Bank Capital Exposure to Private Credit Funds: A Kinked-Loss Model with Correlated Limited Partners

[Author]

April 24, 2026

Abstract

A bank exposed to a gated private-credit fund through a direct loan book, a warehouse NAV facility, and a subscription line holds three channels of risk that an additive channel-by-channel VaR aggregation measures incorrectly. This paper proves three results. First, expected bank loss is continuous but non-differentiable in the aggregate credit shock at the fund's gate-break threshold, and channel-additive VaR strictly under-reads joint VaR on a non-trivial quantile range—an informational gap arising from an atom at zero in the subscription-line marginal. Second, when the zero-correlation baseline wedge is strictly positive, whether the wedge at the full-correlation endpoint $\rho \to 1^-$ strictly exceeds its zero-correlation endpoint $\rho \to 0^+$ value is characterized by an explicit integral inequality $(C^\star)$ on primitives; when the zero-correlation baseline is zero, this endpoint strict dominance is unconditional. Third, endogenizing the subscription-line haircut and the fund-level gate as the choice of a fund facing a Bertrand-competitive, Basel-constrained bank, the equilibrium covenant pair differs from the first-best pair because the channel-additive Basel capital rule prevents the bank from pricing the joint-VaR wedge into spreads; the fund picks a tighter haircut and a *stricter* gate than social optimality requires, and a supervisory floor on the haircut, on gate laxity, or on both restores the first-best. The formal correlation results are at the two endpoints of $\rho \in (0, 1)$; full-range interior monotonicity is stated as a conjecture. The endpoint characterization is verified by direct evaluation at a policy calibration and across a 97-point grid of primitives and under scale-matched Student-t aggregate shocks. The mechanism differs structurally from tail-dependence super-additivity: the source is the gate-induced atom, not regular variation in the tails.

1 Introduction

In the first two quarters of 2026 the non-traded private-credit sector absorbed the largest redemption shock of its post-2008 history. Blackstone's BCRED received redemption requests totaling \$3.7 billion in the first quarter; the median listed business development company traded at a price-to-NAV ratio near 0.74, the steepest discount since October 2020; JPMorgan, Goldman Sachs, and Barclays widened spreads on back-leverage facilities by 50 to 150 basis points and exercised previously dormant collateral markdown rights. The Office of Financial Research measured \$123 billion of committed bank exposures to private credit in its Brief 26-02 [Office of Financial Research, 2026]; the Basel Committee released a discussion paper on bank-NBFI interconnections [Basel Committee on Banking Supervision, 2025]; the IMF flagged the sector as a first-order source of banking-system tail risk [International Monetary Fund, 2025]. A common concern runs through these episodes and the policy literature. A large bank is typically exposed to a single private-credit fund through several distinct channels simultaneously: a direct loan book, a warehouse facility secured by fund NAV, and a subscription line secured by the fund's uncalled capital commitments. The three channels are independent contributors to aggregate loss with shared dependence on the credit shock; the direct-loan book enters as an additive baseline rather than as a mechanism driver. The policy question is whether assessing these channel risks one at a time gives the right answer for joint tail risk.

This paper characterizes the expected bank-loss function $L(\epsilon)$ induced by this three-channel exposure when the underlying fund has a redemption gate and a correlated limited-partner pool. I establish that L has a kink at the gate-break threshold, that the derivative jump at the kink is strictly increasing in LP correlation with a closed-form expression, and that channel-additive VaR strictly under-reads joint VaR on a characterizable quantile range. For the interior regime where the zero-correlation baseline wedge is already strictly positive, I supply an exact integral inequality $(C^\star)$ that is necessary and sufficient for the wedge at $\rho \to 1^-$ to strictly exceed the wedge at $\rho \to 0^+$, together with a direct numerical certificate at a policy calibration and a 97-point robustness grid. Formal results compare the two endpoints of $\rho \in (0, 1)$; full-range interior monotonicity is stated as Conjecture 1.

Three distinct mechanical ingredients drive the results. The kink (Proposition 1) follows from the continuous-onset multiplier convention that converts a step activation at $\bar{\epsilon}^\star$ into a C^0-but-not-C^1 event. The strict positivity of the channel-additive-VaR wedge (Proposition 3(a)) follows structurally from the subscription-line marginal having an atom at zero of mass $p_S > F(\bar{\epsilon}^\star)$: on a quantile range strictly interior to this atom, the marginal VaR of the subscription channel is zero, so the sum of channel-wise VaRs misses the subscription contribution even though the joint event already assigns positive loss. This strict-positivity property is preserved in any two-or-more-channel model with the same atom-and-continuous-marginal structure. The correlation monotonicity of the expected kink jump (Proposition 2) is the classical Vasicek-Gordy convex-payoff mechanism [Vasicek, 2002, Gordy, 2003]: higher ρ spreads the conditional distribution of the LP default fraction, and the bank's subscription-line payoff is convex in that fraction, so a mean-preserving spread strictly raises expected loss. The endpoint characterization $(C^\star)$ of the wedge's $\rho \to 1^-$ vs. $\rho \to 0^+$ comparison (Proposition 3(c)) has a structurally different source: it is governed by branch-displacement integrals of f against the no-fire and fire conditional masses over windows of widths Δ_N and Δ_F on either side of $F^{-1}(\alpha)$. Proposition 4 establishes that the local Vasicek-Gordy convex-payoff intuition predicts the wrong global sign of $(C^\star)$ in the subscription-dominated regime under both Gaussian and Student-t aggregate shocks; the sign is determined globally by how f distributes mass across the branch-displacement windows.

The paper then endogenizes the covenant pair $(q_c, \bar{\xi}^\star)$ as the choice of a fund facing a Bertrand-competitive, Basel-constrained bank and an atomistic limited-partner pool. This layer directly addresses the concern that Propositions 3–5 describe a measurement problem rather than an economic mechanism: if the agents who set $(q_c, \bar{\xi}^\star)$ internalize the wedge, the gap evaporates at the contracting stage. They do not. Under a Basel capital rule that aggregates the bank's counterparty exposures channel-by-channel (Assumption 8, motivated by Basel Committee on Banking Supervision [2019]), the bank cannot price the joint-VaR wedge through spreads. The fund optimizes against spreads that omit the wedge and chooses a covenant pair that omits it too. Proposition 7 shows the equilibrium wedge strictly exceeds the first-best wedge. Proposition 7(c) pins down the direction: the fund sets a *tighter* haircut and a *stricter* gate than the first-best. The stricter-gate result is a surprising direction, because gates are commonly argued to be too generous to sponsors; here the externality runs the other way. The wedge-reducing effect of gate laxity is internalized by the supervisor but not by the fund, so the supervisor's instrument is a *floor* on gate laxity, not a cap. An alternative microfoundation through a manager reputation cost of gate-break realizations delivers the same direction (see Remark 3 and the accompanying text in Section 4).

