Spending the Anchor: Modeling Monetary Policy Response to the Covid Inflation*

Lukasz A. Drozd
Federal Reserve Bank of Philadelphia
lukasz.drozd@phil.frb.org

Marina M. Tavares
International Monetary Fund
marinamendestavares@gmail.com

August 2026

Abstract

We model the monetary policy response to the COVID recession. The recession destroys workforce-embodied organizational capital, which must be rebuilt at a convex cost before potential output recovers; the rebuilding pressure acts as a persistent cost-push shock. The central bank operates an inherited targeting rule and communicates future inflation under private information: it receives a signal that the shock will fade quickly and truthfully announces a fast return of inflation toward target. The model replicates the salient features of the U.S. recovery—high realized inflation alongside anchored one-year expectations and little marginal traction from meeting-day communication. We show that the strategy is fragile: if the public partly anticipates the shock's persistence, the same announcement loses almost all credibility and expected inflation jumps.

Keywords: COVID inflation, central-bank credibility, forward guidance, inflation expectations, organizational capital, optimal monetary policy.
JEL codes: E31, E32, E52, E58, D82, D83.

*Corresponding author: Lukasz A. Drozd, Federal Reserve Bank of Philadelphia, Research Department, 10 N. Independence Mall West, Philadelphia, PA 19106. We thank Urban Jermann, João Gomes, Jonas Arias, Makoto Nakajima, and Alessandro Dovis for comments. The views expressed herein are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia, the Federal Reserve System, the International Monetary Fund, or their respective management. The quantitative exercise is illustrative and is not an estimate of the historical optimum.

1 Introduction

The Federal Reserve entered the post-COVID inflation with a strong inflation-fighting record, delayed conventional tightening, and repeatedly communicated that inflation would return toward two percent as supply disruptions faded. We treat that communication as part of policy: it moved expectations, and we model how. Three questions organize the paper. Can private information rationalize communication that downplays the persistence of inflation? Does such communication meaningfully change the recovery path, bringing it closer to the data? And what policy risks does the strategy carry?

The model combines two mechanisms in a simple New Keynesian optimal policy problem. The supply side follows Drozd and Tavares (2026, hereafter DT26b): the COVID recession unevenly destroys organizational capital embodied in workforces, and restoring it requires a convex rebuilding investment whose cost pressure reaches the Phillips curve as a persistent cost-push term. Recovery is therefore an investment process. Tight policy lowers current demand and marginal cost, but it also lowers the return to rebuilding and prolongs the recovery—a reason to partly look through the shock. The communication side follows Drozd (2026, hereafter DT26a): an inflation-target-constrained authority can shift its rule by announcing future inflation, at an authorization cost. The public rationally discounts a promise the institution is unlikely to deliver, which makes the traction of communication endogenous. Delivery of the outstanding commitment is disciplined by a graduated, bounded public audit rather than by an all-or-nothing waiver decision.

Timing is central. At date t the central bank announces the conditional center of inflation at $t + 1$. The announcement is a promise, not a forecast, and it embeds a privately observed signal v_t about the persistence of the supply shock. The baseline calibration makes the 2021 signal optimistic: its implied one-year persistence is 0.344, whereas realized supply pressure falls by only a factor of 0.883 between 2021 and 2022. The bank therefore announces a low future center. The promise comes due against a graduated public audit: the institution delivers the center that balances its desired accommodation against the expected institutional consequence of shading the announced letter, and delivering the letter in a high-supply-pressure state requires a costly output contraction.

Pricing the promise this way rules out a costless expectations gain: expectations price the posterior-weighted inflation centers that the two hidden institutional modes

will deliver, and the planner's date-$(t+1)$ continuation contribution weights the same delivered centers by the same posterior,

$$(1-q_t)L^G_{t+1} + q_t L^R_{t+1},$$

together with the expected institutional cost of shading the letter—the same objects that determine private expectations. The planner cannot obtain the expectations benefit of a low announcement while independently selecting a high-inflation continuation path.

Policy also feeds back into credibility. The public cannot observe whether the institution currently operates in its ordinary mode or in an override mode—a low-resilience institution persistently less exposed to the public audit. Both modes deliver; they share the target and the stabilization objective and differ only in their exposure to the same bounded graduated sanction, $\chi_G > \chi_R$. This hidden-mode structure drives both how policy spends credibility over time and how much static traction a given announcement has on expectations.

The baseline reproduces the central post-COVID observation—the wide gap between realized inflation and short-horizon expectations. Inflation reaches 7.40 percent in 2021, while one-year expected inflation is 3.12 percent. The bank announces a 3.01-percent future center—an elevated glide path, not a mechanical two-percent promise. The prospective delivered share of the promise bottoms out at 0.687. The delivery accounting makes the cost of the promise transparent: conditional on the signal-implied supply path, enforcing the letter of the 2021 promise requires a 2022 output gap of about -2.07 percent, while the weak-mode delivered center requires far less contraction.

The private signal is essential to the interpretation. The bank expects supply pressure to fade and chooses policy accordingly; the adverse 2022 supply innovation arrives only after the conditional center and the conventional action are in place, so it appears as an implementation surprise around an otherwise delivered plan. The realized 2022 state is therefore evidence that the bank's internal assessment proved too optimistic, not that it knowingly announced a false path.

