j$ includes beliefs, λ', and the full continuation value

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

Result 1: The Local Value of a Waiver

For inherited g_{-1} near the target state:

$$\theta_j^W(g_{-1}) = \bar{\theta} + \omega_j g_{-1} + o(g_{-1}),$$
$$\Delta\mathcal{V}_j(g_{-1}) = \frac{1}{2}\Psi_j g_{-1}^2 + o(g_{-1}^2),$$
$$\Psi_j = \mathcal{A}_j(1-\omega_j)^2 \geq 0$$

- $\mathcal{A}_j$: continuation-adjusted reset curvature after g is reoptimized
- ω_j: pass-through of the inherited promise into the waiver center

The Waiver Gain Is Quadratic

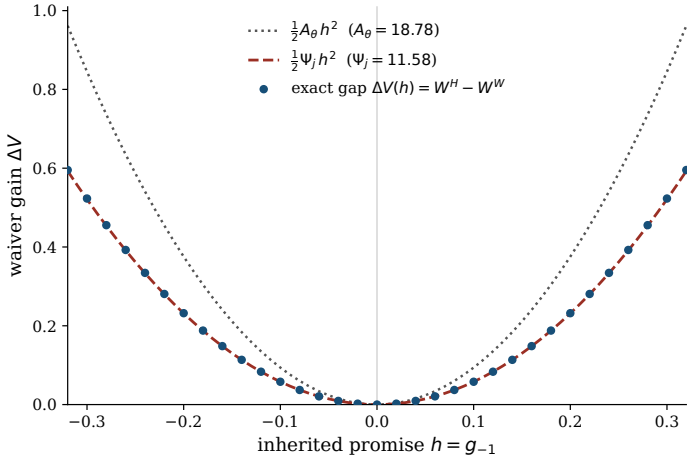

Quadratic by the LQ structure: the exact gap tracks $\frac{1}{2}\Psi_j h^2$, not the steeper $\frac{1}{2}A_\theta h^2$.

What Does Ψ_j Measure?

Reset curvature

$\mathcal{A}_j$

- current stabilization cost
- belief transition
- continuation value
- new announcement already reoptimized

Promise reversal

$1 - \omega_j$

- fraction of g_{t-1} undone by waiver
- if $\omega_j = 1$: no local temptation
- larger reversal $\Rightarrow$ larger waiver value

What Does Ψ_j Measure?

Reset curvature

$\mathcal{A}_j$

- current stabilization cost
- belief transition
- continuation value
- new announcement already reoptimized

Promise reversal

$1 - \omega_j$

- fraction of g_{t-1} undone by waiver
- if $\omega_j = 1$: no local temptation
- larger reversal $\Rightarrow$ larger waiver value

$\Psi_j = \mathcal{A}_j(1 - \omega_j)^2$**: reset curvature times the square of the promise reversal.**

From Waiver Value to the Cost Anticipated at Issuance

$$\widehat{W}_j = W_j^W + \mathbb{E}_F \min\{\Delta\mathcal{V}_j, F_j\}$$

At the target reference state:

$$\boxed{\Xi_j = \chi_{j0}\Psi_j}, \qquad \chi_{j0} = \Pr(F_j \geq 0)$$

- Ψ_j: cost of validating a marginal promise if borne
- χ_{j0}: probability the institution actually bears it by honoring
- Ξ_j: **expected implementation curvature at issuance**

Result 2: Expected Implementation Costs Make Speech Cautious

Local issuance problem:

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

$$\boxed{g_j(\widetilde{v}) = K_j \widetilde{v}, \qquad K_j = \frac{\zeta_j}{\gamma_j + \beta \Xi_j}}$$

- Larger $\Xi_j \Rightarrow$ weaker response to private urgency
- Under common (ζ, γ): $\Xi_{\mathrm{G}} > \Xi_{\mathrm{R}} \Rightarrow K_{\mathrm{G}} < K_{\mathrm{R}}$

Conditional ordering – not implied by Gaussianity alone.

