Reserve Demand under Mandatory Treasury Clearing: A Two-Channel Decomposition of the Operational Threshold

[Author Name]
[Affiliation]

April 2026

Abstract

The Federal Reserve's operating framework requires a rule for switching from passive (Standing Repo Facility backstop) to active (reserve management purchases) tooling as reserves decline. Mandatory central clearing of Treasury repo and the recalibration of the supplementary leverage ratio shift the reserve level at which that switch should occur. We decompose this shift into two analytically distinct channels. Channel A is bank-side and level: the kink in reserve demand, set by a value-at-risk rule for intraday liquidity, moves rightward by $\Delta R^k = z_\alpha \, \Delta\sigma_\eta$, distribution-free in the underlying outflow shock. Calibration to April 2026 conditions delivers $\Delta R^k \approx \$21\text{B}$ for a one-standard-deviation reform-induced increase in intraday outflow volatility. Channel B is dealer-side and slope: the local elasticity of SOFR volatility to reserve drains is amplified by a factor $\rho = (1 + c\gamma)/(c\gamma) \approx 3$ when the dealer leverage constraint is near binding, and is silent when dealer headroom is large. The two channels are conceptually distinct and arise in different empirical regimes: Channel A operates uniformly across days; Channel B activates at quarter-ends, on tax days, and during stress episodes such as September 2019. A planner microfoundation pins down when the bank-side level shift coincides with the planner's optimal switching threshold shift, namely on a positive-measure cone in the planner's preference parameters that contains the empirically realistic Federal Reserve. The decomposition yields a concrete operational rule: monitor the intraday SOFR volatility percentile, expect the kink to shift by $z_\alpha \Delta\sigma_\eta$, and trigger active tooling at $R \leq R^k_{\text{post}}$.

1 Introduction

The Federal Reserve currently runs an ample-reserves regime in which the policy rate is administered by interest on reserve balances and the corridor is bracketed by overnight reverse repurchases below and the Standing Repo Facility (SRF) above. As reserves fall along the post-2022 runoff, two questions become operational. At what reserve level does the Fed switch from passive backstop tooling (the SRF used only as a release valve) to active reserve management (open-market purchases that re-supply reserves and prevent SOFR from leaving the corridor)? And how does the answer change as two regulatory reforms phase in: the Securities and Exchange Commission's mandate that interdealer Treasury repo move into central clearing through the Fixed Income Clearing Corporation (FICC), with a deadline of June 30, 2027, and the federal banking agencies' recalibration of the supplementary leverage ratio (SLR), which excludes Treasuries from the eSLR exposure measure effective April 1, 2026?

This paper derives a closed-form decomposition of the reform-induced shift in the operational reserve threshold. The decomposition has two analytically distinct channels operating on different parts of the reserve-demand curve.

Channel A (bank-side, level). Mandatory clearing fragments end-of-day settlement timing across CCP cycles, initial-margin posting, and variation-margin draws. Under a value-at-risk (VaR) rule for intraday liquidity, banks size reserve buffers proportional to the standard deviation σ_η of net intraday outflows. Clearing raises σ_η, so the kink $R^k = \sigma_\eta z_\alpha$ in reserve demand — the level at which precautionary scarcity begins to bind, shifts right by

$$\Delta R^k \;=\; z_\alpha \, \Delta\sigma_\eta.$$

The shift is distribution-free in the family of intraday outflow shocks: it requires only a scale family and the VaR rule, not normality, not a particular calibration of the dealer side, and not the SLR weight. At the calibration anchored to April 2026 reserves and a one-standard-deviation reform-induced volatility increase of $\Delta\sigma_\eta =$ \$9B, the implied shift is $\Delta R^k \approx$ \$21B — the same order of magnitude as a single Federal Open Market Committee reserve adjustment. *This is the operational level shift.*

Channel B (dealer-side, slope). Off the kink, the equilibrium SOFR responds to bank shortfall borrowing through dealer intermediation. When the dealer SLR constraint is far from binding (slack-dealer regime), incremental bank repo demand is met along the dealer's smooth marginal-cost schedule, with slope $a_u = c/(1 + c\gamma)$. When the dealer constraint is at the bank's VaR threshold (kink-aligned regime), incremental demand is met by end-user substitution, with slope $a_c = 1/\gamma$. The ratio

$$\rho \;=\; \frac{a_c}{a_u} \;=\; \frac{1 + c\gamma}{c\gamma}$$

is the regime-shift amplification of the local elasticity of SOFR volatility to reserve drains. At the calibration $c\gamma = 0.5$, $\rho = 3$. *Channel B is regime-dependent.* It activates at quarter-ends, on tax days, around large Treasury settlement dates, and during the September 2019-style configurations in which dealer headroom is exhausted. It is silent in the empirically modal slack-dealer state.

Microfoundation: when does the bank-side level shift equal the operational threshold shift? The Federal Reserve does not directly observe R^k. It observes SOFR volatility and chooses the reserve target that minimizes a planner loss balancing balance-sheet cost ($\lambda_R R$) against money-market disruption ($\lambda_\sigma \sigma^2_{\text{SOFR}}$). We show (Corollary 4) that the planner's optimal switching threshold $R^\star$ coincides with the bank's VaR kink R^k on a positive-measure cone of preference parameters, namely $\lambda_R/\lambda_\sigma > \max(K_{\text{pre}}, K_{\text{post}})$, where K_θ is a closed-form regime-specific constant. Calibration places $\max(K_{\text{pre}}, K_{\text{post}}) \approx 0.027$ while the Fed's revealed preference (inferred from the December 2025 Reserve Management Purchases activation threshold of approximately \$3.0T) sits at $\lambda_R/\lambda_\sigma \approx 1000$ — four orders of magnitude inside the cone. The bank-side Channel A shift therefore *is* the operational threshold shift the Fed should target.

Contribution and positioning. Three bodies of work bear on the question. The reserve-demand literature [Lopez-Salido and Vissing-Jorgensen, 2023, Afonso et al., 2022, 2024, Haubrich, 2023] estimates the kink and its time variation but treats it as exogenous to clearing structure and the SLR state. The dealer-balance-sheet literature [Du et al., 2023b,a, Duffie et al., 2023] prices SLR as a binding intermediary constraint but does not link it to aggregate reserve supply or to a Fed operational rule. The clearing literature [Copeland and Kahn, 2024, Hempel et al., 2024, Chang et al., 2025, Liang and Zhu, 2025] models OTC-versus-CCP wiring and accounting netting benefits but contains no aggregate reserve variable and no SOFR comparative static. No existing paper jointly delivers $R^k(\theta, Q)$ and the local SOFR-volatility slope as closed-form functions of the regulatory state. The closest predecessor, Lopez-Salido and Vissing-Jorgensen [2023], is the special case $Q = 0$ with β fixed.

Roadmap. Section 2 positions the paper against three adjacent literatures. Section 3 sets up the two-period environment with three agents (a VaR-rule bank, an SLR-constrained dealer, and a price-elastic end-user). Section 4 establishes existence, uniqueness, and the closed-form SOFR map. Section 5 states the headline Channel A result. Section 6 characterizes Channel B and proves the regime-shift amplification. Section 7 ties the two channels into the operational-threshold decomposition and presents the planner microfoundation. Section 8 reports the numerical robustness exercises. Section 9 states the testable implications and the operational rule. Section 11 concludes. Proofs and the off-locus implicit-function formula are in the appendix.

2 Related literature

Three positioning moves locate the contribution.

Reserve demand and the operational tipping point. Lopez-Salido and Vissing-Jorgensen [2023] estimates a structural reserve-demand function with an IORB spread, a convenience yield, and a balance-sheet cost, and uses it to identify the QT level beyond which rate volatility resurges. Afonso et al. [2022, 2024] estimate a time-varying reserve-demand elasticity in real time. Haubrich [2023] derives an inventory-theoretic optimal buffer of about \$60B but holds structural parameters fixed. None of these papers carries clearing intensity or the SLR weight as a state variable. The present paper nests Lopez-Salido and Vissing-Jorgensen [2023] as the special case in which clearing intensity $Q = 0$ and the SLR weight β does not vary across regimes (Section 8). The contribution is to make the reserve-demand kink and the local SOFR-volatility slope explicit functions of the regulatory state.

Dealer balance sheets and intermediary constraints. He and Krishnamurthy [2013], Adrian and Shin [2010], Du et al. [2023a], and Du et al. [2023b] show that intermediary leverage is empirically a binding pricing constraint across asset classes. Duffie et al. [2023] document the four-fold decline in dealer capacity per dollar of Treasuries. Carlson et al. [2025] show heterogeneity in dealer SLR slack at the bank-holding-company level. We use this literature for the price-taking dealer with an SLR balance-sheet constraint and a quadratic intermediation cost. We do not model dealer market power; Chang et al. [2025] study the imperfect-competition channel separately. Assumption A3 below imposes price-taking to isolate the regulatory channel; in a richer model with both, the two channels would interact, and the present paper is silent on that interaction.

Clearing, sponsored repo, and netting. Copeland and Kahn [2024] document that dealers shift volume to FICC sponsored clearing when balance sheets are tight. Liang and Zhu [2025] estimate that joint clearing and SLR recalibration release on the order of \$1.3T of dealer balance-sheet capacity. Hempel et al. [2024] and Bowman et al. [2024] caution that netting-based SLR relief is more limited than headline accounting suggests. The OFR [Office of Financial Research, 2025] estimates that mandatory clearing modestly raises SOFR tail volatility relative to the 2022 pilot baseline. The present paper takes the cash-drain channel (mQ) and the intraday-volatility channel ($\sigma'_\eta(Q) > 0$) as the two operative effects of clearing. We do not model netting-induced SLR relief as a third clearing channel; were $\sigma'_\eta(Q) \leq 0$ to obtain empirically (the netting case), Channel A would reverse sign and the rest of the model would be unchanged. Section 8 reports a sensitivity exercise on this assumption.