The fragility experiment moves part of that persistence into the public information set before implementation. If the public partly "sees through" the transitory-shock assessment, it anticipates that the promise will come due in a high-supply-pressure state, where the cost of delivering the letter is severe. Holding the robust announcement and the conventional policy path fixed, the prospective delivered share falls to

0.013, expected inflation rises to 3.87 percent, and inflation rises to 8.12 percent. The arithmetic behind the collapse is stark: enforcing the letter in the realized 2022 state would require an output gap of -18.4 percent, while the delivered center requires only -1.0 percent, so the public expects the promise to be shaded to discretion. The cliff is the delivery problem's own enforcement threshold—past it, the bounded sanction no longer holds delivery to the letter—not a reporting adjustment. The strategy is fragile because it relies on the central bank correctly reading the public's inference problem—a demanding requirement in practice.

The remainder of the paper proceeds as follows. Section 2 presents the communication and supply blocks and states the branch-consistent policy problem. Section 3 describes the calibration, including the private persistence signal. Section 4 reports the baseline and fragility exercises. Section 5 presents the market diagnostic for marginal communication traction. Section 6 concludes. The online appendix gives the full numerical algorithm, branch equations, validation tests, and data construction.

2 Model Framework

2.1 Communication, private information, and a binding future promise

Inflation targeting and the guidance wedge. Normalize the long-run inflation target to zero. Conditional on the disclosed supply outlook, the inherited targeting regime specifies a default future inflation center $\bar{a}_t$. In the annual application,

$$\bar{a}_t = \phi_{\mathrm{IT}}\, \widehat{u}^{CB}_{t+1|t}, \tag{1}$$

where u_t denotes supply pressure and $\widehat{u}^{CB}_{t+1|t}$ is the bank's signal-implied forecast disclosed with the announcement. The central bank announces

$$a_t = \bar{a}_t + g_t, \tag{2}$$

where g_t is the forward-guidance wedge. A negative wedge promises a faster return toward target than the inherited rule.

Private signal. Supply pressure is publicly observed at date t. The central bank additionally observes a scalar signal

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

which maps into perceived persistence through

$$\rho_t^u = \Lambda(\alpha_v + \gamma_v v_t), \qquad \widetilde{\mathbb{E}}_t^{CB} u_{t+1} = \rho_t^u u_t. \tag{4}$$

Here $\Lambda(z) \equiv (1 + e^{-z})^{-1}$. The signal is informative about the conditional distribution, not the future realization. A new innovation can therefore leave u_{t+1} far above the bank's forecast without making the date-t announcement dishonest.

For transparency, the application treats the signal-implied supply forecast as disclosed through the bank's macroeconomic projections. The Delphic component of communication is therefore observed, and the one-dimensional guidance wedge g_t is the remaining Odyssean message evaluated against the targeting rule. The companion theory allows the public to infer the signal jointly with the institutional mode; the present full-disclosure specialization isolates the delivery mechanism without adding a second message dimension.

Message inference and hidden enforcement exposure. The public is uncertain whether the ordinary mode G or override mode R governs the policy episode. Conditional message densities are locally Gaussian in the wedge,

$$g_t \mid j \sim \mathcal{N}(0, \sigma_j^2), \qquad \sigma_R > \sigma_G > 0. \tag{5}$$

If λ_t is the prior probability of mode R, the posterior after words is

$$q_t = \frac{\lambda_t f_R(g_t)}{(1 - \lambda_t) f_G(g_t) + \lambda_t f_R(g_t)}. \tag{6}$$

Delivery is disciplined by continuous public monitoring, following the companion theory. Shading the announced letter by w draws a graduated, bounded expected sanction

$$\mathcal{K}_j(w) = \chi_j \bar{F} \left[1 - \frac{\tau}{\sqrt{V}} e^{-w^2/2V} \right] + \frac{\bar{\varrho}}{2} w^2, \qquad V = \tau^2 + \sigma_\pi^2, \qquad \chi_G > \chi_R > 0, \tag{7}$$

where $\bar{F}$ is the bounded formal consequence, τ the audit tolerance, σ_π control noise, and $\bar{\varrho}$ a reputational curvature: formal and reputational discipline add, and posterior damage keeps growing where the bounded formal audit saturates. The hidden modes differ only in their exposure χ_j to the same audit.

The promise comes due. The announcement a_t is the letter of the promise for date $t+1$. Enforcing the letter, given the expected supply state and the next-period expectation e_{t+1}, requires

$$x_{t+1}^H = \frac{a_t - \beta e_{t+1} - u_{t+1}}{\kappa}. \tag{8}$$

Let d_{t+1} denote the static discretionary center—what the institution would deliver with no promise outstanding. Mode j delivers the center that settles the tension between the letter and the desire,

$$\theta_{t+1}^j = \operatorname*{argmin}_\theta \ \frac{\Psi}{2}\big(\theta - d_{t+1}\big)^2 + \mathcal{K}_j\big(\theta - a_t\big) = d_{t+1} + \operatorname{sign}(a_t - d_{t+1})\, P_j\big(|a_t - d_{t+1}|\big), \tag{9}$$

where Ψ is the continuation-adjusted temptation curvature—the reduced (Schur-complement) curvature of the branch objective once the new announcement is re-optimized, below the period-loss value $A_\theta = 1 + \vartheta/\kappa^2$—and P_j is the delivered-pull map generated by the first-order condition of (9): the same branch condition, delivered share, and promise ceiling as the companion theory. Because the formal sanction is bounded, (9) has a global-minimum switch: once the letter–desire gap crosses an enforcement threshold, delivery collapses from the high-share branch to near-complete shading, leaving only the reputational tether. The mode-j *delivered share* and its posterior blend,

$$s_{j,t} = \frac{P_j\big(|a_t - d_{t+1}|\big)}{|a_t - d_{t+1}|}, \qquad s_t = (1 - q_t)s_{G,t} + q_t s_{R,t}, \tag{10}$$

replace the predecessor technology's honoring probabilities and plan survival. The delivered-center losses are

$$L_{t+1}^j = \frac{1}{2}\big(\theta_{t+1}^j\big)^2 + \frac{\vartheta}{2}\big(x_{t+1}^j - \bar{x}_{t+1}\big)^2, \qquad j \in \{G, R\}, \tag{11}$$

where x_{t+1}^j implements θ_{t+1}^j through the conditional-mean Phillips curve and $\bar{x}_{t+1}$ is the capacity-related activity adjustment defined in Section 2.3.