From Caution to a Statistical Target Sieve

$$g \mid j \sim \mathcal{N}(\mu_j, \Sigma_j^2), \qquad \Sigma_j^2 = K_j^2 \sigma_v^2 + \sigma_{g,j}^2$$

- common intercept and communication noise
- $K_{\mathrm{G}} < K_{\mathrm{R}} \Rightarrow \Sigma_{\mathrm{G}} < \Sigma_{\mathrm{R}}$

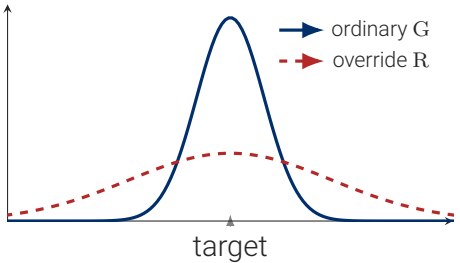

Distance from the rule is immediate Bayesian evidence.

The Closed-Form Bellman–Gaussian Benchmark

Additional restrictions:

- centered local Gaussian messages with $\Sigma_{\mathrm{G}} < \Sigma_{\mathrm{R}}$
- override mode supplies no incremental commitment to the off-rule announcement
- $\Delta\mathcal{V}_{\mathrm{G}}(g) = \Psi_{\mathrm{G}} g^2/2$ extended over the analytical range
- logistic waiver cost with constant location; honored increment $= g$, fallback $= 0$

The Closed-Form Bellman–Gaussian Benchmark

Additional restrictions:

- centered local Gaussian messages with $\Sigma_{\mathrm{G}} < \Sigma_{\mathrm{R}}$
- override mode supplies no incremental commitment to the off-rule announcement
- $\Delta\mathcal{V}_{\mathrm{G}}(g) = \Psi_{\mathrm{G}} g^2/2$ extended over the analytical range
- logistic waiver cost with constant location; honored increment $= g$, fallback $= 0$

The frontier theorem is exact in this approximate economy.

Result 3: A Finite Communication Frontier

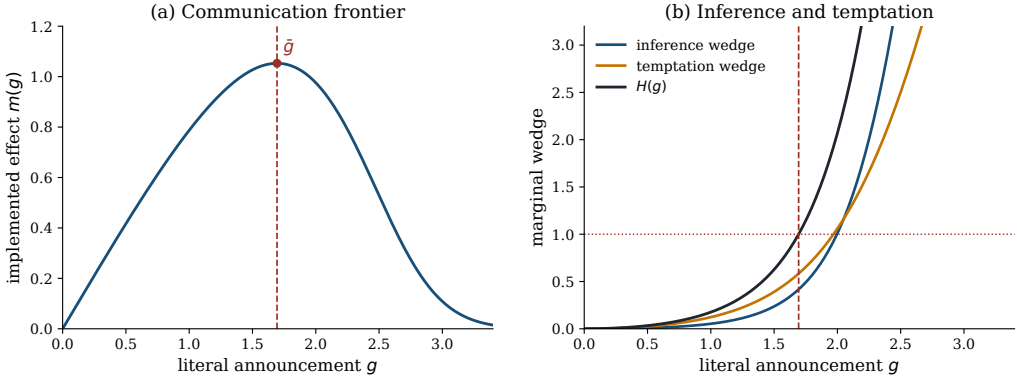

$$m(g) = \frac{g}{[1 + C_\lambda e^{\eta_g g^2}][1 + D_F e^{\kappa_F g^2}]}, \qquad \kappa_F = \frac{\Psi_{\mathrm{G}}}{2\sigma_F}$$

Why the Frontier Is Unique

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

- **Inference**: extreme speech shifts mass toward override
- **Temptation**: a larger promise raises the optimized waiver value

Why the Frontier Is Unique

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

- **Inference**: extreme speech shifts mass toward override
- **Temptation**: a larger promise raises the optimized waiver value

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

Beyond $\bar{g}$, saying more shifts expectations by less.