September 2019 and the kink-aligned regime. Anbil et al. [2020], Afonso et al. [2019], and Copeland et al. [2025] converge on a diagnosis of the September 2019 SOFR spike as the joint product of intraday timing concentration, dealer balance-sheet tightness, and a corporate-tax outflow. The model's kink-aligned regime ($z_c \downarrow z_\alpha$) is the abstract structural counterpart of that configuration. We use September 2019 as an illustrative reference point in the kink-aligned regime; we do not claim to rederive the spike's quantitative magnitude (the calibrated $\rho \approx 3$ is too small to produce the observed 500-basis-point excursion, which would require $\rho \approx 50$ and additional concentration mechanisms outside the present model). The mapping is structural, not quantitative.

Operational framework analogs. The Bank of England's transition to a repo-led operating framework and the ECB's corridor-versus-floor debate confirm the universality of the active-versus-passive switching question. Duygan-Bump and Kahn [2026] frames the three-way trade-off (balance-sheet size, rate volatility, market intervention) as a trilemma. Kashyap et al. [2025] make the case for structured central-bank intervention in Treasury markets. The present paper provides a closed-form rule for the active-passive switch that the trilemma framing leaves at the conceptual level.

3 Model

3.1 Environment

A two-period economy with three agents and three regulatory state variables. Periods are $t \in \{0, 1\}$. At $t = 0$ the bank chooses its reserve holding given the conditional law of the intraday outflow. At $t = 1$ the shock realizes, the bank borrows any shortfall in repo, and markets clear.

3.2 Three agents

Bank (B). Holds r reserves at the interest rate on reserve balances i_R. Funds reserves at the outside money-market rate i_M. Faces a net intraday outflow $\eta \sim \mathcal{N}(0, \sigma_\eta(Q)^2)$ at $t = 1$. If $\eta > r$, borrows the shortfall in repo at the equilibrium SOFR s. The bank's decision rule is the value-at-risk floor $\Pr(\eta > r) \leq \alpha$, which under normality binds at

$$r \;=\; R^k(\theta, Q) \;\equiv\; \sigma_\eta(Q)\, z_\alpha,$$

where $z_\alpha \equiv \Phi^{-1}(1 - \alpha)$. We adopt the VaR rule as the bank's primitive: it is the operational implementation of liquidity coverage ratio (LCR) regulation under 12 CFR 249, and it is the closed-form solution to $\max_r (i_R - i_M) r$ subject to $\Pr(\eta > r) \leq \alpha$ when $\tau_0 \equiv i_R - i_M \leq 0$. A smooth-penalty FOC variant is reported in Appendix C; the comparative statics are qualitatively identical.

Dealer (D). Lends $L_D \geq 0$ at SOFR s, funds at i_M. Pays a quadratic intermediation cost $\frac{c}{2} L_D^2$ and faces the SLR balance-sheet constraint

$$\beta(\theta) L_D \;\leq\; \tilde{k}(\theta, Q) \;\equiv\; k - mQ,$$

where $\tilde{k}$ reflects the cash drain from CCP margin posting at rate m. The dealer is a price-taker.

End-user composite (E). Contributes a price-elastic component $-\gamma(s - i_R)$ to repo demand, capturing hedge-fund and leveraged-investor substitution. The collapse to a representative-E formulation with constant γ is standard.

3.3 Regulatory state

The exogenous regulatory state $(\theta, Q, m, \beta(\theta), k)$ is known at $t = 0$. The reform comprises three movements: $\beta_{\text{post}} < \beta_{\text{pre}}$ (Treasury exclusion from SLR), $Q_{\text{post}} > Q_{\text{pre}}$ (FICC mandate phase-in), and a declining aggregate reserve supply R along QT. Notation is summarized in Table 1.

3.4 Standing assumptions

A1. Normal scale family. $\eta \sim \mathcal{N}(0, \sigma_\eta(Q)^2)$ with $\sigma_\eta : [0, 1] \to (0, \infty)$ continuously differentiable and $\sigma_\eta'(Q) > 0$. The sign of $\sigma_\eta'(Q)$ is a maintained reduced-form input motivated by settlement-

Table 1: Notation.

Symbol	Meaning
i_R, i_M, s	IORB, outside money rate, equilibrium SOFR
$\tau \equiv s - i_R$, $\tau_0 \equiv i_R - i_M$	SOFR-IORB spread, IORB-floor spread
$R \geq 0$	Aggregate reserve supply (Fed-controlled)
$\eta \sim \mathcal{N}(0, \sigma_\eta(Q)^2)$	Net intraday outflow
$Q \in [0, 1]$	Mandatory clearing intensity
$m > 0$	Per-unit cash margin posted to CCP
$\beta(\theta)$	SLR weight on dealer Treasury repo
$k > 0$	Dealer balance-sheet limit
$\tilde{k} = k - mQ$, $\Lambda = \tilde{k}/\beta$	Effective headroom, free repo capacity
c, γ	Dealer cost coefficient, end-user elasticity
$\alpha, z_\alpha = \Phi^{-1}(1 - \alpha)$	Bank tail tolerance, standardized quantile
$R^k(\theta, Q) = \sigma_\eta(Q) z_\alpha$	Bank VaR-rule reserve target (the kink)
$H(\theta, Q)$, $\eta^c = R + H$	Dealer headroom, dealer-binding cutoff
$z_R = R/\sigma_\eta$, $z_c = \eta^c/\sigma_\eta$	Standardized thresholds
$\sigma_{\text{SOFR}}(R, \theta, Q)$	Cross-η standard deviation of $s^\star$
$\bar{\sigma}$	Fed-tolerated SOFR volatility
$R^\star(\theta, Q, \bar{\sigma})$	Operational threshold (passive-active switch)
$\rho = (1 + c\gamma)/(c\gamma)$	Regime-shift slope ratio

timing fragmentation under mandatory clearing. The opposite case (clearing reduces σ_η via netting; Bowman et al. 2024) reverses Channel A; Section 8 reports the sensitivity of ΔR^k to $\sigma'_\eta(Q)$ over the range $[0, 2 \times \text{baseline}]$.

A2. Linear cash drain. $\tilde{k}(\theta, Q) = k - mQ$ with $k > mQ$ on the relevant range.

A3. Dealer price-taking. Imposed for tractability, isolating the regulatory-constraint channel.

A4. Reserve representativeness. $r = R$ in equilibrium.

A5. Sign of carry. $\tau_0 = i_R - i_M \leq 0$.

A6. Interior dealer headroom. $H(\theta, Q) > 0$ in both regimes, equivalently $(1 + c\gamma)\Lambda > \gamma|\tau_0|$. This bounds the size of the joint reform; under the calibration of Section 8, the binding k is approximately 0.014 versus a calibrated $k = 75$, so A6 holds with four orders of magnitude of slack.

A7. LCR tail tolerance. $z_\alpha \geq 0.2$. The empirical LCR range $\alpha \in (0.005, 0.1)$ gives $z_\alpha \in [1.28, 2.58]$.

3.5 Equilibrium

Definition 1 (Equilibrium)**.** Given (R, θ, Q), an equilibrium is a pair $(r^\star, s^\star(\cdot))$ such that

1. $r^\star = R^k(\theta, Q)$ (bank VaR rule);

2. for each realization η, $s^\star(\eta; R, \theta, Q)$ clears the dealer supply against the bank shortfall plus end-user demand;
3. representativeness: $r^\star = R$.

The bank's optimal $r^\star = R^k$ and the Fed-controlled aggregate R are reconciled only on the locus $R = R^k(\theta, Q)$. Off-equilibrium $R \neq R^k$ corresponds to the cross-section of Fed policy choices we analyze.

4 Equilibrium SOFR and basic results

4.1 The SOFR map

The dealer's first-order condition delivers an unconstrained supply $L_D^u(s) = (s - i_M)/c$ on $s \leq \bar{s}(\theta, Q) \equiv i_M + c\tilde{k}/\beta$ and a constrained supply $L_D^c = \tilde{k}/\beta$ above. The aggregate residual repo demand at $t = 1$ is $D_B(s; R, Q, \eta) = (\eta - R)^+ - \gamma(s - i_R)$. Equating supply and demand in each regime delivers the SOFR map.

Lemma 1 (Existence and uniqueness of the SOFR map)**.** *For each (R, θ, Q, η) there exists a unique $s^\star(\eta; R, \theta, Q) \in [i_M, \infty)$ given by*

$$s^\star(\eta; R, \theta, Q) \;=\; s_0 + \begin{cases} 0, & \eta \leq R, \\ a_u(\eta - R), & R < \eta \leq \eta^c, \\ a_u(\eta^c - R) + a_c(\eta - \eta^c), & \eta > \eta^c, \end{cases}$$

where $s_0 = (i_M + c\gamma i_R)/(1 + c\gamma)$, $a_u = c/(1 + c\gamma)$, $a_c = 1/\gamma$, and

$$\eta^c(R, \theta, Q) \;=\; R + (1 + c\gamma)\,\frac{\tilde{k}(\theta, Q)}{\beta(\theta)} - \gamma\tau_0 \;\equiv\; R + H(\theta, Q).$$

$s^\star$ is continuous, weakly increasing in η, with slope a_u on (R, η^c) jumping to a_c on (η^c, ∞). The slope ratio is $\rho = a_c/a_u = (1 + c\gamma)/(c\gamma) > 1$.