Rational expectations and delivery-consistent welfare. Private expectations price the posterior-weighted delivered centers, equivalently the delivered-share blend of the letter and the desire:

$$e_t \equiv \widetilde{\mathbb{E}}_t \pi_{t+1} = (1 - q_t)\theta^G_{t+1} + q_t \theta^R_{t+1} = s_t a_t + (1 - s_t) d_{t+1}. \tag{12}$$

Accordingly, the date-$(t+1)$ contribution to the planner's expected loss weights the same delivered centers by the same posterior and adds the expected institutional cost of shading,

$$\big(1 - q_t\big) L^G_{t+1} + q_t L^R_{t+1} + \big(1 - q_t\big)\mathcal{K}^{\text{net}}_G + q_t \mathcal{K}^{\text{net}}_R, \tag{13}$$

where $\mathcal{K}^{\text{net}}_j$ is the sanction (7) at the delivered shading, net of its zero-shading level. Equations (12) and (13) are the central consistency restriction: the same delivered centers and posterior that discount the promise in expectations also price the continuation in welfare. The sanction term matters: without it the planner could issue a letter so far from the desire that delivery collapses, neutralizing its own communication for free.

Inflation outcomes update the institutional posterior through the mode-specific delivered densities

$$L_j(\pi_{t+1}) = \phi_{\sigma_H}\big(\pi_{t+1} - \theta^j_{t+1}\big), \tag{14}$$

so the delivery log-likelihood ratio is linear in the realized outcome. The post-outcome posterior and next prior satisfy

$$\operatorname{logit} q^D_t = \operatorname{logit} \lambda_t + \log \frac{f_R(g_t)}{f_G(g_t)} + \log \frac{L_R(\pi_{t+1})}{L_G(\pi_{t+1})}, \tag{15}$$

$$\lambda_{t+1} = (1 - \rho_\lambda)\bar{\lambda} + \rho_\lambda q^D_t. \tag{16}$$

2.2 Supply: organizational capital, destruction, and rebuilding

The supply side is a compact annual version of DT26b. We first state the environment and the equations used in the application. Online Appendix Section A maps the compact objects to the parent production model and derives the rebuilding conditions.

Environment. Time is annual and the economy has three layers. *Wholesale producers* come in types $j = 1, \ldots, N_J$ with gross-output weights ω_j, representing industries that differ only in how hard COVID hit them. Each type produces a common intermediate

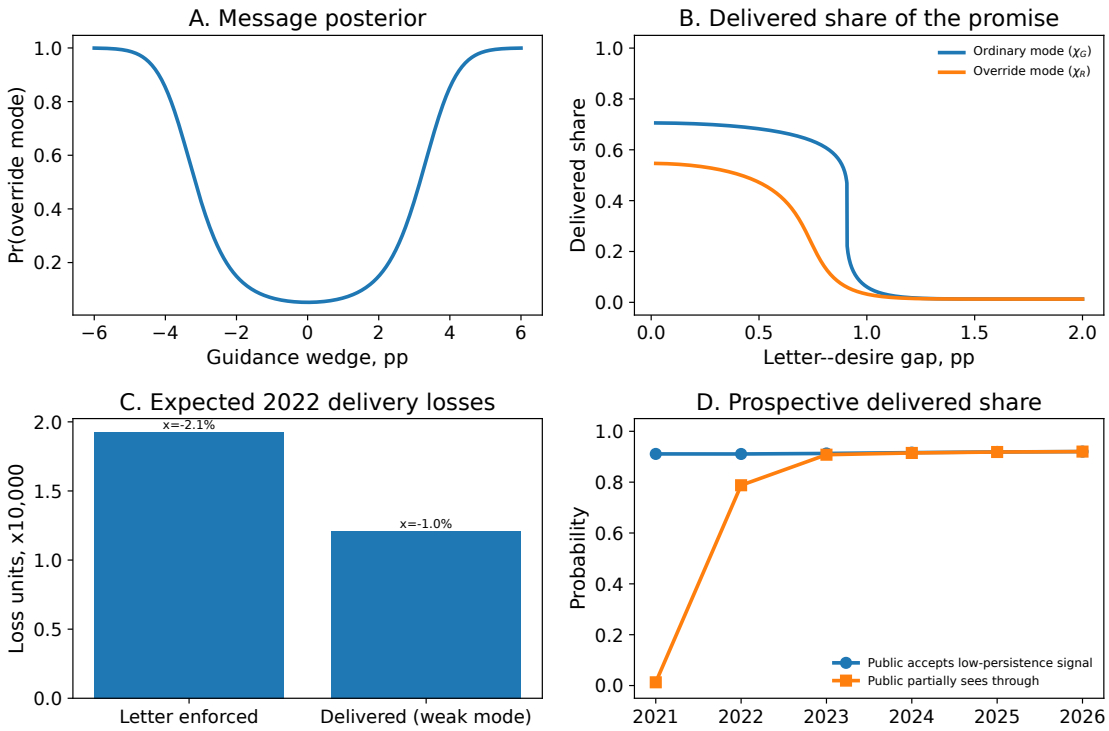

Figure 1: The delivery-consistent communication mechanism

Notes: Panel A shows message-stage inference from the guidance wedge. Panel B plots each mode's delivered share of the promise as the letter–desire gap widens; the collapse is the enforcement threshold of the bounded audit. Panel C compares the 2022 losses from enforcing the letter and from the weak-mode delivered center, priced by the 2021 planner. Panel D shows the prospective delivered share when the public accepts the low-persistence signal and when it partly anticipates the subsequent persistence surprise.