The Frontier Bounds Optimal Speech

- Every $g > \bar{g}$ is weakly dominated by a smaller announcement to the left

  $\rightarrow$ same implemented shift, smaller literal promise, lower waiver risk

- Hence optimal speech stays interior, below the frontier:

$$0 \leq g^{\star} < \bar{g}$$

- As stabilization urgency grows, $g^{\star}$ approaches $\bar{g}$ from below

The Frontier Bounds Optimal Speech

- Every $g > \bar{g}$ is weakly dominated by a smaller announcement to the left

  $\rightarrow$ same implemented shift, smaller literal promise, lower waiver risk

- Hence optimal speech stays interior, below the frontier:

$$0 \leq g^{\star} < \bar{g}$$

- As stabilization urgency grows, $g^{\star}$ approaches $\bar{g}$ from below

A deeper crisis exhausts communication capacity; it does not create more.

Result 4: Credibility in Levels vs. at the Margin

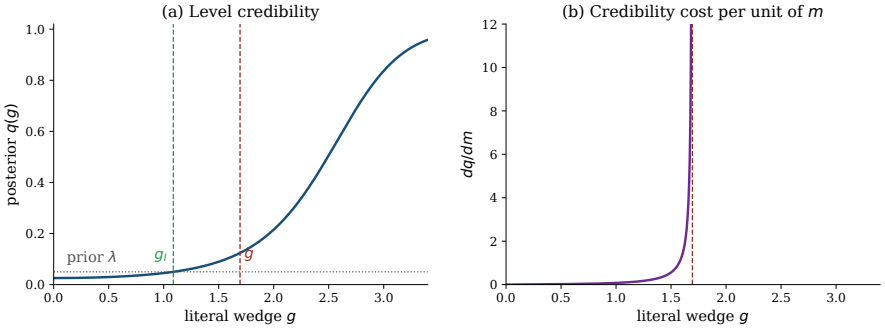

- $|g| < g_I$ may **build in levels**; increasing any $g > 0$ **spends at the margin**
- as $g \uparrow \bar{g}$, credibility cost per unit of m diverges

Actions Matter: A Delivery Failure Damages Beliefs

Let Θ^H be the honored center and θ the waiver reset.

$$\mathbb{E}_\theta(q^\pi - q) = \underbrace{\mathcal{B}(q, \chi_{\mathrm{G}}, \chi_{\mathrm{R}}, \sigma_\pi)}_{>0}(\theta - \Theta^H)^2 + o((\theta - \Theta^H)^2)$$

- larger miss $\Rightarrow$ outcome is more diagnostic of the low-resistance mode
- belief damage is second order because small misses are obscured by control noise
- persistence ρ_λ carries the damage into the next prior

Purely Bayesian: no policy sign yet.

A Delivery Failure Can Tether the Reset

Under $W^c_{j,\lambda} > 0$ and $W^c_{j,\lambda\lambda} \geq 0$:

$$\boxed{\Theta^W_j = \frac{A_\theta d(e,u) + \varrho_j \Theta^H}{A_\theta + \varrho_j}}$$

- $d(e,u)$: static center; ϱ_j: reputational curvature
- larger value of credibility $\Rightarrow$ reset stays closer to the promise

A Delivery Failure Can Tether the Reset

Under $W^c_{j,\lambda} > 0$ and $W^c_{j,\lambda\lambda} \geq 0$:

$$\boxed{\Theta^W_j = \frac{A_\theta d(e,u) + \varrho_j \Theta^H}{A_\theta + \varrho_j}}$$

- $d(e,u)$: static center; ϱ_j: reputational curvature
- larger value of credibility $\Rightarrow$ reset stays closer to the promise

$$\omega_j = \frac{\varrho_j}{A_\theta + \varrho_j}$$

Result 1 leaves ω_j unsigned; the value restriction signs it.

Result 5: Low-Noise Institutional Fragility

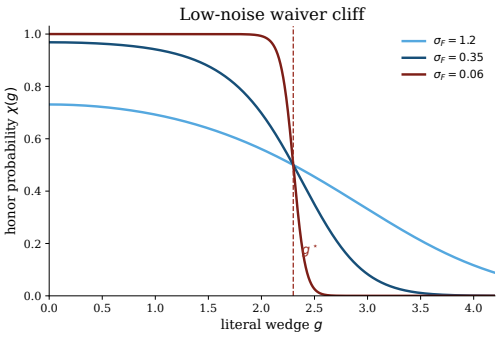

$$g^\star = \sqrt{\frac{2\bar{F}_{\mathrm{G}}}{\Psi_{\mathrm{G}}}}$$

- $\sigma_F \downarrow 0$: threshold rule
- below $g^\star$: predictable honor
- at $g^\star$: order-one collapse

$$\bar{g}(\sigma_F) \longrightarrow \min\{\bar{g}_I, g^\star\}$$

When Is a Positive Gate Desirable?