Proof. Setting $L_D^u(s) = (\eta - R)^+ - \gamma(s - i_R)$ in the unconstrained regime gives the first arm. Setting $L_D = \tilde{k}/\beta$ in the constrained regime gives the second. Continuity at $\eta = \eta^c$ pins H. $\square$

A useful re-expression: writing $b \equiv a_c - a_u = 1/[\gamma(1 + c\gamma)]$,

$$s^\star(\eta) - s_0 \;=\; a_u(\eta - R)^+ + b(\eta - \eta^c)^+. \tag{1}$$

4.2 Closed-form variance of SOFR

Let $Y_t \equiv (\eta - t)^+$ for $t \in \{R, \eta^c\}$. Standard truncated-normal moments deliver

$$\mathbb{E}[Y_t] = \sigma_\eta M(z_t), \qquad \operatorname{Var}(Y_t) = \sigma_\eta^2 V(z_t),$$

with $z_t = t/\sigma_\eta$,

$$M(z) = \phi(z) - z[1-\Phi(z)], \qquad V(z) = J(z) - M(z)^2, \qquad J(z) = (1+z^2)[1-\Phi(z)] - z\phi(z),$$

and cross-moment

$$\mathrm{Cov}(Y_{t_1}, Y_{t_2}) = \sigma_\eta^2 C(z_1, z_2), \qquad C(z_1,z_2) \equiv (1+z_1 z_2)[1-\Phi(z_2)] - z_1\phi(z_2) - M(z_1)M(z_2).$$

Lemma 2 (Closed-form SOFR variance)**.** *Under A1–A6, the cross-η variance of $s^\star$ is*

$$\sigma^2_{\mathrm{SOFR}}(R,\theta,Q) \;=\; \sigma_\eta(Q)^2 \cdot \Psi(z_R, z_c; c, \gamma), \tag{2}$$

where $z_R = R/\sigma_\eta$, $z_c = (R + H(\theta,Q))/\sigma_\eta$, and

$$\Psi(z_R, z_c; c, \gamma) \;\equiv\; a_u^2 V(z_R) + b^2 V(z_c) + 2a_u b\, C(z_R, z_c).$$

Proof. From (1), $s^\star - s_0 = a_u Y_R + bY_{\eta^c}$. Variance is bilinear; the closed forms above deliver (2). ☐

4.3 Monotonicity of SOFR volatility in reserves

Lemma 3 (σ_{SOFR} strictly decreasing on $[R^k, \infty)$)**.** *Under A1–A6, $\sigma^2_{\mathrm{SOFR}}(R,\theta,Q)$ is strictly decreasing in R on $R \geq R^k(\theta,Q)$.*

Proof. Differentiating (2) in R, both z_R and z_c rise by $1/\sigma_\eta$, so $d\sigma^2_{\mathrm{SOFR}}/dR = \sigma_\eta[\Psi_1 + \Psi_2]$. We show $\Psi_1 + \Psi_2 < 0$ on $\{(z_R, z_c) : z_R \leq z_c\}$, which contains the operational range $R \geq R^k$. Two pieces. (i) $\Psi_1 < 0$: using $V'(z) = -2\Phi(z)M(z)$ and $C_1(z_1,z_2) = -\Phi(z_1)M(z_2)$,

$$\Psi_1 \;=\; -2\Phi(z_R)\big[a_u^2 M(z_R) + a_u b\, M(z_c)\big] \;<\; 0.$$

(ii) $\Psi_2 < 0$: closed form (Lemma 5 below) gives $\Psi_2 = -2b[a_u p\, g + a_c(1-p)m]$ where $p = 1-\Phi(z_c)$, $m = M(z_c)$, and $g = (z_c - z_R) - M(z_R) + M(z_c) \geq 0$ for $z_c \geq z_R$. Both summands are non-negative and the second is strictly positive for any finite z_c. ☐

Corollary 1 (Existence and uniqueness of $R^\star$)**.** *Fix (θ, Q). For any $\bar\sigma \in (0, \sigma_{\mathrm{SOFR}}(R^k,\theta,Q)]$, the operational threshold*

$$R^\star(\theta,Q,\bar\sigma) \;\equiv\; \min\{R \geq R^k(\theta,Q) : \sigma_{\mathrm{SOFR}}(R,\theta,Q) \leq \bar\sigma\}$$

exists, is unique, and satisfies the point condition $\sigma_{\mathrm{SOFR}}(R^\star,\theta,Q) = \bar\sigma$. For $\bar\sigma > \sigma_{\mathrm{SOFR}}(R^k)$, the threshold is $R^\star = R^k$.

Proof. By Lemma 3, $\sigma_{\mathrm{SOFR}}(\cdot,\theta,Q)$ is strictly decreasing on $[R^k,\infty)$. Since $V \to 0$ and $C \to 0$ as $z_R, z_c \to \infty$, $\lim_{R\to\infty} \sigma_{\mathrm{SOFR}} = 0$. The intermediate value theorem delivers the threshold. ☐

Lemma 3 and Corollary 1 establish that the operational threshold $R^\star$ is well-defined as an object before any planner microfoundation. The next two sections analyze how R^k and the local slope $G_\theta = |\partial\sigma_{\text{SOFR}}/\partial R|_{R^k}$ each move under the joint reform.

5 Channel A: bank-side level shift

This section states the headline result. The reserve-demand kink, set entirely by the bank's VaR rule, shifts rightward by a quantity that depends only on the change in intraday outflow volatility and the bank's tail-tolerance quantile. The dealer side, the SLR weight, and the cash-drain coefficient do not enter.

Proposition 1 (Kink shift, distribution-free)**.** *Let $R^k(\theta, Q)$ be the VaR-rule reserve target at tail tolerance $\alpha \in (0, 1/2)$ for any baseline distribution $F(\cdot; Q) = F_0(\cdot/\sigma_\eta(Q))$ in the scale family with strictly positive median-to-quantile gap. Then R^k is strictly increasing in σ_η, and*

$$R^k(\theta, Q_{post}) \;>\; R^k(\theta, Q_{pre}) \quad \textit{for all } \theta, \alpha.$$

The result holds independently of $\beta(\theta)$.

Proof. $R^k = \sigma_\eta(Q)F_0^{-1}(1-\alpha)$, so $\partial R^k/\partial Q = \sigma'_\eta(Q)F_0^{-1}(1-\alpha) > 0$ by A1. □

Corollary 2 (Channel A magnitude)**.**

$$\Delta R^k \;=\; z_\alpha\,[\sigma_\eta(Q_{post}) - \sigma_\eta(Q_{pre})] \;>\; 0.$$

Economic content. Mandatory clearing fragments end-of-day settlement timing across CCP cycles, initial-margin posting windows, and variation-margin draws. The bank's tail-control rule for intraday liquidity requires reserves to scale linearly with the standard deviation of net outflows, so the demand-side kink shifts right by exactly $\sigma'_\eta(Q)z_\alpha \cdot \Delta Q$. The supplementary leverage ratio reform does not enter this margin: β governs dealer balance-sheet capacity, which prices the SOFR *above* the kink but does not change the bank's own VaR floor.

Calibration magnitude. At $\alpha = 1\%$ ($z_\alpha = 2.326$) and a reform-induced volatility increase $\Delta\sigma_\eta = \$9\text{B}$ (from $\sigma_\eta = \$26.5\text{B}$ at $Q_{\text{pre}} = 0.10$ to $\sigma_\eta = \$35.5\text{B}$ at $Q_{\text{post}} = 0.70$, calibrated to FRBNY Liberty Street intraday SOFR percentiles and the OFR 2025 estimate of post-mandate clearing share),

$$\Delta R^k \;=\; 2.326 \times 9 \;\approx\; \$20.94\text{B}.$$

Compared to the FRBNY Reserve Management Purchases threshold of approximately \$3.0T (December 2025 activation; Federal Reserve Bank of New York 2026), the shift is 0.70% of the threshold. In absolute terms, it is the same order of magnitude as a single FOMC reserve adjustment (\$25–100B). Sensitivity to $\sigma'_\eta(Q)$ ranges from $\Delta R^k \approx \$7\text{B}$ (at $\sigma'_\eta(Q) = 5$) to \$56B (at $\sigma'_\eta(Q) = 40$), all approximately linear in $\sigma'_\eta(Q)$.

What this delivers. Proposition 1 is what the Federal Reserve should monitor as the leading signal for the operational threshold shift. Three features are worth emphasizing. First, the result is *distribution-free*: it requires only that the intraday outflow shock lie in a scale family with $\sigma'_\eta(Q) > 0$. Normality is not used. Second, the result is *regime-independent*: it does not require knowing whether the dealer is at or near its SLR constraint. Third, the result is *closed-form*: $\Delta R^k = z_\alpha \, \Delta\sigma_\eta$ contains exactly two empirically identifiable inputs — the bank's tail tolerance and the change in intraday volatility — both of which are observable.

6 Channel B: dealer-side slope amplification

Channel A pins down the level shift of the kink. Channel B governs the local slope of σ_{SOFR} at the kink, which determines how reserve drains translate into volatility on days when reserves are near the operational threshold. Channel B is governed by the dealer's SLR headroom; it is regime-dependent.