good with a constant-elasticity-of-substitution technology (elasticity ρ, returns γ to flexible inputs) over two inputs: effective labor and materials. Production is roundabout, in the spirit of Basu (1995): materials are units of the aggregate final good bought back from retail, so every producer is simultaneously every other producer's customer, and this circularity stands in for the input–output network of the parent model. *Retailers* are monopolistically competitive: they buy the wholesale composite, differentiate it with demand elasticity ϵ_p (gross markup $\epsilon_p/(\epsilon_p - 1)$), and set prices subject to Rotemberg quadratic adjustment costs with coefficient Φ_p—the source of the Phillips curve in Section 2.3, with slope κ_p. *Households* supply labor and consume; their behavior is summarized by a standard Euler equation, so the planner below chooses the activity gap x_t directly and the implementing nominal rate is recovered afterward in (25). Feasibility requires final output to cover consumption, materials, and the real resources spent on rebuilding—the $\Gamma h^3/3$ costs that also appear in the

planner's objective.

Labor services are the key input, and they require more than bodies: production runs on *organizational capital*—trained teams, routines, and supplier relationships embodied in each firm's assembled workforce. The workforce inherited from the previous year acts as *capacity*: a firm can produce up to that capacity at the going wage, but operating with part of its organization destroyed raises its marginal cost, and expanding employment beyond inherited capacity requires recruiting, screening, and training—activities whose cost rises with the pace of expansion and which raise the effective wage of the marginal hire. Wholesale firms maximize the present value of profits, revenue net of the wage bill, materials, and rebuilding costs; the pricing layer is linearized below, while the rebuilding problem is kept in its exact nonlinear form. The compact model does not carry a separate wage Phillips curve: the activity slope κ_x summarizes labor-cost pass-through at the annual frequency.

The compact model tracks this environment with two variables per producer type. A type's *shortfall* $q_{jt} \geq 0$ is the fraction of its pre-pandemic organizational capital still missing at the end of year t; its *rebuilding* $h_{jt} \geq 0$ is the fraction restored during year t by hiring and training above inherited capacity. The 2020 collapse sets the initial shortfall in proportion to each type's job exposure e_j, so $q_{j0} = \Delta e_j$ with $\sum_j \omega_j e_j = 1$, where Δ is the economy-wide mean destruction. A fraction ψ of the destroyed stock is a *warm pool*—laid-off workers on recall who return without retraining—so only the remainder must pass through costly rebuilding:

$$q_{j1} = (1 - \psi) q_{j0} - h_{j1}, \tag{17}$$

$$q_{jt} = q_{j,t-1} - h_{jt}, \qquad t \geq 2. \tag{18}$$

Rebuilding is an investment that firms choose, and its cost is the model's key friction. Restoring organizational capital at pace h_{jt} burns real resources—recruiting, training, and the disruption of incumbent teams—with total cost $\Gamma h_{jt}^3/3$ and hence marginal cost Γh_{jt}^2: rebuilding twice as fast costs four times as much at the margin, the convexity documented by the hiring-cost evidence in Section 3. The benefit of a restored unit is a year of operating with a smaller shortfall, with flow value $\xi(q_{jt} + b x_t)$: larger when the remaining shortfall q_{jt} is big, because scarce capacity carries a high shadow price, and when aggregate demand x_t is strong, with b governing the demand sensitivity. Each type chooses its rebuilding path to minimize the present value of rebuilding costs plus shortfall losses. The optimality condition compares rebuilding a

unit this year with postponing it one year: acting now costs Γh_{jt}^2, saves the discounted marginal cost $\beta\Gamma h_{j,t+1}^2$ of acting later, and buys one year of the flow benefit. Writing this net marginal cost as F_{jt}, optimal rebuilding satisfies the complementarity system

$$F_{jt} \equiv \Gamma h_{jt}^2 - \beta\Gamma h_{j,t+1}^2 - \xi(q_{jt} + bx_t) \geq 0, \tag{19}$$

$$h_{jt} \geq 0, \qquad h_{jt}F_{jt} = 0 : \tag{20}$$

either rebuilding is active and its net marginal cost is zero, or the benefit never covers the cost and the shortfall is left standing. Differentiating $F_{jt} = 0$ shows how demand raises the private return to reconstruction:

$$\frac{\partial h_{jt}}{\partial x_t} = \frac{\xi b}{2\Gamma h_{jt}} > 0. \tag{21}$$

Equation (21) is the assisting role of demand, and the convexity of the cost is what makes it useful for policy. The response is *self-limiting*: it is large when rebuilding is slow and fades as the rebuild accelerates, so the first units of demand buy the most repair. And it is cheapest in inflation terms just where it is most useful: because the marginal wedge Γh_{jt}^2 is convex, accommodating demand while rebuilding is still slow raises marginal cost only to second order, so the planner can speed the early recovery at little inflationary cost, paying the convex rebuilding premium only once the rebuild runs hot. A merely linear marginal cost—a quadratic total cost—would deliver neither property.