I do not claim to resolve the three-channel aggregation problem as a complete regulatory framework. Four scope restrictions shape what the paper delivers. First, the continuous-onset multiplier that converts the gate step into a C^0-but-not-C^1 kink is a tractability convention; a bang-bang gate variant yields the same correlation-amplification comparative statics with a value jump in place of a derivative jump. Second, the paper's super-additivity headline applies to channel-additive aggregation as used in internal bank MIS and in some legacy supervisory exercises, not to Expected Shortfall or to any coherent aggregation scheme; I return to this in Section 5. Third, the dollar figures I report are illustrative calibrations of a 97-point robustness grid whose total missed-capital range spans \$0.5 billion to \$76 billion across reasonable primitive perturbations; the point estimate is not the headline. Fourth, I report one negative robustness finding that repairs a weaker earlier version of the theory: under scale-matched Student-t aggregate shocks with $\nu \in \{3, 5, 10\}$ and subscription-line-to-direct ratios up to 200, the super-additivity direction does not reverse, contrary to what a leading-order Taylor prediction at the kink would suggest. I register the reason as a stand-alone observation: in this class of mixed-discrete-continuous loss distributions, the sign of the super-additivity wedge is governed by global tail integration, not by local geometry at the kink.

The paper sits at the intersection of three literatures. He and Li [2026] develop a dynamic credit-chain model with multi-layer rollover risk; my focus is instead on the bank's channel-aggregation problem for a single fund and my distinguishing ingredients (gate censoring, LP correlation, NAV-versus-commitment collateral distinction) are absent in theirs. Matta and Perotti [2024] characterize optimal gate and settlement rules from a within-fund perspective; I take gate parameters as given and characterize their pass-through into bank-side capital. The VaR super-additivity literature [Daníelsson et al., 2013, Zhu et al., 2023] identifies tail regular variation as a sufficient condition; my mechanism is structurally different, with atom-at-zero mass from the gate structure rather than heavy-tail dependence as the driver. Proposition 2 uses the Gaussian one-factor machinery introduced by Vasicek [2002] and developed by Gordy [2003], applied to a new object (the subscription-line expected excess-shortfall coefficient). The empirical backdrop in Acharya et al. [2025] documents bank credit-line drawdowns from NBFIs as a capital risk channel; Chernenko et al. [2025] measures bank exposure to BDCs; Haque et al. [2025] documents indirect credit supply from banks through NBFIs. I provide a theoretical mechanism consistent with these empirical patterns. Policy documents including Office of Financial Research [2026], International Monetary Fund [2025], and Basel Committee on Banking Supervision [2025] flag the aggregation question but contain no formal model.

Section 2 sets up the environment, agents, and loss decomposition. Section 3 states the measurement results: the kink theorem, the correlation-monotonicity of the jump, the super-additive wedge and its conditional characterization, the global-tail proposition, and the gate-strictness comparative static. Section 4 endogenizes $(q_c, \bar{\xi}^\star)$ under competitive bank pricing, proves the first-best benchmark (Proposition 6), establishes the equilibrium wedge externality (Proposition 7), and delivers the supervisory-restoration result (Proposition 8). Section 5 collects extensions, limitations, the relationship to coherent risk measures, and the mapping to He and Li [2026] and Zhu et al. [2023]. Section 6 develops the policy implications with an illustrative dollar calibration and a statement of the counter-intuitive gate-floor policy lever. Section 7 concludes. Proofs not in the main text are in the appendix.

2 Model

2.1 Environment and timing

There are two dates, $t = 0$ and $t = 1$. All random variables live on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

An aggregate credit shock $\epsilon \in \mathbb{R}$ is drawn at $t = 1$ from a distribution with CDF F and strictly positive, continuously differentiable density f on $\mathbb{R}$. Larger ϵ denotes worse stress. A continuum of limited partners indexed by $i \in [0, 1]$ has idiosyncratic liquidity shocks

$$\eta_i = \sqrt{\rho}\, Z + \sqrt{1-\rho}\, u_i, \qquad Z, u_i \overset{iid}{\sim} N(0, 1), \tag{1}$$

with $\rho \in (0, 1)$ the LP-correlation parameter. Z is an aggregate liquidity factor independent of ϵ. Limited partner i defaults on a capital call if $\eta_i > \tau(\epsilon)$, where $\tau(\epsilon) = \tau_0 - \beta\epsilon$ with $\beta > 0$. Conditional on Z, the default fraction is

$$q(\epsilon, Z; \rho) = \Phi\left(\frac{\sqrt{\rho}\, Z - \tau(\epsilon)}{\sqrt{1-\rho}}\right), \tag{2}$$

with unconditional mean $\mathbb{E}_Z[q(\epsilon, Z; \rho)] = \Phi(-\tau(\epsilon))$, independent of ρ.

The fund's redemption gate breaks once the expected default fraction $\Phi(-\tau(\epsilon))$ exceeds a threshold $\bar{\xi}^\star \in (0, 1)$, equivalently once $\epsilon > \bar{\epsilon}^\star$, where

$$\tau^\star := \Phi^{-1}(1 - \bar{\xi}^\star), \qquad \bar{\epsilon}^\star := (\tau_0 - \tau^\star)/\beta. \tag{3}$$

A warehouse LTV covenant triggers once ϵ exceeds a threshold $\bar{\epsilon}_{\text{LTV}}$. The bank holds three claims: a direct loan book of size B_L with per-unit exposure coefficient $\gamma_D > 0$; a warehouse NAV facility of size B_W with LTV haircut loading $\delta_1 > 0$; and a subscription line of size B_S with collateral cushion $q_c \in (0, 1)$.

2.2 Assumptions

Assumption 1 (Coincident triggers)**.** The gate and warehouse LTV thresholds coincide: $\bar{\epsilon}_{\text{LTV}} = \bar{\epsilon}^\star$.

Assumption 2 (Aggregate-shock density)**.** F admits a strictly positive, C^1 density f on $\mathbb{R}$ and satisfies $\mathbb{E}_F[(\epsilon - \bar{\epsilon}^\star)^+] < \infty$.

Assumption 3 (LP one-factor structure)**.** The LP pool satisfies (1) with $\rho \in (0, 1)$.

Assumption 4 (Interior cushion and gate)**.** $\bar{\xi}^\star, q_c \in (0, 1)$.

Assumption 5 (Non-trivial channels)**.** $B_W, B_S, \gamma_D, \delta_1, \beta > 0$.

Assumption 6 (Gate strictly below cushion)**.** $\bar{\xi}^\star < q_c$.