6.1 The slope object

Let $G_\theta(R) \equiv |\partial\sigma_{\mathrm{SOFR}}/\partial R|_{(\theta,Q_\theta);R=R^k(\theta,Q_\theta)}$. From Lemma 2,

$$G_\theta(R^k) \;=\; \frac{|\Psi_1(z_\alpha, z_c^\theta) + \Psi_2(z_\alpha, z_c^\theta)|}{2\sqrt{\Psi(z_\alpha, z_c^\theta)}}, \qquad z_c^\theta \;=\; z_\alpha + h_\theta, \quad h_\theta \;=\; \frac{H(\theta, Q_\theta)}{\sigma_\eta(Q_\theta)}.$$

Define the univariate function

$$\mathcal{G}(z_c) \;\equiv\; \frac{|\Psi_1(z_\alpha, z_c) + \Psi_2(z_\alpha, z_c)|}{2\sqrt{\Psi(z_\alpha, z_c)}}, \qquad z_c \geq z_\alpha,$$

so that $G_\theta(R^k) = \mathcal{G}(z_\alpha + h_\theta)$.

6.2 A covariance representation

For $u \sim \mathcal{N}(0,1)$, define $X(u) := a_u(u - z_\alpha)^+ + b(u - z_c)^+$ and $S(u) := a_u\mathbf{1}\{u > z_\alpha\} + b\mathbf{1}\{u > z_c\}$. Then X is convex with a.e.-derivative $X'(u) = S(u)$, and $\mathrm{Var}(X) = \Psi(z_\alpha, z_c)$.

Lemma 4 (Covariance representation of the slope)**.**

$$\mathrm{Cov}(X, S) \;=\; -\tfrac{1}{2}[\Psi_1(z_\alpha, z_c) + \Psi_2(z_\alpha, z_c)],$$

and consequently (since $\Psi_1 + \Psi_2 < 0$ *on* $\{z_R \leq z_c\}$ *by Lemma 3),*

$$\mathcal{G}(z_c) \;=\; \frac{\mathrm{Cov}(X, S)}{\sigma_X} \;=\; \frac{\mathrm{Cov}(X, S)}{\sqrt{\Psi(z_\alpha, z_c)}}.$$

The proof is in Appendix A. Lemma 4 interprets $\mathcal{G}$ as the regression coefficient of S on standardized X: a Cauchy–Schwarz bound gives $\mathcal{G} \leq \sigma_S$.

6.3 Closed form and unconditional monotonicity of Ψ

Lemma 5 (Closed-form Ψ')**.** *For* $z_c \geq z_\alpha$*, with* $h := z_c - z_\alpha$*,* $g(z_c) := h - M(z_\alpha) + M(z_c)$*,* $p(z_c) := 1 - \Phi(z_c)$*, and* $m(z_c) := M(z_c)$*,*

$$\Psi'(z_c) \;:=\; \frac{d}{dz_c}\Psi(z_\alpha, z_c) \;=\; -2b\big[a_u\, p\, g + a_c\,(1-p)\, m\big] \;<\; 0.$$

The proof is in Appendix A. The closed form is purely algebraic in (z_a, z_c) and the proof uses only $z_a \leq z_c$; substituting $z_a := z_R$ delivers $\Psi_2(z_R, z_c)$, which is how Lemma 3 invokes it.

6.4 Boundary regimes

Lemma 6 (Boundary values of $\mathcal{G}$)**.** $\mathcal{G} : [z_\alpha, \infty) \to \mathbb{R}_+$ *has two regime-defining boundary values:*

$$\underbrace{\mathcal{G}(z_\alpha)}_{\textit{kink-aligned regime}} = a_c K_0, \qquad \underbrace{\lim_{z_c \to \infty} \mathcal{G}(z_c)}_{\textit{slack-dealer regime}} = a_u K_0,$$

where $K_0 := \Phi(z_\alpha) M(z_\alpha)/\sqrt{V(z_\alpha)}$*.*

Proof. Kink-aligned regime. When $z_R = z_c = z_\alpha$, $\eta^c = R$, so $X = a_c(\eta - z_\alpha)^+$ and $S = a_c \mathbf{1}\{u > z_\alpha\}$. Hence $\Psi = a_c^2 V(z_\alpha)$. From Lemma 5-style algebra at $z_c = z_\alpha$ (with $g = 0$, $m = M(z_\alpha)$, $1 - p = \Phi(z_\alpha)$),

$$\mathrm{Cov}(X, S) \;=\; \Phi(z_\alpha) M(z_\alpha)[a_u^2 + a_u b + a_c b] \;=\; a_c^2 \Phi(z_\alpha) M(z_\alpha).$$

So $\mathcal{G}(z_\alpha) = a_c^2 \Phi(z_\alpha) M(z_\alpha)/(a_c\sqrt{V(z_\alpha)}) = a_c K_0$. *Slack-dealer regime.* As $z_c \to \infty$, Mills-ratio gives $1 - \Phi(z_c) \sim \phi(z_c)/z_c$ and $M(z_c) \sim \phi(z_c)/z_c^2$, both vanishing exponentially. Hence $V(z_c) \to 0$, $C(z_\alpha, z_c) \to 0$, $\Psi \to a_u^2 V(z_\alpha)$, and $\mathrm{Cov}(X, S) \to a_u^2 \Phi(z_\alpha) M(z_\alpha)$. So $\mathcal{G}(z_c) \to a_u K_0$. The ratio is $a_c/a_u = \rho$. ☐

The two boundary values correspond to two distinct dealer regimes. In the **kink-aligned regime**, $\eta^c \downarrow R^k$, the dealer SLR constraint binds at exactly the bank's VaR threshold, and any shortfall borrowing meets the constrained-dealer market. The slope is $a_c = 1/\gamma$, set entirely by end-user elasticity. In the **slack-dealer regime**, $\eta^c \gg R^k$, the dealer has so much headroom that the constraint almost never binds when the bank borrows. The slope is $a_u = c/(1 + c\gamma)$, set by dealer marginal cost.

6.5 Boundary derivatives

Lemma 7 (Boundary derivative at the kink-aligned end)**.** *Under A7,* $\mathcal{G}'(z_\alpha^+) < 0$ *—* $\mathcal{G}$ *is strictly decreasing in a right-neighborhood of* $z_c = z_\alpha$*.*

Lemma 8 (Boundary derivative at the slack-dealer end)**.** $\mathcal{G}'(z_c) < 0$ *for* z_c *sufficiently large; the leading-order asymptotic threshold is* $z_c \gtrsim \Phi(z_\alpha)M(z_\alpha)/V(z_\alpha)$*.*

Proofs in Appendix A.

6.6 Regime classification

Theorem 1 (Regime classification of slope amplification)**.** *Under A1–A6:*

1. *(Kink-aligned, exact.) If both* $h_{pre} \downarrow 0$ *and* $h_{post} \downarrow 0$*, then* $G_\theta \to a_c K_0$ *and* $G_{post}/G_{pre} \to 1$*.*
2. *(Slack-dealer, exact.) If both* $h_{pre} \to \infty$ *and* $h_{post} \to \infty$*, then* $G_\theta \to a_u K_0$ *and* $G_{post}/G_{pre} \to 1$*.*
3. *(Cross-regime amplification, exact.) If* $h_{pre} \to \infty$ *and* $h_{post} \downarrow 0$*, then*

$$\frac{G_{post}}{G_{pre}} \to \rho = \frac{1 + c\gamma}{c\gamma}.$$

Symmetrically, $G_{post}/G_{pre} \to 1/\rho$ *if pre is kink-aligned and post is slack-dealer.*

Proof. Direct from Lemma 6 applied to the covariance representation of Lemma 4. □

Honest scope. Theorem 1 delivers exact statements at three corners of the $(h_{\text{pre}}, h_{\text{post}})$ plane. Cross-regime amplification at any specific empirical interior point requires the interpolation result Corollary 3 below, under the maintained Conditions (1)–(2). The corner statements stand without conjecture; the interpolation is what extends them to interior parameter pairs.

Condition 1 (No interior maximum)**.** Under A1–A7, $\mathcal{G} : [z_\alpha, \infty) \to \mathbb{R}_+$ has no critical point at any $z_c \in (z_\alpha, \infty)$ with $\mathcal{G}(z_c) > a_c K_0$.

Condition 2 (No interior minimum)**.** Under A1–A7, $\mathcal{G}$ has no critical point at any $z_c \in (z_\alpha, \infty)$ with $\mathcal{G}(z_c) < a_u K_0$.

Conditions 1 and 2 are verified numerically across the calibration grid (see Section 8); analytic proof of strict-interior monotonicity of $\mathcal{G}$ remains open.

Corollary 3 (Slope-ratio envelope)**.** *Under A1–A7 and Conditions 1, 2, for any two regulatory states* θ', θ'' *with* $h_{\theta'}, h_{\theta''} \geq 0$*,*

$$\frac{G_{\theta'}}{G_{\theta''}} \in [1/\rho, \rho].$$

Proof. Under 1, $\mathcal{G}(z_c) \leq a_c K_0$ on $[z_\alpha, \infty)$. Under 2, $\mathcal{G}(z_c) \geq a_u K_0$. So both $\mathcal{G}(z_\alpha + h_{\theta'})$ and $\mathcal{G}(z_\alpha + h_{\theta''})$ lie in $[a_u K_0, a_c K_0]$, and their ratio lies in $[1/\rho, \rho]$. □

Interpretation. Channel B describes the slope amplification on stress days when dealer headroom is low. The model shows that the regime-shift factor $\rho = (1 + c\gamma)/(c\gamma) = 1 + 1/(c\gamma)$ is large when $c\gamma$ is small — modest dealer marginal cost combined with inelastic end-user demand. At the calibration $c\gamma = 0.5$, $\rho = 3$. Channel B is silent when both regimes sit deep in the slack-dealer regime: the calibrated base case has $h_{\text{pre}} = 84.9$ and $h_{\text{post}} = 633.7$, both far above the kink-aligned boundary, and the ratio $G_{\text{post}}/G_{\text{pre}} = 1.00$ in numerical evaluation. Channel B fires when at least one regime sits in the kink-aligned right-neighborhood: an alternative calibration with $h_{\text{pre}} = 0.498$ delivers $G_{\text{post}}/G_{\text{pre}} = 0.526$, a 47% slope reduction post-reform. The Section 8 numerical exercises bound the empirical regime in which Channel B contributes.