Two aggregates summarize the supply side:

$$Q_t = \sum_j \omega_j q_{jt}, \qquad \mathcal{W}_t = \sum_j \omega_j \Gamma h_{jt}^2. \tag{22}$$

Q_t is the economy's missing capacity—the *capacity wall*; $\mathcal{W}_t$ is the average rebuilding premium currently being paid—the wall's imprint on marginal cost. Because rebuilding today lowers the shortfall at every future date, an accommodative stance has an intertemporal supply effect: it raises inflation through demand today, but it removes capacity pressure—and therefore lowers inflation and raises output—tomorrow.

2.3 The conditional policy problem

The communication problem is solved conditional on the nonlinear capacity and rebuilding transition generated by the companion supply model. Let u_t denote the resulting supply pressure and let $\bar{x}_t = (1-\gamma)Q_t$ be the activity adjustment associated with missing capacity. Conditional on expected inflation e_t, the mean Phillips curve is

$$\theta_t = \beta e_t + \kappa x_t + u_t. \tag{23}$$

At date zero there is no inherited promise, so current policy is selected directly. Thereafter the current center is the one each mode delivers against the outstanding letter. The planner chooses the announcement path and prices the expected delivered losses and sanction in (13), rebuilding resources, and the authorization cost of the guidance wedge:

$$\mathcal{L} = \sum_{t=0}^{T} \beta^t \left[\left(1 - q_{t-1}\right) L_t^G + q_{t-1} L_t^R + \mathcal{K}_t^{\text{net}} + s_N \mathcal{R}_t + \frac{\varphi_g}{2} g_t^2 \right], \tag{24}$$

where $q_{-1} = 0$, the initial delivered centers coincide with discretion, $\mathcal{K}_t^{\text{net}}$ is the posterior-weighted net sanction of the due promise, $\mathcal{R}_t$ is the real resource cost of rebuilding, and s_N is the labor share of gross output that converts it into loss units. The constraint set includes (6)–(12) and the target-state terminal conditions.

The nominal-rate path implementing the conditional allocation is recovered from the IS curve,

$$i_t - r_t^n = e_t + \sigma \left(\widetilde{\mathbb{E}}_t x_{t+1} - x_t \right). \tag{25}$$

The rate is therefore an implementation object rather than an additional constraint in the conditional communication problem.

The signal-implied supply path used in (24) is certainty equivalent with respect to the delivery outcome. A fully joint solution would carry mode-specific organizational-capital distributions. The present formulation instead uses the calibrated nonlinear rebuilding path and reoptimizes the entire communication problem conditional on it. This is the principal quantitative approximation. It is preferable to ignoring delivered losses: the announcement is costly because the planner evaluates what delivering it would require.

3 Calibration

The supply path and rebuilding-resource series are taken from the calibrated annual transition in DT26b. The model starts in 2021 and ends in 2033. The inherited targeting rule uses $\phi_{IT} = 0.35$.[1] Message dispersions are $\sigma_G = 1.20$ and $\sigma_R = 3.00$ percentage points. The delivery block is the continuous graduated audit of DT26a: bounded formal sanction $\bar{F} = 4.50$ with hidden exposures $\chi_G = 1$ and $\chi_R = 0.50$, audit tolerance $\tau = 0.24$ and control noise $\sigma_\pi = 0.24$ percentage points, continuation-adjusted temptation curvature $\Psi = 11.6$, and reputational curvature $\bar{\varrho} = 0.15$. The long-run override-mode prior is $\bar{\lambda} = 0.05$, with persistence $\rho_\lambda = 0.871$.

The private persistence signal. The 2021 private signal is $v_0 = -0.28$. Under (4), it implies a one-year persistence of 0.344. The actual ratio of 2022 to 2021 supply pressure is 0.883. The difference is a new adverse innovation realized after the announcement. This calibration deliberately separates a truthful low-persistence assessment from ex-post forecast error.

Targets. The communication block—message dispersions, the graduated-audit delivery technology, the override-mode prior, and the reputational curvature—is shared in structure with the cost-push application of DT26a, which runs the same delivery technology on the same episode calibrated to December-over-December core CPI with the 2020–21 lower-bound calendar; here the supply path is the calibrated organizational-capital transition and the comparison is against headline CPI on the same December-over-December convention. No target requires the announcement itself to remain at two percent. The optimal announcement is allowed to be an elevated glide path because its future delivery cost is priced explicitly.

[1] The tight loading is consequential. DT26a shows, in the same environment, that a default target generous enough to carry the required accommodation itself drives the optimal guidance wedge to approximately zero, making communication quantitatively irrelevant. Guidance matters because the inherited regime is stricter than the policy the bank must deliver.

Table 1: Baseline delivery-consistent transition

Year	Supply	Forecast	Announcement	$\widetilde{E}_t \pi_{t+1}$	Inflation	Del. share
2021	4.48	4.48	3.01	3.12	7.40	0.687
2022	3.96	1.54	1.65	1.77	5.53	0.689
2023	2.58	0.58	1.26	1.37	3.78	0.693
2024	1.53	0.23	1.21	1.31	2.68	0.696
2025	0.74	0.09	1.29	1.37	1.96	0.699
2026	0.24	0.04	1.40	1.46	1.57	0.701

Notes: Supply is the calibrated realized pressure; Forecast is the path implied by the central bank's private signal at the outbreak. Announcement and expected inflation refer to the following year's center. Del. share is the prospective delivered share of the outstanding promise. Inflation is the realized annualized rate. All inflation entries are levels in percent.