Assumption 1 is a single-threshold convenience. The two-threshold generalization, in which the gate and the warehouse LTV trigger at distinct thresholds $\bar{\epsilon}^\star$ and $\bar{\epsilon}_{\mathrm{LTV}}$, is an unproved generalization in this paper: the Section 5 double-kink extension sketches the two-threshold case, but we do not prove Proposition 3 in that setting. We discuss decoupling in Section 5. Assumption 6 orders the gate trigger strictly below the subscription-line collateral cushion; empirically, BDC 10-K disclosures place gate thresholds near 15–25% of NAV and subscription-line haircuts near 20–40% of uncalled commitments, consistent with 6.

2.3 Bank-loss decomposition

Total bank loss is

$$L(\epsilon, Z; \rho) = L_D(\epsilon) + L_W(\epsilon) + L_S(\epsilon, Z; \rho), \tag{4}$$

with channel components

$$L_D(\epsilon) = \gamma_D\, \epsilon, \tag{5}$$

$$L_W(\epsilon) = B_W\, \delta_1\, (\epsilon - \bar{\epsilon}^\star)^+, \tag{6}$$

$$L_S(\epsilon, Z; \rho) = B_S\, (\epsilon - \bar{\epsilon}^\star)^+ \left(q(\epsilon, Z; \rho) - q_c\right)^+. \tag{7}$$

The direct-loan channel is linear in stress. The warehouse channel is piecewise linear, activating at the LTV trigger. The subscription-line channel has a product structure: a positive-part multiplier $(\epsilon - \bar{\epsilon}^\star)^+$ gates the channel on the gate event, and the conditional excess LP shortfall $(q - q_c)^+$ determines the loss magnitude, averaged over the systematic factor Z.

Remark 1 (On the continuous-onset multiplier)*.* The multiplier $(\epsilon - \bar{\epsilon}^\star)^+$ is a tractability convention. An alternative bang-bang specification replaces it with an indicator $\mathbf{1}\{\epsilon > \bar{\epsilon}^\star\}$, producing a value jump rather than a derivative jump at $\bar{\epsilon}^\star$. As Appendix D shows, the super-additivity direction and the monotonicity of the wedge in ρ survive under the bang-bang variant; only the local differential structure at $\bar{\epsilon}^\star$ differs. A microfoundation for continuous onset via a continuum of LPs with stress-dependent call timing is straightforward but not needed for the results that follow.

2.4 Equilibrium object

The object of analysis is the distribution of L induced by (4) under the primitives above. Throughout the paper I define

$$\bar{L}(\epsilon; \rho) := \mathbb{E}_Z[L(\epsilon, Z; \rho)] = \gamma_D \epsilon + B_W \delta_1 (\epsilon - \bar{\epsilon}^\star)^+ + B_S (\epsilon - \bar{\epsilon}^\star)^+\, G(\epsilon; \rho), \tag{8}$$

where

$$G(\epsilon; \rho) := \mathbb{E}_Z\left[\left(\Phi(W(\epsilon, Z)) - q_c\right)^+\right], \qquad W(\epsilon, Z) = \frac{\sqrt{\rho}\, Z - \tau(\epsilon)}{\sqrt{1-\rho}}, \tag{9}$$

and I write $G^\star(\rho) := G(\bar{\epsilon}^\star; \rho)$ for the value of G at the kink.

Lemma 1. *Under Assumptions 2–5, $\bar{L}(\cdot; \rho)$ is continuous on $\mathbb{R}$.*

Proof. L_D and L_W are continuous. Continuity of $G(\cdot; \rho)$ follows from dominated convergence, using the integrable envelope supplied by 2. □

3 Results

3.1 The kink theorem

Proposition 1 (Kink at the gate-break threshold)**.** *Under 1–5, $\bar{L}(\cdot;\rho)$ is C^1 on $\mathbb{R}\setminus\{\bar{\epsilon}^\star\}$. At $\bar{\epsilon}^\star$, left- and right-derivatives exist and satisfy*

$$J(\rho) := \lim_{\epsilon\downarrow\bar{\epsilon}^\star} \bar{L}'(\epsilon;\rho) - \lim_{\epsilon\uparrow\bar{\epsilon}^\star} \bar{L}'(\epsilon;\rho) = B_W\delta_1 + B_S\,G^\star(\rho) > 0. \tag{10}$$

Proof. The left-derivative equals γ_D because only L_D is active for $\epsilon < \bar{\epsilon}^\star$. For $\epsilon > \bar{\epsilon}^\star$,

$$\bar{L}'(\epsilon;\rho) = \gamma_D + B_W\delta_1 + B_S G(\epsilon;\rho) + B_S(\epsilon-\bar{\epsilon}^\star)\,\partial_\epsilon G(\epsilon;\rho).$$

The Leibniz step on G is justified by the C^1 integrand in ϵ and dominated convergence. As $\epsilon\downarrow\bar{\epsilon}^\star$, the $(\epsilon-\bar{\epsilon}^\star)$ term vanishes and $G(\epsilon;\rho)\to G^\star(\rho)$. Subtracting the left-derivative yields (10). Strict positivity uses $B_W\delta_1 > 0$. □

The economic content is that below $\bar{\epsilon}^\star$ the bank absorbs linear direct-loan losses only. At $\bar{\epsilon}^\star$ the warehouse LTV covenant and the fund-level gate activate simultaneously, and the subscription line exposes the bank to the convex positive-part tail of LP excess default. The jump $J(\rho)$ is the marginal stress sensitivity that channel-additive aggregation would report below the threshold and miss until ϵ crosses it.

3.2 Correlation monotonicity of the jump

Proposition 2 (Monotonicity in LP correlation)**.** *Under 3–5, $\rho\mapsto G^\star(\rho)$ is strictly increasing on $(0,1)$, with closed-form derivative*

$$\frac{dG^\star}{d\rho} = \frac{1}{4\pi\sqrt{\rho(1-\rho)}}\exp\left(-\frac{(\tau^\star)^2}{2} - \frac{u_c(\rho)^2}{2(1-\rho)}\right) > 0, \tag{11}$$

where $u_c(\rho) = \sqrt{(1-\rho)/\rho}\big(\sqrt{1-\rho}\,\tau^\star + \Phi^{-1}(q_c)\big)$. Under 6, $G^\star(0^+) = 0$ and $G^\star(1^-) = (1-q_c)\bar{\xi}^\star$.

Proof. See Appendix A.2. □

The proof uses a Gaussian u-substitution in the spirit of the one-factor credit-risk literature [Vasicek, 2002, Gordy, 2003]. The economic mechanism is classical: higher ρ spreads the conditional distribution of the LP default fraction, and because the bank's subscription-line payoff is a convex function of that fraction via the positive-part $(q-q_c)^+$, a mean-preserving spread in the conditional distribution raises expected loss.