Bidirectional amplification. Theorem 1.3 shows that the kink-aligned-to-slack transition produces a factor-ρ *reduction* in the slope (the $1/\rho$ end). Channel B is therefore bidirectional: clearing-induced kink-alignment amplifies the slope, while SLR-induced dealer slack reduces it. Whether the joint reform raises or lowers the slope at a particular calibration depends on which regime each side sits in, not on the sign of the reform per se.

7 The two-channel decomposition and the operational threshold

This section assembles the two channels into the operational threshold and answers the conceptual question: when is the bank-side level shift the operational threshold shift?

7.1 The decomposition identity

By Corollary 1, $R^\star(\theta, Q, \bar{\sigma})$ is well-defined and uniquely characterized by $\sigma_{\text{SOFR}}(R^\star, \theta, Q) = \bar{\sigma}$ (or by $R^\star = R^k$ if $\bar{\sigma} > \sigma_{\text{SOFR}}(R^k)$). Total differentiation across regimes at fixed $\bar{\sigma}$ delivers the identity

$$\Delta R^\star \;=\; \underbrace{\Delta R^k}_{\text{Channel A}} \;+\; \underbrace{[R^\star - R^k]_{\text{post}} - [R^\star - R^k]_{\text{pre}}}_{\text{Channel B}}.$$

Channel A is closed-form (Corollary 2). Channel B depends on the gap between the operational threshold and the kink in each regime, which in turn depends on G_θ and the higher-order curvature of $\sigma_{\text{SOFR}}(R)$ above the kink. The implicit-function formula for the off-locus contribution is in Appendix B.

7.2 The calibration locus

Define the calibration locus $\mathcal{C} \equiv \{(\bar{\sigma}, \theta, Q) : R^\star(\theta, Q, \bar{\sigma}) = R^k(\theta, Q)\}$: the locus where the Fed's tolerance for SOFR volatility is exactly tight at the kink. On $\mathcal{C}$, the gap $[R^\star - R^k]_\theta = 0$, so Channel B contributes zero to $\Delta R^\star$.

Proposition 2 (Threshold shift on the calibration locus)**.** *Suppose* $(\bar{\sigma}_{pre}, \theta_{pre}, Q_{pre}) \in \mathcal{C}$ *and* $(\bar{\sigma}_{post}, \theta_{post}, Q_{post}) \in \mathcal{C}$ *(allowing* $\bar{\sigma}$ *to be regime-dependent). Then*

$$\Delta R^{\star} \;=\; \Delta R^{k} \;=\; z_{\alpha}\left[\sigma_{\eta}(Q_{post}) - \sigma_{\eta}(Q_{pre})\right] \;>\; 0,$$

and the operational-threshold shift is purely Channel A.

Proof. On $\mathcal{C}$, $R^{\star} = R^{k}$ in each regime by definition. Apply Corollary 2. □

Proposition 2 is not a result that uses Channel B. It is the Channel A shift in clean form, derived under the assumption that the Fed sits on the calibration locus in both regimes. The remaining task is to show that the empirically realistic Federal Reserve sits on this locus.

7.3 Planner microfoundation

The Fed selects reserve supply R on the operational domain $R \in [R^{k}(\theta, Q), \infty)$ to minimize the loss

$$\mathcal{L}(R; \theta, Q) \;=\; \lambda_{R} \cdot R \;+\; \lambda_{\sigma} \cdot \sigma^{2}_{\mathrm{SOFR}}(R, \theta, Q),$$

where $\lambda_{R} > 0$ is the marginal opportunity cost of expanding the Fed balance sheet (foregone Treasury issuance, fiscal-monetary tension) and $\lambda_{\sigma} > 0$ is the marginal cost of money-market volatility (financial-stability disruption, transmission impairment). Both are deep parameters of Fed preferences, not regime-dependent. The domain restriction $R \geq R^{k}$ reflects the operational stance of never deliberately running reserve supply below the aggregate VaR floor.

Strict convexity of $\sigma^{2}_{\mathrm{SOFR}}$ in R on the operational range is the substantive condition for uniqueness of the planner's optimum.

Condition 3 (Strict convexity of $\sigma^{2}_{\mathrm{SOFR}}$ in R)**.** On $\{(z_{R}, z_{c}) : z_{\alpha} \leq z_{R} \leq z_{c}\}$, $\Psi_{11} + 2\Psi_{12} + \Psi_{22} > 0$.

Condition 3 is verified numerically on a 192,000-point grid across $(c\gamma, \alpha) \in [0.01, 2] \times [0.005, 0.1]$ with no violations (Section 8). The diagonal-collapse identity $\Psi_{11} + 2\Psi_{12} + \Psi_{22}\big|_{z_{c}=z_{R}} = a_{c}^{2} V''(z_{R})$ holds to machine precision.

Proposition 3 (Planner FOC under Condition 3)**.** *On the planner's domain* $R \in [R^{k}(\theta, Q), \infty)$*, under Condition 3, the planner's optimum* $R^{P}(\theta, Q; \lambda_{R}/\lambda_{\sigma})$ *exists, is unique, and is characterized as follows.*

1. *(Corner.) If* $\lambda_{R} \geq -\lambda_{\sigma}\sigma_{\eta}(Q)[\Psi_{1}(z_{\alpha}, z_{\alpha} + h_{\theta}) + \Psi_{2}(z_{\alpha}, z_{\alpha} + h_{\theta})]$*, then* $R^{P} = R^{k}$*.*
2. *(Interior.) Otherwise,* R^{P} *is the unique root in* (R^{k}, ∞) *of the FOC*

$$\lambda_{R} \;=\; 2\lambda_{\sigma}\sigma_{\eta}(Q) \cdot \mathrm{Cov}(X, S),$$

evaluated at $z_{R} = R^{P}/\sigma_{\eta}$*,* $z_{c} = (R^{P} + H(\theta, Q))/\sigma_{\eta}$*.*

Proof. $\mathcal{L}$ is continuous and goes to infinity as $R \to \infty$, so attains a minimum on $[R^k, \infty)$. By Condition 3, $\mathcal{L}$ is strictly convex on $[R^k, \infty)$; the FOC has at most one root, and any local minimum is the global minimum. In Case 1, $\mathcal{L}'(R^k) \geq 0$ combined with strict convexity gives $\mathcal{L}' > 0$ on (R^k, ∞), so $\mathcal{L}$ is strictly increasing and the corner is the global minimum. In Case 2, $\mathcal{L}'(R^k) < 0$ combined with $\mathcal{L}'(R) \to \lambda_R > 0$ as $R \to \infty$ delivers a unique interior zero by the intermediate value theorem. ☐

7.4 When does $R^\star = R^k$ in both regimes?

Define the regime-specific constant

$$K_\theta \;\equiv\; 2\sigma_\eta(Q_\theta) \cdot \mathrm{Cov}(X(\cdot; z_\alpha, z_\alpha + h_\theta), S(\cdot; z_\alpha, z_\alpha + h_\theta)).$$

Corollary 4 (Both-corners cone). *Under fixed deep parameters $(\lambda_R, \lambda_\sigma)$ and Condition 3, the calibration-locus condition $R^P(\theta, Q_\theta) = R^k(\theta, Q_\theta)$ holds with strict slack in both pre- and post-reform regimes if and only if*

$$\lambda_R \;>\; \max\{K_{pre}, K_{post}\} \cdot \lambda_\sigma.$$

This condition defines a positive-measure open cone in the $(\lambda_R, \lambda_\sigma)$ plane. In this cone, Proposition 2 delivers the operational-threshold shift, and Channel B does not enter the cross-regime $\Delta R^\star$.

Proof. By Proposition 3, the corner $R^P = R^k$ holds in regime θ if and only if $\lambda_R \geq K_\theta \lambda_\sigma$. Strict slack in both regimes is the open intersection $\{\lambda_R > K_{\mathrm{pre}}\lambda_\sigma\} \cap \{\lambda_R > K_{\mathrm{post}}\lambda_\sigma\} = \{\lambda_R > \max(K_{\mathrm{pre}}, K_{\mathrm{post}})\lambda_\sigma\}$, which is a positive-measure open cone since $K_{\mathrm{pre}}, K_{\mathrm{post}}$ are finite positive constants. The boundary $\lambda_R = K_\theta \lambda_\sigma$ for one regime is a measure-zero knife-edge included in the closure of the cone but not in its interior. ☐

Empirical position. At the base calibration, $K_{\mathrm{pre}} = 0.0198$ and $K_{\mathrm{post}} = 0.0265$, so the both-corners cone is $\lambda_R/\lambda_\sigma > 0.0265$. Backing out the implied Fed preference ratio from the December 2025 RMP activation threshold ($R \approx \$3.0\text{T}$) and the calibrated parameters places the Fed at $\lambda_R/\lambda_\sigma \approx 1000$ — four orders of magnitude inside the cone. The both-corners regime is empirically generic in the sense that any plausible Fed objective function with a tighter aversion to Treasury issuance opportunity cost than to a basis-point of money-market volatility lies inside it. The corner case $\lambda_R/\lambda_\sigma < 0.03$ would require a Fed willing to expand reserves by more than \$30B per basis point of volatility reduction, which is empirically extreme.