4 Results

4.1 Baseline: a truthful glide path under a strong anchor

The branch-consistent baseline reproduces the inflation–expectations separation. Inflation peaks at 7.40 percent in 2021, while the one-year expectation is 3.12 percent. The bank announces a future center of 3.01 percent. The inherited default center is 2.54 percent, so the guidance wedge is 0.47 percentage point.

The announcement is therefore neither a mechanical two-percent forecast nor cheap talk. It is a realistic elevated glide path that promises a faster return than the inherited rule. The prospective delivered share in 2021 is 0.687. The planner obtains the current expectations benefit only by accepting that a corresponding part of the center will bind in 2022.

Figure 2 shows the timing. The private signal predicts rapid decay of supply pressure, but the realized 2022 state remains high. Planned inflation and expectations fall quickly; realized inflation remains elevated because of the new supply innovation. This is the coherent interpretation of the post-pandemic miss: the bank's conditional assessment proves too optimistic, rather than the bank knowingly promising a path it never intended to deliver.

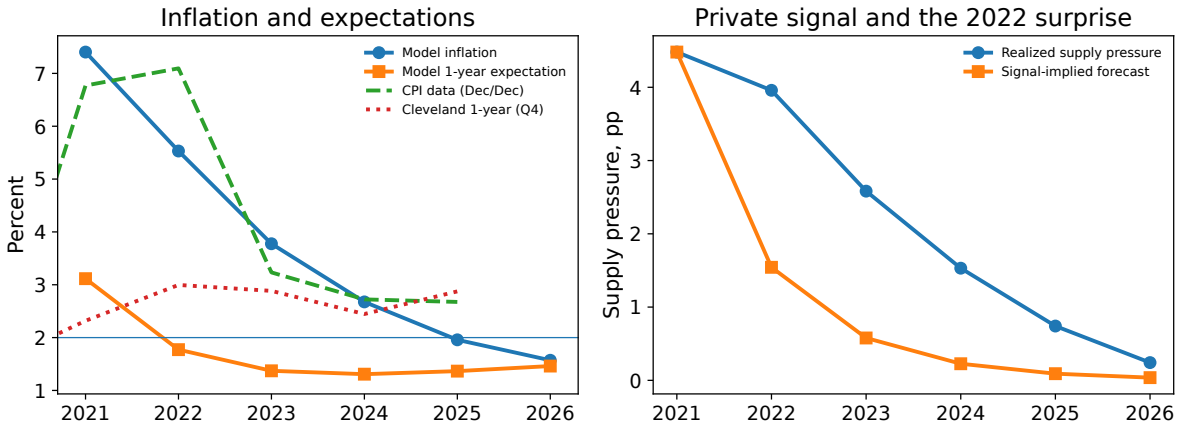

Figure 2: Inflation, expectations, and the private persistence signal

Notes: Panel A compares the model with December-over-December headline CPI inflation and the fourth-quarter Cleveland Fed one-year expectation, which forecasts headline CPI; the inflation–expectations comparison is therefore on a consistent headline basis. Panel B separates the signal-implied supply forecast from the realized calibrated path. The 2022 gap is an innovation that arrives after the 2021 announcement.

4.2 What the 2021 planner pays for the promise

The delivery accounting is the main substantive difference from the cost-push application of DT26a: here the 2021 planner explicitly evaluates what each mode will deliver in 2022. Under the signal-implied supply path, enforcing the letter of the 3.01-percent center requires an output gap of -2.07 percent. The weak mode delivers a center much closer to discretion and requires a much smaller adjustment. The corresponding losses are reported in Table 2.

Table 2: The letter and the delivered center priced in 2021

	Letter enforced	Delivered (weak mode)
2022 inflation center (percent)	3.01	3.17
2022 output gap (percent)	-2.07	-1.03
Delivery loss ($\times 10^4$)	1.93	1.21

Notes: The centers and output gaps are conditional on the signal-implied 2022 supply state. Losses include inflation and activity terms and are expressed in the sanction normalization of (7).

The same plan becomes much more expensive when the 2022 supply shock remains close to its 2021 level. Under that realized state, enforcing the letter would require a far deeper output contraction than under the baseline conditional forecast. This is the state in which the enforcement threshold matters. Delivery enters both expectations

and welfare from the start, so the wedge reflects a promise whose future delivery cost the planner has already priced.

4.3 Institutional fragility and the hidden low-resilience switch

Fragility has two ingredients. First, the audit sharpens: a tighter tolerance and lower control noise concentrate the delivery technology, raising delivered shares on small letter–desire gaps while bringing the enforcement threshold closer in gap units. Second, the public becomes less willing to accept the bank's low-persistence signal. We model the latter as partial weight on the subsequently realized persistence of supply pressure. The counterfactual also allows the actual, hidden mode to switch to low exposure in 2021; the public does not observe the switch directly.

The two ingredients play different roles. If the audit sharpens but the public continues to accept v_0, the letter–desire gap stays small, the announcement remains credible, and expected inflation is little changed. In the main fragility exercise, the public assigns modest weight to the possibility that the supply disruption will remain high. The gap at which the robust 2021 promise must be delivered then lies just across the enforcement threshold. The prospective delivered share falls to 0.013, expected inflation rises to 3.87 percent, and current inflation rises to 8.12 percent with the robust conventional policy held fixed.