Corollary 1. *$J(\rho)$ is strictly increasing on $(0,1)$, with $J(0^+) = B_W\delta_1$ and $J(1^-) = B_W\delta_1 + B_S(1-q_c)\bar{\xi}^\star$.*

3.3 The super-additive wedge

Let $L^\Sigma = L_D + L_W + L_S$, and let $h(\epsilon) := \gamma_D\epsilon + B_W\delta_1(\epsilon-\bar{\epsilon}^\star)^+$. Define the *channel-additive wedge*

$$W(\rho;\alpha) := \mathrm{VaR}_\alpha(L^\Sigma) - \sum_{i\in\{D,W,S\}} \mathrm{VaR}_\alpha(L_i). \tag{12}$$

The marginal loss distributions are characterized in the next lemma.

Lemma 2 (Marginals)**.** *L_D has a continuous CDF. L_W has an atom at zero of mass $F(\bar{\epsilon}^\star)$; for $\alpha > F(\bar{\epsilon}^\star)$, $\mathrm{VaR}_\alpha(L_W) = B_W\delta_1(F^{-1}(\alpha) - \bar{\epsilon}^\star)$. L_S has an atom at zero of mass*

$$p_S(\rho) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\infty} f(\epsilon)\,\Phi\big(Z^*(\epsilon;\rho)\big)\,d\epsilon, \quad Z^*(\epsilon;\rho) := \frac{\sqrt{1-\rho}\,\Phi^{-1}(q_c) + \tau(\epsilon)}{\sqrt{\rho}}. \tag{13}$$

Under 6, $p_S(\rho) > F(\bar{\epsilon}^\star)$ strictly, the map $\rho \mapsto p_S(\rho)$ is continuous on $(0,1)$, and

$$p_S(1^-) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\infty} f(\epsilon)(1 - p(\epsilon))\,d\epsilon,$$

where $p(\epsilon) := \Phi(-\tau(\epsilon))$. For $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$, $\mathrm{VaR}_\alpha(L_S) = 0$.

Proof. See Appendix A.3. The $\rho \to 1^-$ limit uses pointwise convergence $\Phi(Z^*(\epsilon;\rho)) \to \Phi(\tau(\epsilon)) = 1 - p(\epsilon)$ and dominated convergence with envelope $f(\epsilon)$. □

The quantile range $(F(\bar{\epsilon}^\star), p_S(\rho))$ is where the subscription-line marginal VaR is exactly zero while the joint event assigns positive loss with positive probability. This range is the site of the wedge.

Proposition 3 (Super-additive wedge)**.** *Fix $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$. Define $p(\epsilon) := \Phi(-\tau(\epsilon))$ and $\epsilon_c := (\tau_0 - \Phi^{-1}(1-q_c))/\beta$.*

(a) Strict positivity. *$W(\rho;\alpha) > 0$. The strict positivity of the wedge follows from the subscription-line marginal's atom at zero alone, and is preserved in any two-or-more-channel model with the same atom-and-continuous-marginal structure.*

(b) Continuity and zero-correlation baseline. *$W(\cdot;\alpha)$ is continuous on $(0,1)$, and*

$$\lim_{\rho\to 0^+} W(\rho;\alpha) = W_0(\alpha) := B_S\big(F^{-1}(\alpha) - \bar{\epsilon}^\star\big)\big(p(F^{-1}(\alpha)) - q_c\big)^+ \ge 0.$$

Under 6, $W_0(\alpha) = 0$ iff $F^{-1}(\alpha) \le \epsilon_c$ (Case 1)*; $W_0(\alpha) > 0$ iff $F^{-1}(\alpha) > \epsilon_c$* (Case 2)*.*

(b′) Case 1 strict amplification. *If $W_0(\alpha) = 0$, then $\lim_{\rho\to 1^-} W(\rho;\alpha) > 0 = W_0(\alpha)$.*

(c) Case 2 exact iff. *If $W_0(\alpha) > 0$, then $\lim_{\rho\to 1^-} W(\rho;\alpha) > W_0(\alpha)$ if and only if*

$$\int_{F^{-1}(\alpha)}^{F^{-1}(\alpha)+\Delta_N} f(\epsilon)(1-p(\epsilon))\,d\epsilon \;<\; \int_{F^{-1}(\alpha)-\Delta_F}^{F^{-1}(\alpha)} f(\epsilon)\,p(\epsilon)\,d\epsilon, \tag{$C^\star$}$$

with

$$\Delta_N = \frac{B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)(p(F^{-1}(\alpha)) - q_c)}{\gamma_D + B_W\delta_1}, \quad \Delta_F = \frac{B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))}{\gamma_D + B_W\delta_1 + B_S(1-q_c)}.$$

Proof. Parts (a), (b), (b′) and (c) are proved in Appendix A.4. The argument for (a) constructs a lower bound on L^Σ on the set $\{\epsilon > F^{-1}(\alpha)\} \cap \{q(\epsilon, Z;\rho) > q_c\}$, which has positive probability under 6. Part (c) uses the branch-displacement Lemma 4 (Appendix) to reduce $\mathbb{P}(L_1^\Sigma \le \ell_0^*) - \alpha$ to the integral difference appearing in ($C^\star$). □

The strict-positivity result (a) captures the informational gap. The subscription-line marginal VaR is zero on the stated quantile range because the atom at zero of mass $p_S(\rho) > \alpha$ swallows the α-quantile of L_S. The joint distribution of L^{Σ} does not inherit this atom on the tail event $\{\epsilon > F^{-1}(\alpha)\}$, because the LP factor Z independently pushes $q(\epsilon, Z; \rho)$ above q_c with positive probability. Summing channel VaRs therefore omits a contribution that the joint tail picks up.

The Case 2 characterization (c) is the paper's main formal content. When $W_0(\alpha) > 0$, positive LP correlation does not unconditionally raise the wedge further; whether it does is pinned down exactly by the integral inequality ($C^{\star}$). The quantities Δ_N and Δ_F are the displacements of the $\rho \to 1^-$ VaR-boundary to the right and left, respectively, measured along the no-fire and fire branches of the conditional loss distribution; the inequality compares the probability mass that leaks out of the no-fire branch against the mass that comes in from the fire branch. Unlike the comparative static of Proposition 2, which is driven locally by Vasicek-Gordy convex-payoff dispersion, the sign of the Case-2 endpoint comparison is determined globally by the branch-displacement integrals Δ_N and Δ_F against the full density f; see Proposition 4 for a concrete manifestation in which the local convex-payoff intuition predicts the wrong global sign.