Implication for the headline. Corollary 4 closes the conceptual loop. The bank-side Channel A level shift $\Delta R^k = z_\alpha \Delta\sigma_\eta$ is the operational-threshold shift the Fed should target whenever its preference parameters lie in the empirically generic both-corners cone. Outside that cone, the cross-regime $\Delta R^\star$ requires the off-locus implicit-function formula in Appendix B, which carries additional terms in m and β (Channel B). The two channels therefore enter the operational threshold separately: Channel A is the unconditional level shift; Channel B is a regime-dependent slope effect that can amplify or offset Channel A off the calibration locus.

7.5 Comparative statics

Table 2: Two-channel comparative statics.

Parameter	Channel A: R^k	Channel B: $\mathcal{G}$ near kink	$\Delta R^\star$ on locus	$\Delta R^\star$ off locus
$\sigma_\eta \uparrow$ (clearing volatility)	$+$	$+$ via $h \downarrow$	$+$	$+$
$m \uparrow$ (clearing cash drain)	0	$+$ via $\Lambda \downarrow$	0	$+$
$\beta \downarrow$ (SLR relaxation)	0	$-$ via $\Lambda \uparrow$	0	$-$
$c \uparrow$ (intermediation cost)	0	$-$ via $\rho \downarrow$	0	depends
$\gamma \uparrow$ (price elasticity)	0	$-$ via $\rho \downarrow$	0	$-$
$\alpha \downarrow$ (tighter tail tol.)	$+$ via $z_\alpha \uparrow$	$+$ on the LCR range	$+$	$+$

The Channel-B comparative statics rely on $\mathcal{G}$ being monotone in h over the relevant range. The boundary values (Lemma 6) and the corner statements of Theorem 1 hold without any conjecture; the interpolation is conjecture-conditional on Conditions 1 and 2.

8 Calibration and robustness

This section reports the calibration anchors, the numerical verification of Conditions 1–3, the A1 sensitivity exercise, and the A6 boundary check. All exercises use the empirically grounded April 2026 calibration; reproduction scripts are in `code/explore/`.

8.1 Calibration

Two calibrations anchor the analysis. The base calibration places both regimes deep in the slack-dealer regime; an alternative calibration with deliberately tight pre-reform dealer headroom maps to the September 2019 configuration.

The April 2026 FRED anchors used to ground the calibration are: WRESBAL = \$2.90T (April 22, 2026); SOFR = 3.65% (April 23); IORB = 3.65% (April 24); DGS3MO = 3.69% (April 22); RRPONTSYD = \$0.114B (April 23). The model predictions are near-insensitive to τ_0 in this empirical range.

8.2 Channel A magnitude at calibration

By Corollary 2,

$$\Delta R^k \;=\; z_\alpha\,[\sigma_\eta(Q_{\text{post}}) - \sigma_\eta(Q_{\text{pre}})] \;=\; 2.326 \times 9.0 \;\approx\; \$20.94\text{B}.$$

This is 0.70% of the FRBNY Reserve Management Purchases activation threshold (\$3.0T as of December 2025). Sensitivity to $\sigma'_\eta(Q)$: varying the slope from 5 to 40 yields ΔR^k from \$7B to

Table 3: Calibration parameters (base, slack-dealer in both regimes).

Parameter	Pre	Post	Source
Q (clearing intensity)	0.10	0.70	OFR 2025: 60–80% post-mandate
β (SLR weight on UST repo)	0.05	0.005	BIS Basel III SLR; eSLR Treasury exclusion
m (FICC margin)	0.02	0.02	FICC sponsored repo schedule midpoint
k (\$B SLR budget)	75.0	75.0	Liang and Zhu [2025]: dealer capacity
$\tau_0 = i_R - i_M$	-0.0010	-0.0010	FRED IORB-SOFR ($\sim$ 0bp Apr 2026)
c (dealer cost)	0.5	0.5	calibration midpoint
γ (end-user elasticity)	1.0	1.0	calibration midpoint
α (LCR tail)	0.01	0.01	LCR-style
σ_0 (intraday std baseline)	25.0	25.0	FRBNY Liberty Street
$\sigma'_\eta(Q)$ slope	15.0	15.0	reduced-form
σ_η (\$B)	26.50	35.50	implied
$\Lambda = (k - mQ)/\beta$	1499.96	14997.20	implied
$H = (1 + c\gamma)\Lambda - \gamma\tau_0$	2249.94	22495.80	implied
$h = H/\sigma_\eta$	84.90	633.68	implied
$R^k = \sigma_\eta z_\alpha$ (\$B)	61.65	82.59	implied
$\rho = (1 + c\gamma)/(c\gamma)$	3.00	3.00	regime-invariant

\$56B linearly.

8.3 Channel B at calibration

Table 4 reports the cross-regime slope ratio under the two calibrations.

Table 4: Channel B activity by calibration.

Calibration	h_{pre}	h_{post}	G_{pre}	G_{post}	$G_{\text{post}}/G_{\text{pre}}$
Base (slack/slack)	84.90	633.68	2.44×10^{-2}	2.44×10^{-2}	1.000
Kink-aligned pre	0.498	3.617	4.63×10^{-2}	2.44×10^{-2}	0.526

In the base calibration, both regimes sit deep in the slack-dealer regime, so $G_{\text{pre}} = G_{\text{post}} = a_u K_0$ and Channel B is silent. In the kink-aligned-pre / slack-post calibration, Channel B reduces the slope post-reform by 47% — a Theorem 1.3 cross-regime case approaching the asymptote $1/\rho = 0.333$. The economic reading: SLR relief moves dealer headroom from binding-at-the-kink to slack, and the regulated-constraint mechanism stops amplifying volatility off the kink.

8.4 Conditions 1, 2, 3: numerical verification

Condition 3 is verified on a $(c\gamma, \alpha, z_R, z_c)$ grid with $(c\gamma, \alpha) \in \{0.01, \ldots, 5\} \times \{0.1, \ldots, 10\} \times \{0.005, \ldots, 0.1\}$ and z_R, z_c on $[z_\alpha, z_\alpha + 5] \times [z_R, z_R + 5]$, totaling 192,000 grid points. *Zero violations.* The diagonal-collapse identity $\Psi_{11} + 2\Psi_{12} + \Psi_{22}|_{z_c = z_R} = a_c^2 V''(z_R)$ holds with maximum relative error 4.15×10^{-16} (machine precision).

Conditions 1 and 2 are verified across a 120-case standard grid plus 6 stress-test extreme cases with ρ up to 10^4. No real violations beyond floating-point noise (the single apparent violation of 2 has magnitude 1.2×10^{-9} against a boundary value of 5.16×10^{-4}, a relative size of 2.3×10^{-6}).

8.5 A1 sensitivity

Sweeping $\sigma'_\eta(Q)$ from -20 to $+40$:

- At $\sigma'_\eta(Q) = 0$: Channel A vanishes ($\Delta R^k = 0$); Channel B is unchanged.
- At $\sigma'_\eta(Q) < 0$ (the netting case): Channel A reverses sign and contributes negatively.
- $|\Delta R^k|$ scales linearly in $\sigma'_\eta(Q)$.

The structural decomposition is unaffected by the sign of $\sigma'_\eta(Q)$. If empirical work — for instance, on FRBNY Liberty Street intraday SOFR percentiles around the FICC mandate phase-in — finds $\sigma'_\eta(Q) \leq 0$ (consistent with the netting interpretation in Bowman et al. [2024]), the model's structure stands and Channel A simply reverses sign.

8.6 A6 boundary check

A6 requires $(1 + c\gamma)\Lambda > \gamma|\tau_0|$, equivalently $k > mQ + \beta\gamma|\tau_0|/(1 + c\gamma)$. At the base calibration, the binding k is approximately 0.014, against the calibrated $k = 75$ — four orders of magnitude of slack. Even in the kink-aligned calibration ($k = 0.442$), $H_{\text{post}} = \$128.4\text{B}$ is roughly 9,000 times the binding boundary. A6 is comfortably satisfied throughout the empirically relevant parameter range.

8.7 Special cases and benchmarks

$Q = 0$ (no clearing). $\sigma_\eta = \bar{\sigma}_0$, Channel A vanishes, and the model collapses to a single-regulatory-state Lopez-Salido and Vissing-Jorgensen [2023]: a reserve-demand function with no clearing or SLR state variable.

Frictionless dealer ($c \to 0$, $\beta \to 0$). All sub-cases of $c/\beta \to \kappa \in [0, \infty]$ deliver $H \to \infty$ and $\sigma_{\text{SOFR}} \to 0$. The model collapses to the ample-reserves limit.

Pure clearing ($\beta_{\text{post}} = \beta_{\text{pre}}$, Q varying). Both $\Lambda_{\text{post}} - \Lambda_{\text{pre}} < 0$ and $\sigma_\eta^{\text{post}} > \sigma_\eta^{\text{pre}}$, so $h_{\text{post}} < h_{\text{pre}}$ and (under Conditions 1–2) $G_{\text{post}} > G_{\text{pre}}$.