The mechanism is straightforward. If the letter is enforced in the realized 2022 state, the output gap must fall to -18.4 percent. The delivered center of the low-exposure institution requires only -1.0 percent. The public therefore anticipates near-complete shading: past the enforcement threshold the bounded audit no longer holds delivery to the letter. The expectation effect of the promise disappears before delivery, and the current Phillips curve transmits the higher expected inflation into the outbreak year.

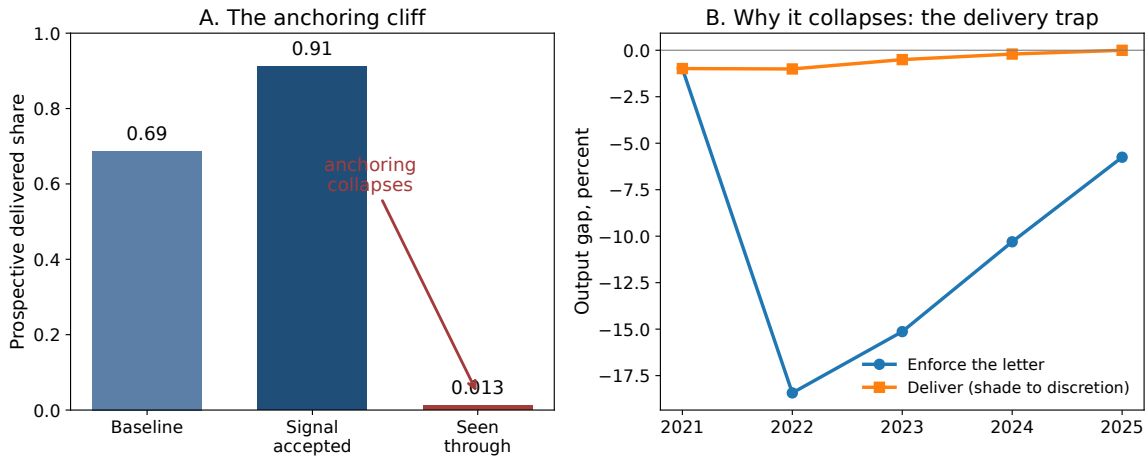

Figure 3: The fragility of a low-persistence promise: the anchoring cliff and the delivery trap

Notes: Panel A plots the prospective delivered share of the 2021 promise under three readings of the private signal: the baseline mix, a public that accepts the low-persistence signal, and a public that partly sees through it. Delivered pull collapses across the sharp enforcement threshold, from about 0.9 when the signal is accepted to 0.013 when it is seen through. Panel B shows why: enforcing the inherited letter after the 2022 surprise requires a severe output contraction, whereas the low-exposure institution's delivered center permits a much smaller adjustment.

Table 3: The same robust policy under alternative interpretations of persistence

Environment	Del. share	Expected inflation	Inflation	x_{2022}, letter	x_{2022}, delivered
Robust optimum	0.687	3.12	7.40	-2.1	-1.0
Sharp audit; public accepts v_0	0.911	3.04	7.33	-17.6	-1.0
Partial see-through; hidden switch	0.0128	3.87	8.12	-18.4	-1.0
Full see-through (polar)	0.0128	9.24	13.28	-44.2	-2.7

Notes: The robust row is optimized under the baseline audit and the private low-persistence signal. The second row sharpens the audit but leaves that signal accepted by the public. The third row is the main fragility counterfactual: the public partly anticipates the realized persistence and the actual hidden mode switches to low exposure. The final row uses the realized persistence fully and is a polar sensitivity, not a historical calibration.

In the weak-anchor exercise the planner raises its glide path endogenously when credibility is poor—peak inflation is 7.64 percent and expected inflation 3.37 percent—because delivery consistency prevents an arbitrarily aggressive promise from substituting for a weak institution. Our fragility rows vary only what the public infers from words, so even the polar full-see-through row (13.28 and 9.24 percent) describes a public that already knows the realized state. Through a different channel, it reaches the darkest reading in DT26a, in which announcements are read as volatile policy preferences rather than information (near 11.8 and 9.8 percent); the companion theory ranks that reading as the more adverse, ordering interpretations rather than peak inflation.

5 The Market Footprint of Traction

The theory distinguishes the average weight on an outstanding promise from the marginal effect of making the message more aggressive. A simple market diagnostic is whether scheduled Federal Reserve communication produces excess volatility in short-horizon inflation compensation. For series z and era E, define

$$R_z(E) = \frac{\mathrm{sd}(\Delta z_t \mid t \text{ is an FOMC day in } E)}{\mathrm{sd}(\Delta z_t \mid t \text{ is an ordinary trading day in } E)}. \tag{26}$$

A ratio of one means that meeting days contain no excess news relative to ordinary trading days. During the inflation surge, the nominal-yield ratio remained elevated while the inflation-compensation ratio fell sharply. This is consistent with a policy that retained average anchoring power but had little remaining marginal effect on inflation expectations.