- **Aligned objectives, full enforcement**: the optimal gate is zero
  - → $g = 0$ already reproduces silence inside the feasible set
- **Two rationales for a positive gate**:
  - → *Screening*: a delegated authority may overvalue moving expectations and speak when society prefers the default
  - → *Opacity commitment*: opacity can be optimal ex ante, yet the "perfect-offset" type breaks pooling interim

Opacity optimal when $\beta^2\sigma_\mu^2 > \frac{h}{D}(y + \beta\bar{\mu})^2$; the gate that locks it is $\mathcal{F}^0 = J^0$. A commitment device even under aligned objectives, not a bias correction.

Communication and the Interest-Rate Margin

Local cost-push allocation:

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

Current rate / output gap

$$\vartheta x = -\alpha \kappa \pi$$

Forward guidance

$$\alpha \beta \bar{\tau} \pi = 0 \;\Rightarrow\; \pi = 0$$

Costless guidance offsets the shock ($\pi = x = 0$); the rate is reserved for natural-rate disturbances.

Credibility Also Relaxes the Implementation Boundary

Suppose target-reaffirming communication anchors fraction $\bar{\tau}$ of the forecast. Under $i_t = \phi_\pi \pi_t$:

$$\boxed{\text{local determinacy} \iff \phi_\pi > 1 - \bar{\tau}}$$

- $\bar{\tau} \to 0$: standard Taylor principle $\phi_\pi > 1$
- $\bar{\tau} \to 1$: strong inherited credibility substitutes for current-rate aggression

An implementation condition – not an extra constraint on the institutional allocation.

What Theory Can – and Cannot – Say About Myopia

- Myopia omits future effects of today's announcement on:
  - → public beliefs and future communication capacity
  - → honoring incentives and the inherited promise
- These continuation effects need not have the same sign in every state

What Theory Can – and Cannot – Say About Myopia

- Myopia omits future effects of today's announcement on:
  - → public beliefs and future communication capacity
  - → honoring incentives and the inherited promise
- These continuation effects need not have the same sign in every state
- No general theorem that a myopic central bank always talks more
- The stationary ranking is a **quantitative result** of the recurring-ZLB application

Applications

What the Applications Compute – and What They Take as Primitive

Post-pandemic mechanism

- computes local optimized $\widehat{\Psi}$
- calibrated / frozen Gaussian message likelihoods
- misreading exercises price a fixed policy path

Recurring-ZLB mechanism

- primitive temptation $\psi g^2/2, \psi = 0.18$
- frozen local message coefficients
- forward-looking vs. repeated one-off policy

Separate mechanism studies – not one cross-application calibration.

The Pandemic Application Now Uses the Optimized Gap

Object	Value
A_θ	18.78
$\widehat{\mathcal{A}}$	11.72
$\widehat{\Psi}$	11.58
$\widehat{\Psi}/A_\theta$	0.62

- $g_{2021} = 0.40$
- $\bar{g} = 1.49$
- box upper bound $= 2.00$
- $M_E/M_0 = 0.976$

Reoptimization lowers temptation; the baseline is not at the frontier.