Pure SLR relaxation ($Q_{\text{post}} = Q_{\text{pre}}$, β relaxing). σ_η unchanged, Λ rises, $h_{\text{post}} > h_{\text{pre}}$, and $G_{\text{post}} < G_{\text{pre}}$.

September 2019 (illustrative). The kink-aligned regime ($z_c \downarrow z_\alpha$) of Theorem 1 corresponds structurally to the September 2019 configuration documented by Anbil et al. [2020] and Copeland et al. [2025]: corporate-tax outflow plus large Treasury settlement plus dealer balance-sheet tightness, all pushing $\eta^c \downarrow R^k$. We use this mapping as illustrative, not quantitative: the calibrated $\rho = 3$ is too small to reproduce the observed 500-basis-point SOFR excursion, which would require $\rho \approx 50$ and additional concentration mechanisms outside the present model. The kink-aligned regime captures the structural classification; a quantitative replication would require extending the model to dealer-bank concentration as in Copeland et al. [2025].

9 Implications and the operational rule

9.1 Testable implications

The model delivers four testable predictions, ordered from most direct to most regime-dependent.

I1. Channel A kink shift. As FICC clearing phases in, the reserve threshold at which the SOFR-IORB spread first widens shifts right by $\Delta R^k = z_\alpha \, \Delta\sigma_\eta(Q) \approx \$21\text{B}$ at calibration. *Test design:* Event study around FICC repo-clearing phase-in dates (June 2027 deadline). Estimate the kink location using the Afonso et al. [2024] reserve-demand-elasticity approach; regress the estimated kink on a clearing-phase dummy controlling for runoff trajectory and the SLR state. Predicted magnitude at $\alpha = 1\%$ is approximately \$21B per \$9B increase in σ_η.

I2. Channel B amplification on stress days. On low-headroom days (quarter-ends, tax dates, around large Treasury settlements), the elasticity of SOFR conditional volatility to one-day reserve drains is amplified by approximately a factor $\rho = (1 + c\gamma)/(c\gamma) \approx 3$ relative to slack-dealer days. *Test design:* Pool daily SOFR data, estimate the elasticity of SOFR conditional volatility to one-day reserve changes interacted with a dealer-headroom proxy (the Carlson et al. [2025] dealer-balance-sheet headroom index, or a CMTR-spread indicator). Predicted: an approximate factor-3 amplification on low-headroom days.

I3. Operational-threshold shift in the both-corners cone. For Federal Reserve preferences with $\lambda_R/\lambda_\sigma > 0.027$ (a threshold satisfied with four orders of magnitude of slack at the calibrated revealed preference), the optimal switching threshold satisfies $R^\star_{\text{post}} > R^\star_{\text{pre}}$ with magnitude $\Delta R^\star = \Delta R^k \approx \$21\text{B}$. *Test design:* Compute R^k and $R^\star$ at the Fed's revealed preferences (inferred from the December 2025 RMP activation threshold of approximately \$3.0T) and assess whether the implied threshold shift over the next 18 months is consistent with FRBNY Desk projections for σ_η and the eSLR adoption schedule.

I4. Channel B silent in slack/slack. When dealer headroom is large in both regimes (the empirically modal case at calibration), $G_{\text{post}}/G_{\text{pre}} \to 1$ and Channel B contributes zero to the cross-regime $\Delta R^\star$. Channel A then drives the entire shift. *Test design:* Same data as I2 but with a low-headroom dummy interaction; predict that Channel B coefficient is insignificant on high-headroom days and significant (with the predicted factor-ρ amplification) on low-headroom days.

9.2 An operational rule

The decomposition delivers a concrete rule for the Federal Reserve's switch from passive to active reserve-management tooling:

1. Monitor the FRBNY Liberty Street intraday SOFR volatility percentile. Estimate $\sigma_\eta(Q)$ from the cross-section of intraday SOFR draws at each phase of the FICC mandate.
2. Compute the predicted kink shift: $\Delta R^k = z_\alpha \cdot \Delta\sigma_\eta$, with $z_\alpha = 2.326$ at $\alpha = 1\%$.
3. Update the operational threshold: $R^k_{\text{post}} = R^k_{\text{pre}} + \Delta R^k$.
4. Trigger active SRF or open-market purchases when reserves $R \leq R^k_{\text{post}}$.

On low-headroom days, augment the rule with a Channel B contingency: monitor the dealer-headroom proxy and increase the buffer above R^k_{post} when the proxy indicates kink-aligned conditions.

10 Discussion

What the paper delivers. The contribution is a closed-form decomposition of how the Federal Reserve's operational reserve threshold shifts as mandatory Treasury clearing and SLR recalibration phase in. Channel A is the bank-side level shift, distribution-free and approximately \$21B at calibration. Channel B is the dealer-side slope amplification, regime-dependent with factor up to $\rho = 3$. A planner microfoundation pins down when the bank-side level shift coincides with the planner's optimal switching threshold shift: in a positive-measure cone of preference parameters that contains the empirically realistic Federal Reserve.

What the paper does not deliver. Three honest limitations bound the claims.

First, A1 ($\sigma'_\eta(Q) > 0$) is a maintained reduced-form input. The institutional motivation is settlement-timing fragmentation under mandatory clearing, but the opposing case — clearing reduces σ_η via netting [Bowman et al., 2024] — is not ruled out by the model. The Section 8 sensitivity exercise shows that $|\Delta R^k|$ scales linearly in $\sigma'_\eta(Q)$, so the structural decomposition stands but the sign of Channel A is conditional on the empirical sign of $\sigma'_\eta(Q)$. Empirical resolution requires intraday SOFR percentile data around the FICC mandate phase-in, which is not analyzed in this paper.

Second, Conditions 1–3 are verified numerically on a 192,000-point grid covering the empirically relevant parameter range, but an analytic proof of strict-interior monotonicity of $\mathcal{G}$ remains open. The corner statements of Theorem 1 hold without these conditions; the interior-interpolation envelope of Corollary 3 is conjecture-conditional.

Third, the September 2019 mapping is structural, not quantitative. The model's kink-aligned regime captures the qualitative configuration of Anbil et al. [2020] and Copeland et al. [2025], but the calibrated $\rho = 3$ does not produce the observed 500-basis-point spike. A quantitative replication would require extending the model to dealer-bank concentration, which we leave outside the present scope.

Relationship to the literature. The paper nests Lopez-Salido and Vissing-Jorgensen [2023] as the special case $Q = 0$ with β fixed (Section 8). It is complementary to Chang et al. [2025]: Chang–Klee–Yankov isolate the imperfect-competition channel under price-elastic dealers; the present paper isolates the regulatory-constraint channel under price-taking dealers. The two channels would interact in a richer model. The paper differs from Copeland and Kahn [2024] and Liang and Zhu [2025] by adding an aggregate reserve variable and an equilibrium SOFR comparative static; it differs from Copeland et al. [2025] by being structural and forward-looking rather than empirical and backward-looking. Finally, it differs from Duygan-Bump and Kahn [2026] by providing a closed-form rule where the trilemma framing offers a conceptual taxonomy.

11 Conclusion

The Federal Reserve's switch from passive to active reserve-management tooling is governed by the level of reserves at which the kink in reserve demand begins to bind. Mandatory Treasury clearing shifts that level rightward by $\Delta R^k = z_\alpha \, \Delta\sigma_\eta$, distribution-free and approximately \$21B at the April 2026 calibration. The SLR recalibration reshapes the local slope of SOFR volatility above the kink by a regime-dependent factor up to $\rho = (1 + c\gamma)/(c\gamma) \approx 3$, silent on slack-dealer days and active when dealer headroom is exhausted. A planner microfoundation pins down when the bank-side level shift equals the operational threshold shift: in a positive-measure cone that contains the empirically realistic Federal Reserve. The decomposition gives the Fed a structurally-grounded operational rule that updates as the reforms phase in.

References

Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, 2010.

Gara Afonso, Marco Cipriani, Adam Copeland, Anna Kovner, Gabriele La Spada, and Antoine Martin. The market events of mid-september 2019. Technical Report 918, Federal Reserve Bank of New York Staff Reports, 2019.

Gara Afonso, Domenico Giannone, Gabriele La Spada, and John C. Williams. Scarce, abundant, or ample? a time-varying model of the reserve demand curve. Technical Report 941, Federal Reserve Bank of New York Staff Reports, 2022.

Gara Afonso, Domenico Giannone, Gabriele La Spada, and John C. Williams. Tracking reserve ampleness in real time using reserve demand elasticity. Liberty Street Economics, 2024.

Sriya Anbil, Alyssa Anderson, and Zeynep Senyuz. What happened in money markets in september 2019? Technical report, Federal Reserve Board, FEDS Notes, 2020.

David Bowman, Yesol Huh, and Sebastian Infante. Balance-sheet netting and central clearing of treasury repo. Technical Report 2024-057, Federal Reserve Board, FEDS, 2024.

Mark Carlson, Zack Saravay, and Mary-Frances Tian. Fed repo operations and dealer intermediation. Technical Report 2025-052, Federal Reserve Board, FEDS, 2025.

Briana Chang, Elizabeth Klee, and Vladimir Yankov. Rewiring repo. Technical Report 2025-013, Federal Reserve Board, FEDS, 2025.

Adam Copeland and R. Jay Kahn. Repo intermediation and central clearing: An analysis of sponsored repo. Technical Report 1140, Federal Reserve Bank of New York Staff Reports, 2024.

Adam Copeland, Darrell Duffie, and Yilin Yang. Reserves were not so ample after all. *Quarterly Journal of Economics*, 2025.