Corollary 2 (Numerical certificate)**.** *At the calibration* $\gamma_D = 0.3$*,* $B_W = B_S = 1$*,* $\delta_1 = 0.5$*,* $\beta = 2$*,* $\tau_0 = 1.2$*,* $q_c = 0.30$*,* $\bar{\xi}^{\star} = 0.25$*,* $\epsilon \sim N(0, 0.2^2)$*,* $\alpha = 0.99$*, the two sides of* ($C^{\star}$) *evaluate to*

$$\mathrm{LHS}(C^{\star}) = 1.65 \times 10^{-3}, \qquad \mathrm{RHS}(C^{\star}) = 6.31 \times 10^{-3},$$

with margin ratio $3.83\times$*. Hence* $\lim_{\rho \to 1^-} W(\rho; \alpha) > W_0(\alpha)$ *at this calibration. Integrals are evaluated by adaptive Gauss-Kronrod quadrature with relative tolerance* 10^{-8}*. Across a 97-point grid in* $(q_c, \sigma_\epsilon, \bar{\xi}^{\star})$ *around the baseline,* ($C^{\star}$) *holds at every valid grid point, with minimum Case-2 margin* $3.28\times$*.*

Corollary 2 is not an asymptotic statement and does not invoke any Taylor bound. It is a direct evaluation of the two integrals in ($C^{\star}$) at specified primitives. The reproducibility script is `code/explore/iff_certificate.py`; full grid output is in `output/stage3a/corollary2_robustness.json`.

3.4 Global-tail governance of the wedge sign

The asymptotic analysis of ($C^{\star}$) as the subscription book is rescaled $B_S \mapsto \lambda B_S$ yields a tractable primitive condition in the $\lambda \to 0^+$ limit,

$$B_S(1 - q_c)\big(p(F^{-1}(\alpha)) - q_c\big) < q_c(\gamma_D + B_W \delta_1), \tag{C_L}$$

which is the leading-order Taylor approximation of ($C^{\star}$). A leading-order analysis at $\lambda \to \infty$ suggests that (C_L)'s reversal should produce a reversal in ($C^{\star}$). This does not happen under the Gaussian calibration family.

Write $\lambda := B_S$ and hold all other primitives fixed at the Gaussian baseline of Corollary 2. The branch displacements of Lemma 4 then scale asymmetrically in λ: $\Delta_N(\lambda)$ is strictly linear in λ while $\Delta_F(\lambda)$ saturates at $(F^{-1}(\alpha) - \bar{\epsilon}^{\star})(1 - p(F^{-1}(\alpha)))/(1 - q_c)$ as $\lambda \to \infty$, because of the $\lambda(1 - q_c)$ term in the denominator of Δ_F.

Proposition 4 (Global tail dominates local Taylor as $\lambda \to \infty$)**.** *Fix all primitives except* $\lambda := B_S$ *at the Gaussian baseline calibration, and consider the iff condition* ($C^{\star}$) *of Proposition 3(c) parametrized by* λ*.*

(a) $\Delta_N(\lambda) \to \infty$ *linearly in* λ*, while* $\Delta_F(\lambda)$ *is bounded above by* $(F^{-1}(\alpha) - \bar{\epsilon}^{\star})(1 - p(F^{-1}(\alpha)))/(1 - q_c)$ *uniformly in* λ*.*

(b) *Under* 2 *(Gaussian F),* $\mathrm{LHS}(C^\star)(\lambda)$ *is bounded uniformly in* λ *with* $\sup_\lambda \mathrm{LHS}(C^\star)(\lambda) \leq \int_{F^{-1}(\alpha)}^{\infty} f(\epsilon)(1-p(\epsilon))\,d\epsilon < \infty$*, while* $\mathrm{RHS}(C^\star)(\lambda)$ *converges as* $\lambda \to \infty$ *to the finite positive constant* $\int_{\epsilon_\infty}^{F^{-1}(\alpha)} f(\epsilon)\,p(\epsilon)\,d\epsilon > \sup_\lambda \mathrm{LHS}(C^\star)(\lambda)$*, where* $\epsilon_\infty := F^{-1}(\alpha) - (F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))/(1-q_c)$*.*

(c) *Therefore* ($C^\star$) *holds for all sufficiently large* λ *under Gaussian* F*, in contrast to the leading-order Taylor prediction* (C_L)*, which reverses at* $\lambda > \lambda_{crit} = q_c(\gamma_D + B_W\delta_1)/((1-q_c)(p(F^{-1}(\alpha)) - q_c))$*.*

Proof. See Appendix A.6. The numerical certificate for the non-reversal claim is `output/stage3a/iff_certificat` and the heavy-tail robustness is `iff_certificate_heavytail.json`. □

Proposition 4 shows that in the class of mixed-discrete-continuous loss distributions studied here, the leading-order Taylor expansion at the VaR boundary systematically mislocates the sign of the super-additivity wedge when λ is large. The two branch-displacement windows $[F^{-1}(\alpha), F^{-1}(\alpha) + \Delta_N]$ and $[F^{-1}(\alpha) - \Delta_F, F^{-1}(\alpha)]$ scale asymmetrically in λ: the no-fire window grows linearly without bound, while the fire window saturates at a constant determined by $(1-q_c)$ in the fire-branch slope. Because the LHS integrates $f(1-p)$ against the growing window and the underlying Gaussian density f decays faster than the window expands, the LHS remains bounded. The RHS converges to a positive limit. The inequality therefore strengthens rather than reverses as λ grows. The numerical evidence confirms the analytical claim: across the 97-point grid with B_S up to 50, under scale-matched Student-t_ν aggregate shocks with $\nu \in \{3, 5, 10\}$ and B_S up to 200, no reversal is observed, and the minimum margin across the heavy-tail sweep is 1.98× at $\nu = 3$, $B_S \approx 9$.

3.5 Gate strictness and LP correlation: same-sign levers at the endpoints

Proposition 5 (Gate strictness comparative static)**.** *Fix* $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$*. Under* 1–6*:*

(a) $$\lim_{\rho \to 0^+} \frac{\partial W}{\partial \bar{\xi}^\star} = -\frac{B_S(p(F^{-1}(\alpha)) - q_c)^+}{\beta\phi(\tau^\star)} \leq 0,$$ *strict iff* $F^{-1}(\alpha) > \epsilon_c$*.*

(b) $$\lim_{\rho \to 1^-} \frac{\partial W}{\partial \bar{\xi}^\star} = -\frac{1}{\beta\phi(\tau^\star)} \cdot \frac{B_S(1-q_c)E_F}{E_N + E_F} < 0,$$ *where* $E_N, E_F > 0$ *are the no-fire and fire branch density contributions defined in Appendix* A.5*.*

(c) *At both endpoints,* $\partial W/\partial\bar{\xi}^\star \leq 0$*. Numerical evaluation across the 97-point grid confirms the sign throughout, and the magnitude at* $\rho \to 1^-$ *exceeds the magnitude at* $\rho \to 0^+$ *by a factor of 3–5 at the calibration.*

Proof. See Appendix A.5. □

The policy reading is that both levers act in the same direction at the endpoints: at $\rho \to 0^+$ and at $\rho \to 1^-$, tighter gates and lower LP correlation each reduce the wedge. Proposition 5 shows the two endpoint signs match, and numerical verification in Appendix D confirms that the gate-strictness response grows in magnitude as LP correlation increases, with a 3–5× scaling from $\rho \to 0^+$ to $\rho \to 1^-$. I do not claim supermodular economic complementarity, and the interior of $\rho \in (0, 1)$ is covered only conditionally on Conjecture 1; the formal statement is that the two endpoint comparative statics have the same sign and that the gate lever bites harder in the high-correlation endpoint.