Application 1: A Look-Through Path for the Inflation Surge

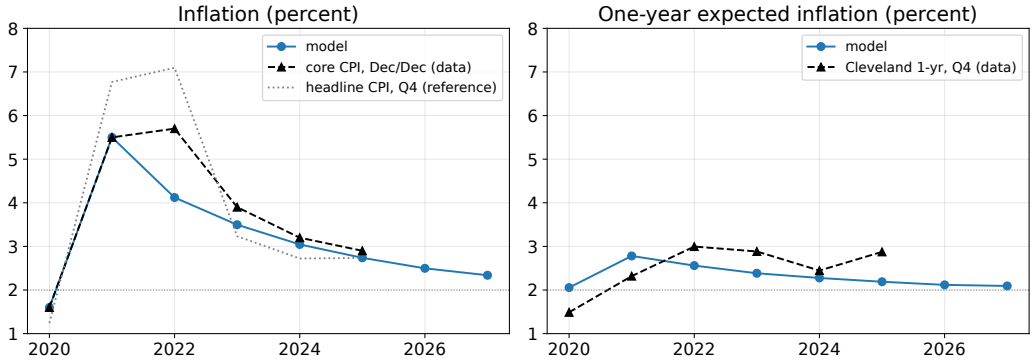

- model peak: inflation 5.50, one-year expectations 2.78
  - → core CPI Dec/Dec; 2020–21 matched, 2022–25 untargeted (RMSE 0.82 pp)
- all promises honored; baseline cumulative loss 43 vs. silent default-target loss 69

The Risk Is the Public's Reading of the Same Words

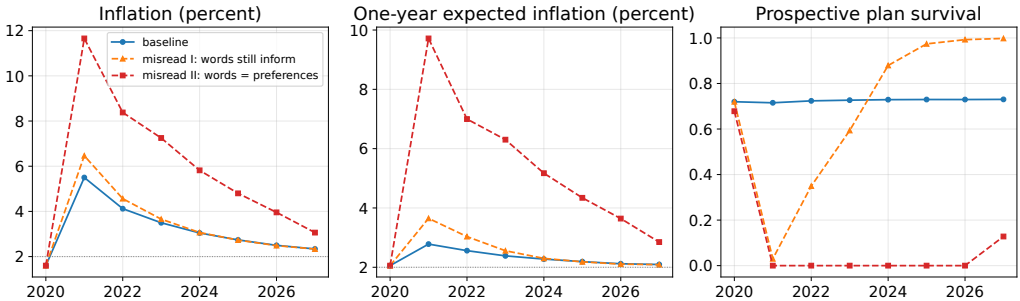

Baseline → commitment discounted → words read as unstable preferences.
The policy path is fixed: these rows price a path; they are not equilibria.

One Path, Six Readings: Immaculate Disinflation to Stagflation

Public reading of the *same* announced words	Peak π	Peak $E\pi$	Cum. loss
Vindicated – shock fades	5.50	2.78	33
Baseline – words inform and commit	5.50	2.78	43
Misreading I – commitment discounted	6.46	3.65	56
Silent default target – no words	5.24	2.55	69
Rule-retreat – anchor abandoned	6.49	3.67	87
Misreading II – words read as preferences	11.66	9.72	217

- Identical announced path; only the public's institutional reading changes
- Peak $E\pi$ spans 2.78 to 9.72 percent – the value of the second anchor

One look-through, priced from immaculate disinflation to 1970s stagflation.

Back to the Data: Traction vs. Surprise Variance

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

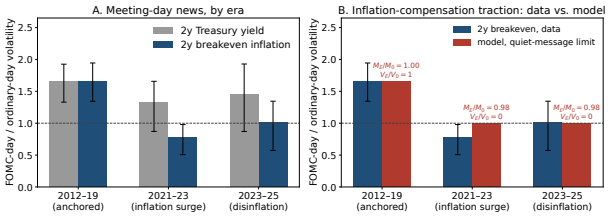

- solved $M_E/M_0 = 0.976$
- traction-only prediction: 1.63
- quiet-message limit: 1.00
- data: $1.66 \rightarrow 0.78$

Most of the collapse is unusually predictable guidance – not zero traction.

▸ Second application: policy myopia at a recurring lower bound

Takeaways

1. Waiver temptation is an **optimized Bellman value**: $\Delta \mathcal{V}_j \simeq \Psi_j g_{-1}^2/2$
2. Expected implementation curvature can make the ordinary mode **speak cautiously**
3. Inference $+$ temptation create a finite frontier in the **Bellman–Gaussian benchmark**
4. Words may **build in levels** while **spending at the margin**; delivery can tether the reset
5. Application: the outcome depends on how the public reads the same words

Takeaways

1. Waiver temptation is an **optimized Bellman value**: $\Delta \mathcal{V}_j \simeq \Psi_j g_{-1}^2/2$
2. Expected implementation curvature can make the ordinary mode **speak cautiously**
3. Inference + temptation create a finite frontier in the **Bellman–Gaussian benchmark**
4. Words may **build in levels** while **spending at the margin**; delivery can tether the reset
5. Application: the outcome depends on how the public reads the same words

The central bank never speaks into a vacuum – it speaks against its own standing rule.