Wenxin Du, Benjamin Hébert, and Amy Wang Huber. Are intermediary constraints priced? *Review of Financial Studies*, 36(4), 2023a.

Wenxin Du, Benjamin Hébert, and Wenhao Li. Intermediary balance sheets and the treasury yield curve. *Journal of Financial Economics*, 2023b.

Darrell Duffie, Michael Fleming, Frank Keane, Claire Nelson, Or Shachar, and Peter Van Tassel. Dealer capacity and u.s. treasury market functionality. Technical Report 1070, Federal Reserve Bank of New York Staff Reports, 2023.

Burcu Duygan-Bump and R. Jay Kahn. The central bank balance-sheet trilemma. Technical report, Federal Reserve Board, FEDS Notes, 2026.

Federal Reserve Bank of New York. The implementation of reserve management purchases to maintain ample reserves. Teller Window, 2026.

Joseph G. Haubrich. Balance-sheet policy in an ample-reserves framework: An inventory approach. Technical Report 23-25, Federal Reserve Bank of Cleveland Working Paper, 2023.

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–770, 2013.

Samuel Hempel, R. Jay Kahn, Robert Mann, and Mark Paddrik. Balance-sheet netting in u.s. treasury markets and central clearing. Technical Report 2024-057, Federal Reserve Board, FEDS, 2024.

Anil K. Kashyap, Jeremy C. Stein, Jonathan Wallen, and Joshua Younger. Treasury market dysfunction and the role of the central bank. Technical report, Brookings Papers on Economic Activity, 2025.

Nellie Liang and Hanming Zhu. Clearing the path for treasury market resilience. Technical Report 103, Brookings Working Paper, 2025.

David Lopez-Salido and Annette Vissing-Jorgensen. Reserve demand, interest rate control, and quantitative tightening. *SSRN Working Paper 4371999*, 2023.

Office of Financial Research. How the treasury clearing rule for repo might affect SOFR. OFR Blog (April 22, 2025), 2025.

A Proofs

A.1 Truncated-normal moments

For $\eta \sim \mathcal{N}(0, \sigma^2)$ and $Y_t = (\eta - t)^+$ with $z = t/\sigma$:

$$\mathbb{E}[Y_t] = \sigma M(z), \qquad M(z) = \phi(z) - z[1 - \Phi(z)],$$

$$\mathbb{E}[Y_t^2] = \sigma^2 J(z), \qquad J(z) = (1 + z^2)[1 - \Phi(z)] - z\phi(z),$$
$$\mathrm{Var}(Y_t) = \sigma^2 V(z), \qquad V(z) = J(z) - M(z)^2.$$

For $t_1 \le t_2$ with $z_i = t_i/\sigma$:

$$\mathbb{E}[Y_{t_1} Y_{t_2}] = \sigma^2[(1 + z_1 z_2)(1 - \Phi(z_2)) - z_1 \phi(z_2)],$$

so $\mathrm{Cov}(Y_{t_1}, Y_{t_2}) = \sigma^2 C(z_1, z_2)$ with

$$C(z_1, z_2) = (1 + z_1 z_2)(1 - \Phi(z_2)) - z_1 \phi(z_2) - M(z_1) M(z_2).$$

Useful identities: $V'(z) = -2\Phi(z) M(z)$, $C_1(z_1, z_2) = -\Phi(z_1) M(z_2)$, $C_2(z_1, z_2) = [z_1 \Phi(z_1) + \phi(z_1)](1 - \Phi(z_2)) - \phi(z_2)$.

A.2 Proof of Lemma 4

Use translation-equivariance: shifting both z_a and z_c up by δ is equivalent to shifting the random variable u down by δ. Define $\tilde{X}(\delta; u) := X(u; z_a + \delta, z_c + \delta) = X(u - \delta; z_a, z_c)$, so

$$\Psi(z_a + \delta, z_c + \delta) = \mathrm{Var}(\tilde{X}(\delta; u)).$$

Differentiating:

$$\frac{d}{d\delta}\mathrm{Var}(\tilde{X}) = 2\mathrm{Cov}(\tilde{X}, d\tilde{X}/d\delta).$$

The pathwise derivative is $d\tilde{X}/d\delta = -X'(u-\delta) = -S(u-\delta)$ a.e. Evaluated at $\delta = 0$:

$$\Psi_1 + \Psi_2 \; = \; -2\mathrm{Cov}(X, S).$$

By Lemma 3, $\Psi_1 + \Psi_2 < 0$ on $\{z_R \le z_c\}$, so $\mathrm{Cov}(X, S) > 0$ and $\mathcal{G}(z_c) = \mathrm{Cov}(X, S)/\sqrt{\Psi(z_\alpha, z_c)}$.

A.3 Proof of Lemma 5

Differentiate $\Psi(z_\alpha, z_c) = a_u^2 V(z_\alpha) + b^2 V(z_c) + 2a_u b\, C(z_\alpha, z_c)$ in z_c. Using $V'(z_c) = -2\Phi(z_c)M(z_c)$ and $C_2(z_\alpha, z_c) = [z_\alpha\Phi(z_\alpha) + \phi(z_\alpha)](1-\Phi(z_c)) - \phi(z_c)$:

$$\Psi'(z_c) = -2b^2\Phi(z_c)M(z_c) + 2a_u b\, C_2(z_\alpha, z_c).$$

Algebraic regrouping (writing $p = 1-\Phi(z_c)$, $m = M(z_c)$, $h = z_c - z_\alpha$, $g = h - M(z_\alpha) + M(z_c)$, and using $\phi(z_c) = (1-\Phi(z_c))\cdot z_c +$ derivative-chain identities) yields the boxed expression. The result is purely algebraic in (z_α, z_c) and uses only $z_\alpha \le z_c$; substituting $z_a := z_R$ delivers $\Psi_2(z_R, z_c)$ used in Lemma 3. Both summands $a_u pg$ and $a_c(1-p)m$ are non-negative for $z_c \ge z_\alpha$, and $a_c(1-p)m > 0$ for any finite z_c, so $\Psi'(z_c) < 0$ strictly.

A.4 Proof of Lemma 7

Direct calculation of $\mathcal{G}'(z_\alpha^+)$ from the closed forms of Lemmas 4–5 reduces the sign question to the single-variable inequality

$$V(z_\alpha)\big[V(z_\alpha) + z_\alpha M(z_\alpha)\Phi(z_\alpha) - (1-\Phi(z_\alpha))^2\big] \; > \; \Phi(z_\alpha)^2 M(z_\alpha)^2.$$

This inequality is verified numerically on the LCR-relevant range $z_\alpha \in [1.28, 2.58]$ (and more broadly on $z_\alpha \ge 0.2$, satisfied by A7) with comfortable margin; a tabular grid is reported in the supplementary numerical scripts.

A.5 Proof of Lemma 8

Apply the Mills-ratio asymptotic $1 - \Phi(z_c) \sim \phi(z_c)/z_c$, $M(z_c) \sim \phi(z_c)/z_c^2$, and the corresponding asymptotic for $C(z_\alpha, z_c)$. The leading-order term in $\mathcal{G}'(z_c)$ for large z_c is negative when $z_c \gtrsim \Phi(z_\alpha)M(z_\alpha)/V(z_\alpha)$.

B Off-locus implicit-function formula

Off the calibration locus, where $\bar{\sigma} < \sigma_{\mathrm{SOFR}}(R^k, \theta, Q)$, the operational threshold $R^\star$ exceeds the kink, and total differentiation of the defining identity $\sigma_{\mathrm{SOFR}}(R^\star, \theta, Q) = \bar{\sigma}$ in σ_η at fixed $\bar{\sigma}$ delivers the implicit-function formula

$$\left.\frac{\partial R^\star}{\partial \sigma_\eta}\right|_{\bar\sigma} = -\frac{(\partial \sigma_{\mathrm{SOFR}}/\partial \sigma_\eta)\big|_{R^\star}}{(\partial \sigma_{\mathrm{SOFR}}/\partial R)\big|_{R^\star}}. \tag{3}$$

The numerator carries Channel A through σ_η's direct entry into the variance representation $\sigma^2_{\mathrm{SOFR}} = \sigma^2_\eta \Psi(z_R, z_c)$. The denominator is $-G_\theta(R^\star)$ (negative by Lemma 3), which is the Channel B object evaluated at $R^\star$ rather than at R^k. The cross-regime $\Delta R^\star$ off the locus is the integral of (3) along the regime-change path, which depends explicitly on m, β, c, γ via h_θ in $G_\theta(R^\star)$. We do not derive a closed-form off-locus $\Delta R^\star$.

C FOC variant of the bank's reserve choice

Under a smooth shortfall penalty $\frac{1}{2}\pi[(\eta - r)^+]^2$ replacing the VaR rule, with $\tau_0 \le 0$, the bank solves

$$\max_r \, (i_R - i_M)r - \tfrac{1}{2}\pi \, \mathbb{E}[(\eta - r)^+]^2.$$

The first-order condition is $-\tau_0 = \pi \cdot \mathbb{E}[(\eta - r)^+ \cdot \mathbf{1}\{\eta > r\}]$, which under normality reduces to a smooth equation in r with a unique solution. The comparative statics in σ_η are qualitatively identical to those of the VaR rule: $r^\star$ is increasing in σ_η, so Proposition 1 survives. The closed-form magnitude $\Delta R^k = z_\alpha \Delta \sigma_\eta$ is replaced by an analogous expression involving the smooth-penalty parameter π and the truncated-normal expectation; the qualitative content of Channel A is preserved.