3.6 Robustness summary

Table 1 summarizes, for each main result, the proof status and the robustness check.

Table 1: Result-by-result verification and robustness.

Result	Status	Robustness check	Locatio
Prop. 1 (kink)	Proved	Bang-bang variant preserves	App. A
Prop. 2 ($G^\star \uparrow \rho$)	Proved, closed form	Verified at $\rho \in \{0.1, 0.5, 0.9\}$	App. A
Prop. 3(a) (strict positivity)	Proved	97-point grid: holds everywhere	App. A
Prop. 3(b′) (Case 1)	Proved	Unconditional	App. A
Prop. 3(c) (Case 2 iff)	Proved	Cor. 2: 3.83×; grid min 3.28×	App. A
Prop. 4 (global tail)	Proved + numerical	Gaussian and $t_\nu, \nu \in \{3, 5, 10\}$; $B_S \leq 200$	App. A
Prop. 5 (gate CS)	Proved	Matches finite-difference at $\rho \in \{0.001, 0.9, 0.99\}$	App. A

4 Endogenous covenants under competitive bank pricing

Sections 2–3 take the subscription-line haircut q_c and the gate threshold $\bar{\xi}^\star$ as exogenous. A natural concern is that the results characterize a measurement problem rather than an economic mechanism: the wedge disappears if the agents setting $(q_c, \bar{\xi}^\star)$ choose them optimally against a bank that passes the wedge's shadow cost into the warehouse and subscription-line spreads. This section closes that gap. I endogenize $(q_c, \bar{\xi}^\star)$ as the choice of a fund facing a competitively-priced, Basel-constrained bank, compare the equilibrium to the first-best contract, and show that a covenant floor restores the first-best. The first result establishes a first-best benchmark (Proposition 6). The second shows that under a Basel capital rule that aggregates counterparty exposures channel-by-channel, the fund's privately optimal covenants leave a strictly positive wedge (Proposition 7). The third shows that a supervisory floor on q_c, on $\bar{\xi}^\star$, or on both restores the first-best (Proposition 8).

4.1 Contracting environment

I add a single contracting date $t = -1$ to the two-date model of Section 2. At $t = -1$ the fund writes the limited-partnership agreement fixing $(q_c, \bar{\xi}^\star)$ and simultaneously negotiates warehouse and subscription-line facilities of sizes (B_W, B_S) with the bank. The $t = 0$ and $t = 1$ primitives are unchanged. Under Assumption 1 the warehouse LTV covenant activates at the same $\bar{\epsilon}^\star$ as the gate, so the contract pins down both through the single covenant pair.

Bank. The bank is risk-neutral and prices each facility at its Bertrand zero-profit spread. It holds regulatory capital equal to its channel-additive Basel Value-at-Risk,

$$\mathrm{RC}^{\mathrm{Basel}}(q_c, \bar{\xi}^\star) := \mathrm{VaR}_\alpha(L_W) + \mathrm{VaR}_\alpha(L_S), \tag{14}$$

against a shadow cost $\kappa > 0$ per unit of regulatory capital. The direct-loan book L_D sits on a separate bank balance sheet and is priced in an independent spread s_D; it plays no role in the contracting problem. The bank's Basel capital computation aggregates the two contracted facilities by summing their standalone VaRs. The joint-VaR wedge $W = \mathrm{VaR}_\alpha(L^\Sigma) - [\mathrm{VaR}_\alpha(L_D) + \mathrm{VaR}_\alpha(L_W) + \mathrm{VaR}_\alpha(L_S)]$ from Proposition 3 does not enter $\mathrm{RC}^{\mathrm{Basel}}$.

Fund. The fund is risk-neutral and chooses $(q_c, \bar{\xi}^\star, B_W, B_S)$ to maximize

$$\Pi_F = R_F(B_W, B_S) - s_W B_W - s_S B_S - \psi(q_c, \bar{\xi}^\star) - \bar{U}, \tag{15}$$

where R_F is fund revenue (strictly increasing, strictly concave, C^2 in (B_W, B_S)), $\psi(q_c, \bar{\xi}^\star)$ is limited-partner compensation for exposure to covenants the LP pool dislikes, (s_W, s_S) are the bank's posted spreads, and $\bar{U}$ is a constant reservation utility absorbed into the fund's outside option.

Limited partners. The LP pool is atomistic and competitive. LP welfare depends on $(q_c, \bar{\xi}^\star)$ through a function ψ with $\partial_{q_c}\psi > 0$ and $\partial_{\bar{\xi}^\star}\psi < 0$: tighter haircuts and stricter gates each reduce LP welfare by reducing ex-ante liquidity.

Informational structure. LPs observe $(q_c, \bar{\xi}^\star)$ in the partnership agreement. They do not observe the warehouse terms (B_W, s_W) or the bank's capital position. Their compensation depends only on $(q_c, \bar{\xi}^\star)$.

I add three assumptions for the contracting layer.

Assumption 7 (Smoothness of revenue and LP compensation)**.** R_F is C^2, strictly increasing, and strictly concave in (B_W, B_S). ψ is C^2 on the interior of its domain with $\partial_{q_c}\psi > 0$ and $\partial_{\bar{\xi}^\star}\psi < 0$.

Assumption 8 (Basel channel-additive capital rule)**.** The bank's regulatory capital on the warehouse-subscription pair is the channel-additive VaR (14), not $\mathrm{VaR}_\alpha(L_W + L_S)$ and not $\mathrm{VaR}_\alpha(L^\Sigma)$. 8 is postulated and is motivated empirically by the Basel III/IV standardized-approach treatment under Basel Committee on Banking Supervision [2019], which places warehouse NAV facilities and subscription lines in distinct counterparty-level risk-weight buckets without a joint-tail netting adjustment.

Assumption 9 (LP informational asymmetry)**.** LP welfare ψ depends on $(q_c, \bar{\xi}^\star)$ only, not on (B_W, s_W, B_S, s_S). 9 is postulated and is motivated by the standard opaqueness of GP-to-bank warehouse terms in LP side-letter disclosures.

8 is the load-bearing friction driving Proposition 7. 9 is supporting: it blocks the LP from inferring the wedge through its covenant-pricing channel. Neither is microfounded from deeper primitives. I treat both as empirical features of the institutional environment and carry them through the mechanism. Remark 2 returns to the role of 8 after the equilibrium result.