Backup

Backup: Second Application – Policy Myopia at a Recurring Lower Bound

- Two-state economy; recession natural-rate gap $r^n = -1.80$
- At the bound: expected inflation is the active stabilization margin
- Same communication institution and public likelihood schedule
- **Forward-looking**: values future credibility and honoring incentives
- **Repeated one-off**: continuation weight set to zero
- Working model takes $\psi = 0.18$ and local message coefficients as primitives

Backup: Myopia Works Through the Stationary Credibility Distribution

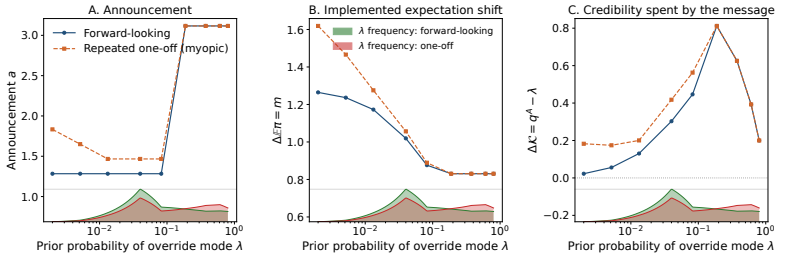

- mean λ: $0.26 \to 0.35$
- loss: $+5.0\%$
- persistent baseline: more talk, less average effect
- $\rho_v = 0$: less talk, still higher loss

Robust result: the welfare ranking, not a universal speech ranking.

Backup: Full Optimized-Gap Formula

Second-order branch objective in reset deviation δ, new announcement g, and inherited announcement h:

$$\tfrac{1}{2}\left[a_j\delta^2 + 2b_j\delta g + c_j g^2 + 2r_j\delta h + 2s_j g h + t_j h^2\right]$$

After optimizing both branches:

$$\boxed{\Psi_j = \frac{(r_j - b_j s_j/c_j)^2}{a_j - b_j^2/c_j}}$$

- denominator: reset curvature after projecting out g
- numerator: inherited-announcement loading after the same projection

Backup: The Pre-Draw Maximum

$$\widehat{W}_j(h) = W_j^W(h) + \mathbb{E}_F \min\{\Gamma_j(h), F_j\}, \qquad \Gamma_j(h) = \Delta\mathcal{V}_j(h)$$

For continuous F_j:

$$\frac{\mathrm{d}}{\mathrm{d}h}\mathbb{E}\min\{\Gamma_j(h), F_j\} = \Gamma_{j,h}(h)\Pr(F_j \geq \Gamma_j(h))$$

At $h = 0$:

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

Backup: Closed-Form Posterior, Honor, and Survival

$$q(g) = \frac{C_\lambda e^{\eta_g g^2}}{1 + C_\lambda e^{\eta_g g^2}}, \qquad C_\lambda = \frac{\lambda}{1-\lambda}\frac{\Sigma_{\mathrm{G}}}{\Sigma_{\mathrm{R}}}$$

$$\chi(g) = \frac{1}{1 + D_F e^{\kappa_F g^2}}, \qquad \kappa_F = \frac{\Psi_{\mathrm{G}}}{2\sigma_F}$$

$$p(g) = [1 - q(g)]\chi(g), \qquad m(g) = gp(g)$$

Backup: Credibility Threshold and Marginal Cost

$$g_I^2 = \frac{2\log(\Sigma_{\mathrm{R}}/\Sigma_{\mathrm{G}})}{\Sigma_{\mathrm{G}}^{-2} - \Sigma_{\mathrm{R}}^{-2}}$$