Panel B of Figure 4 plots the model's counterpart to this ratio against the data, era by era. The mapping is explicit. An unanticipated meeting-day message surprise moves expected inflation by the marginal traction $M_t \equiv \partial \widetilde{E}_t \pi_{t+1} / \partial a_t$; let M_E be its era average along the solved path and V_E the variance of those surprises. Because the surprise moves expectations only on meeting days, it adds $M_E^2 V_E$ to the meeting-day variance of inflation compensation while leaving the ordinary-day variance $\sigma_0^2(E)$ unchanged. Hence

$$R^2(E) = \frac{\sigma_0^2(E) + M_E^2 V_E}{\sigma_0^2(E)} \qquad \Longrightarrow \qquad R^2(E) - 1 = \frac{M_E^2 V_E}{\sigma_0^2(E)}. \tag{27}$$

Normalizing by the anchored pre-COVID era $E_0 = 2012$–19 and letting $\kappa \equiv R^2(E_0) - 1$ absorb the common ordinary-day scale gives the one-constant prediction

$$R^2(E) - 1 = \kappa \left(\frac{M_E}{M_0} \right)^2 \frac{V_E}{V_0}, \tag{28}$$

where M_0 is the anchored-state traction and κ is pinned so that the model reproduces the 2012–19 ratio *by construction*. This makes each plotted model bar transparent:

- **Anchored era (2012–19):** $M_E/M_0 = 1$ and $V_E/V_0 = 1$, so $R = \sqrt{1+\kappa}$ equals the data ratio. This bar is the normalization, not a prediction.
- **Surge and post eras (2021–23, 2023–25):** marginal traction is nearly flat across

eras, so it cannot move the ratio. What moves it is that the optimal announced center is nearly constant once the shock hits—the bank issues no fresh meeting-day guidance—so $V_E \to 0$ and the predicted ratio collapses to one.

The data show the same collapse: meeting-day inflation-compensation news, elevated when guidance was active, disappears during and after the surge.

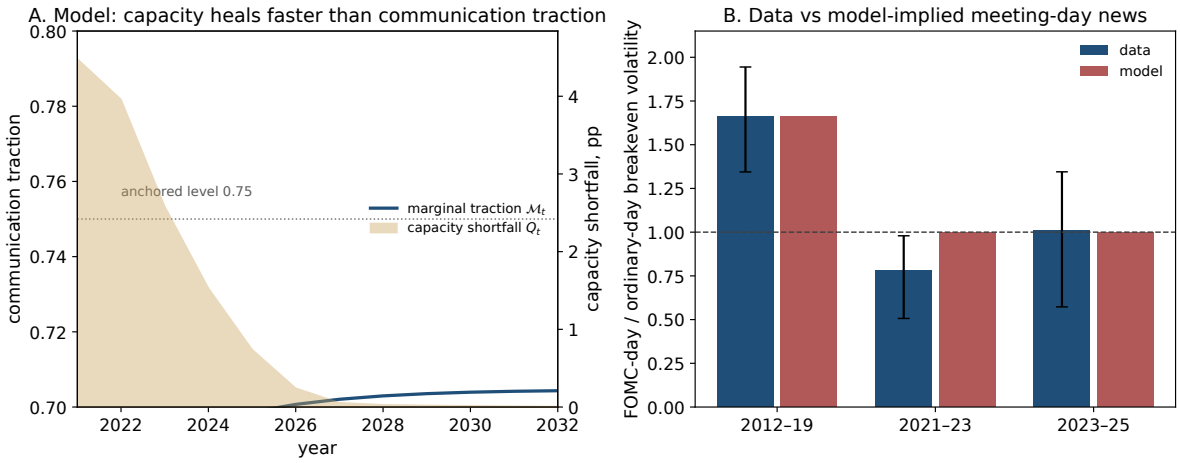

Figure 4: Model traction and the meeting-day diagnostic

Notes: Panel A reports the marginal message traction M_t along the delivery-consistent baseline against the capacity shortfall Q_t. Panel B compares the data volatility ratio in each era (with bootstrap intervals) against the model prediction of (28): the pre-COVID era is matched by construction, and the surge bars fall to one because the optimal announced center is nearly constant, leaving no meeting-day message innovation. The market exercise is descriptive and does not identify the delivery technology.

6 Conclusion

A central bank can use forward guidance to separate current expected inflation from realized inflation during a destructive supply recovery, but only by creating an obligation for the future. Once that obligation is priced correctly, the optimal announcement is an elevated glide path rather than an invariant two-percent promise. The motive to sustain demand is stronger here than under an ordinary cost-push shock: both are partly looked through, but demand now rebuilds the capacity whose absence is causing the inflation. The calibrated branch-consistent application still generates a large post-COVID inflation–expectations wedge: inflation reaches 7.40 percent while one-year expectations remain at 3.12 percent.

The private signal clarifies how this can happen without deception. In 2021 the bank receives information pointing to rapid fading of supply pressure and truthfully

announces a relatively low future center. A new adverse innovation then keeps supply pressure high. The promise becomes costly ex post because enforcing its letter would require a severe contraction. Shading delivery toward discretion is the institution's escape valve—priced by the graduated public audit, not a costless choice that the planner ignores.

This mechanism also makes the policy fragile. If the public accepts the low-persistence signal, even a sharp audit need not destabilize expectations. If the public instead partly anticipates the subsequent persistence and suspects a low-exposure mode, the same policy can lose almost all of its delivered pull. Expected and realized inflation then rise sharply. The relevant policy risk is therefore disagreement about the state in which a future promise will have to be delivered.

The exercise remains illustrative. It conditions communication policy on the capacity-rebuilding path and does not solve the full mode tree for organizational capital. But it satisfies the core rational-expectations restriction: every expectation benefit of guidance is paired with the posterior-weighted cost of what each mode will actually deliver, including the institutional cost of shading. That restriction is indispensable for using the companion communication theory in the post-COVID application.

References

Basu, Susanto, "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 1995, *85* (3), 512–531.

Drozd, Lukasz A., "Built by Actions, Spent by Words: Central Bank Communication under Inflation Targeting," July 2026. Working paper, Federal Reserve Bank of Philadelphia.

_ **and Marina M. Tavares**, "Capacity Wall: Inflation After COVID," July 2026. Working paper, Federal Reserve Bank of Philadelphia.