Assumption 10 (Local strict convexity at interior solutions)**.** At any interior solution of the planner's problem (Proposition 6) and of the fund's problem (Proposition 7), the objectives are strictly convex in $(q_c, \bar{\xi}^\star, B_W, B_S)$. The Hessians of $\mathcal{T}$ at $(q_c^{FB}, \bar{\xi}^{\star,FB}, B_W^{FB}, B_S^{FB})$ and of $-\Pi_F$ at $(q_c^*, \bar{\xi}^{\star,*}, B_W^*, B_S^*)$ are positive-definite.

10 is load-bearing for uniqueness of the first-best and equilibrium contracts in Propositions 6–7 and for the second-order Taylor expansion in Proposition 8. I verify 10 numerically at the Corollary 2 calibration for the covenant pair at fixed leverage. A global analytic verification would require characterizing the Hessians of $\mathcal{L}, \mathrm{RC}^{\mathrm{Basel}}, W, \psi, R_F$ jointly; I do not attempt this.

4.2 Aggregate loss and zero-profit spreads

Write the unconditional expected bank loss on the warehouse-subscription pair as

$$\mathcal{L}(q_c, \bar{\xi}^\star) := \mathbb{E}[L_W + L_S], \tag{16}$$

which is C^1 in $(q_c, \bar{\xi}^\star)$ on the interior by dominated convergence and the smoothness of f, p, Φ. Under Bertrand competition the bank posts zero-profit spreads

$$s_W B_W = \mathbb{E}[L_W] + \kappa \mathrm{VaR}_\alpha(L_W), \qquad s_S B_S = \mathbb{E}[L_S] + \kappa \mathrm{VaR}_\alpha(L_S), \tag{17}$$

so the fund's total financing cost is

$$\mathcal{S}(q_c, \bar{\xi}^\star, B_W, B_S) := s_W B_W + s_S B_S = \mathcal{L}(q_c, \bar{\xi}^\star) + \kappa\, \mathrm{RC}^{\mathrm{Basel}}(q_c, \bar{\xi}^\star). \tag{18}$$

The wedge W does not appear in $\mathcal{S}$: Basel's channel-additive rule (14) prevents the bank from charging for it through spreads. The bank absorbs the wedge against equity at unit shadow cost κ.

4.3 First-best

Definition 1 (First-best contract)**.** The first-best covenant and leverage pair $(q_c^{FB}, \bar{\xi}^{\star,FB}, B_W^{FB}, B_S^{FB})$ minimizes total social cost

$$\mathcal{T}(q_c, \bar{\xi}^\star, B_W, B_S) := \mathcal{L}(q_c, \bar{\xi}^\star) + \kappa\, \mathrm{VaR}_\alpha(L^\Sigma; q_c, \bar{\xi}^\star, \rho) + \psi(q_c, \bar{\xi}^\star) - R_F(B_W, B_S), \tag{19}$$

subject to non-negativity and $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$.

The planner internalizes the true joint-VaR capital cost, $\mathrm{VaR}_\alpha(L^\Sigma) = \mathrm{RC}^{\mathrm{Basel}} + W$, including the wedge.

Proposition 6 (First-best characterization)**.** *Under 1–10, the first-best $(q_c^{FB}, \bar{\xi}^{\star,FB}, B_W^{FB}, B_S^{FB})$ at an interior solution satisfies*

$$\partial_{q_c}\mathcal{L} + \kappa\, \partial_{q_c}\mathrm{RC}^{\mathrm{Basel}} + \kappa\, \partial_{q_c} W + \partial_{q_c}\psi = 0, \tag{20}$$

$$\partial_{\bar{\xi}^\star}\mathcal{L} + \kappa\, \partial_{\bar{\xi}^\star}\mathrm{RC}^{\mathrm{Basel}} + \kappa\, \partial_{\bar{\xi}^\star} W + \partial_{\bar{\xi}^\star}\psi = 0, \tag{21}$$

$$\partial_{B_j} R_F = \partial_{B_j}\mathcal{L} + \kappa\, \partial_{B_j}\mathrm{RC}^{\mathrm{Basel}} + \kappa\, \partial_{B_j} W, \quad j \in \{W, S\}. \tag{22}$$

The first-best is unique by 10.

Proof. See Appendix B.1. □

4.4 Equilibrium covenant pair and the wedge externality

Definition 2 (Equilibrium contract)**.** Given the bank's zero-profit spread schedule (17), the equilibrium contract $(q_c^*, \bar{\xi}^{\star,*}, B_W^*, B_S^*)$ maximizes (15) with $\mathcal{S}$ substituted for $s_W B_W + s_S B_S$.

Proposition 7 (Equilibrium wedge under competitive bank pricing)**.** *Under 1–10, assume the first-best of Proposition 6 admits an interior solution with $W(q_c^{FB}, \bar{\xi}^{\star,FB}; \rho, \alpha) > 0$, which holds whenever $\alpha \in (F(\bar{\epsilon}^{\star,FB}), p_S(\rho))$ by Proposition 3(a).*

(a) Fund covenant FOCs. *The equilibrium contract satisfies*

$$\partial_{q_c}\mathcal{L} + \kappa\, \partial_{q_c}\mathrm{RC}^{\mathrm{Basel}} + \partial_{q_c}\psi = 0, \tag{23}$$

$$\partial_{\bar{\xi}^\star}\mathcal{L} + \kappa\, \partial_{\bar{\xi}^\star}\mathrm{RC}^{\mathrm{Basel}} + \partial_{\bar{\xi}^\star}\psi = 0, \tag{24}$$

$$\partial_{B_j} R_F = \partial_{B_j}\mathcal{L} + \kappa\, \partial_{B_j}\mathrm{RC}^{\mathrm{Basel}}, \quad j \in \{W, S\}. \tag{25}$$

The equilibrium FOCs are identical to the first-best FOCs (20)–(22) except that the $\kappa\,\partial.W$ terms are absent.

(b) Equilibrium wedge exceeds first-best wedge. *Whenever $\partial_{q_c} W \neq 0$ or $\partial_{\bar{\xi}^\star} W \neq 0$,*

$$(q_c^*, \bar{\xi}^{\star,*}) \neq (q_c^{FB}, \bar{\xi}^{\star,FB}), \qquad W(q_c^*, \bar{\xi}^{\star,*}; \rho, \alpha) \geq W(q_c^{FB}, \bar{\xi}^{\star,FB}; \rho, \alpha), \tag{26}$$

with strict inequality in Case 2 (Proposition 3(b)–(c) and Proposition 5(a)).