- $|g| < g_I$: message lowers override probability relative to the prior
- for every $g > 0$: $q_g(g) = 2\eta_g g q(g)[1 - q(g)] > 0$

On the effective branch:

$$\frac{\mathrm{d}q}{\mathrm{d}m} = \frac{2\eta_g g q(g)}{\chi(g)[1 - H(g)]} \longrightarrow \infty \quad \text{as } g \uparrow \bar{g}$$

Backup: Distant Forward Guidance

Suppose an extreme repeated message contributes adverse log-likelihood increment $\ell^g(g) > 0$ each period. If the raw New Keynesian horizon multiplier is bounded by $\bar{M}e^{\gamma n}$ and

$$\ell^g(g) > \gamma,$$

then the credibility-adjusted effect converges to zero as the horizon $n \to \infty$. **Moving an institutionally implausible promise farther out does not rescue it.**

Backup: Delivery Belief Damage

Conditional on an actual waiver and small miss $\Delta_\theta = \theta - \Theta^H$:

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

where

$$\bar{\chi}_q = (1-q)\chi_{\mathrm{G}} + q\chi_{\mathrm{R}}.$$

The next prior inherits the same second-order increase, scaled by ρ_λ.

Backup: Ex-Ante Opacity and the Gate

Opacity vs. revelation, ex ante:

$$\mathbb{E}J_{\mu}^{*} - J^{0} = \frac{\alpha}{2(D+h)}\left[\beta^{2}\sigma_{\mu}^{2} - \tfrac{h}{D}(y+\beta\bar{\mu})^{2}\right]$$

Opacity is optimal when the bracket is positive. The perfect-offset type $\mu_0 = -y/\beta$ then gains $\Delta(\mu_0) = J^0$ by revealing, so pooling unravels unless the gate satisfies $\mathcal{F}^0 = J^0$.

Cost-push mix:

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

Backup: Determinacy Calculation

Let $\delta = 1 - \bar{\tau}$ scale the private forecast in the NKPC and IS equations. Under $i_t = \phi_\pi \pi_t$:

$$\mathbb{E}_t \begin{pmatrix} \pi_{t+1} \\ x_{t+1} \end{pmatrix} = \begin{pmatrix} (\beta\delta)^{-1} & -\kappa(\beta\delta)^{-1} \\ (\phi_\pi - \beta^{-1})/\sigma & 1 + \kappa/(\sigma\beta) \end{pmatrix} \begin{pmatrix} \pi_t \\ x_t \end{pmatrix}$$

The Jury conditions reduce to

$$\phi_\pi > \delta = 1 - \bar{\tau}.$$

Backup: Theory–Application Bridge in 2021

Object	Theory role	Application value
A_θ	current-loss curvature	18.78
$\widehat{\Psi}$	optimized Bellman curvature	11.58
$\widehat{\Psi}/A_\theta$	reoptimization effect	0.62
g_{2021}	issued announcement	0.40
$\bar{g}$	working-model peak	1.49
M_E/M_0	marginal-traction ratio	0.976

- $\widehat{\Psi}$ replaces the invalid current-loss temptation object
- message likelihoods remain calibrated local approximations
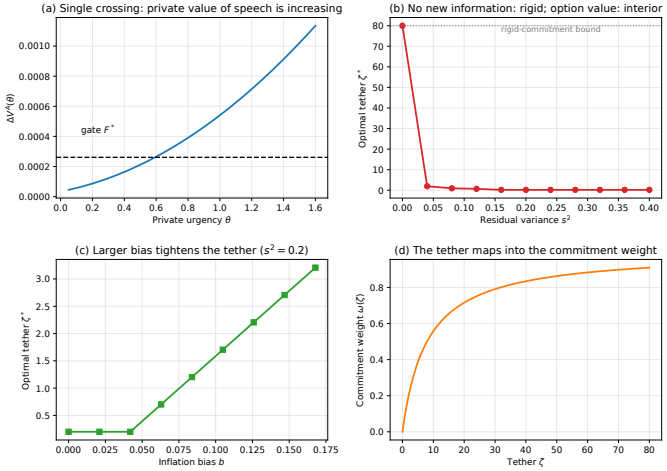
Backup: The Gate and the Tether