(c) Direction of the covenant gap. *In Case 2 the signs* $\partial_{q_c} W < 0$ *and* $\partial_{\bar{\xi}^\star} W \le 0$ *hold (Proposition 3(b) and Proposition 5). Consequently,*

$$q_c^{FB} > q_c^*, \qquad \bar{\xi}^{\star,FB} \ge \bar{\xi}^{\star,*}. \tag{27}$$

The first-best sets a tighter *subscription-line haircut and a* laxer *gate than the equilibrium.*

Proof. See Appendix B.2. □

Proposition 7(c) delivers a result that reverses the conventional reading of gates. The fund picks a *stricter* gate than social optimality requires: the supervisor needs to force it to loosen, not to tighten. The mechanism is that Proposition 5 pins down $\partial_{\bar{\xi}^\star} W \le 0$ (laxer gates reduce the wedge), so the planner's marginal condition on $\bar{\xi}^\star$ includes a force pushing in the direction of laxity that the fund does not internalize.

Remark 2 (The friction is load-bearing on 8). The decisive friction driving Proposition 7(b) is 8. Under full Basel pricing ($\mathrm{RC}^{\mathrm{Basel}}$ replaced with $\mathrm{VaR}_\alpha(L^\Sigma) = \mathrm{RC}^{\mathrm{Basel}} + W$), $\mathcal{S}$ would gain a κW term and the equilibrium FOCs would gain the missing $\kappa\,\partial.W$ terms, so equilibrium would coincide with first-best. The wedge is a pricing-rule externality, not a fundamental information problem. 9 supports this by ensuring LPs cannot infer and internalize the wedge through their own covenant preferences, but 8 alone produces the gap.

Remark 3 (Alternative microfoundation via manager reputation). An alternative microfoundation replaces 9 with a manager reputation cost of gate-break realizations: the manager bears a private cost proportional to the indicator $\mathbf{1}\{\epsilon > \bar{\epsilon}^\star\}$ in addition to the bank spreads. This cost enters the fund's objective with the same sign as $\partial_{\bar{\xi}^\star}\psi$ and pushes in the same direction as 9, so the equilibrium gate is strictly stricter than the first-best gate. The reputation variant delivers the same direction as (27) on $\bar{\xi}^\star$. I omit the full derivation; it follows the same revealed-preference argument as the proof of Proposition 7(b).

Remark 4 (Run-equilibrium interpretation). A coordination-failure variant treats the gate-break event as a global-games equilibrium in LP withdrawal decisions rather than as a scheduled covenant trigger. In that variant the gate threshold $\bar{\xi}^\star$ is pinned down by the LP coordination refinement rather than by the fund's covenant choice. Section 5 returns to this variant and notes that the narrow primitive regime in which the global-games refinement is well-defined excludes the private-credit calibration of interest, so the mechanism cannot be the operative one for this paper's target environment.

4.5 Supervisory restoration

Proposition 8 (Covenant floors restore first-best)**.** *Under 1–10 and the Case-2 sign conditions of Proposition 7(c):*

(i) Minimum haircut. *A floor* $q_c \ge \underline{q_c}$ *with* $\underline{q_c} = q_c^{FB}$ *binds and delivers* $q_c^{**} = q_c^{FB}$*. The gate component* $\bar{\xi}^{\star,**}$ *solves the equilibrium* $\bar{\xi}^\star$ *FOC at* $q_c = q_c^{FB}$ *and generically differs from* $\bar{\xi}^{\star,FB}$*.*

(ii) Minimum gate laxity. *A floor* $\bar{\xi}^\star \ge \underline{\bar{\xi}^\star}$ *with* $\underline{\bar{\xi}^\star} = \bar{\xi}^{\star,FB}$ *binds (since the equilibrium gate is stricter than first-best under Case 2) and delivers* $\bar{\xi}^{\star,**} = \bar{\xi}^{\star,FB}$*.*

(iii) Joint floor. *Imposing* $\underline{q_c} = q_c^{FB}$ *and* $\underline{\bar{\xi}^\star} = \bar{\xi}^{\star,FB}$ *simultaneously restores the first-best exactly.*

The second-best welfare loss from a single-instrument policy (i) or (ii) is

$$\Delta\mathcal{T} = \tfrac{1}{2}\,\partial^2_{\bar{\xi}^\star}\mathcal{T}\cdot(\bar{\xi}^{\star,**} - \bar{\xi}^{\star,FB})^2 + O((\bar{\xi}^{\star,**} - \bar{\xi}^{\star,FB})^3) \tag{28}$$

(for case (i); the case-(ii) expansion is symmetric), with $\partial^2_{\bar{\xi}^\star}\mathcal{T} > 0$ *at first-best by 10.*

Proof. See Appendix B.3. □

The policy reading of Proposition 8 is that the supervisor has two single-instrument levers, either of which corrects one dimension of the covenant pair. The haircut floor raises q_c; the gate-laxity floor raises $\bar{\xi}^\star$. The second of these is counter-intuitive: it is a *floor* on gate laxity, not a cap, because the private covenant-setter is too strict on gates, not too lax. I return to this in Section 6.

4.6 Why the fund is over-cautious on the gate

The economic reading of (27) is that the fund's gate choice is insufficiently lax from a social perspective. Three forces determine the fund's gate choice. First, laxer gates raise $\mathcal{L}$ (a laxer gate fires more often; higher expected loss). Second, laxer gates raise the LP compensation saved through ψ (LPs prefer laxer gates). Third, laxer gates *reduce* the joint-VaR wedge W through Proposition 5: a laxer gate shifts $\bar{\epsilon}^\star$ rightward, which narrows the quantile range $(F(\bar{\epsilon}^\star), p_S(\rho))$ where the atom-at-zero wedge lives.

The bank internalizes the first force through $\mathcal{L}$ and the second through ψ; both are priced into the fund's problem. The bank does not internalize the third force, because the wedge is absorbed against equity rather than charged to the fund. The fund's private calculus therefore weights only the first two forces against each other and omits the wedge-reducing benefit of gate laxity. The fund chooses a stricter gate than would minimize total social cost.

Two observations follow. First, the over-cautious direction is a consequence of the sign of $\partial_{\bar{\xi}^\star}W$, not of the fund's preferences: the fund prefers laxity for ψ-reasons, and even so the equilibrium gate is stricter than first-best. Second, the result does not depend on the fund being risk-averse, on a coordination-failure channel, or on runs. It follows from the wedge-externality channel alone.

4.7 Sanity check and order-of-magnitude bound

At the Corollary 2 calibration with $\rho = 0.99$, $\kappa = 0.10$ (standard shadow-cost-of-capital estimate from the bank-capital literature; see Kashyap and Stein [2004]), and an LP-compensation curvature $\partial^2_{q_c}\psi$ left as a free parameter, the implicit-function-theorem shift from equilibrium to first-best q_c is

$$q_c^{FB} - q_c^* \approx \frac{-\kappa\,\partial_{q_c}W}{\partial^2_{q_c}(\mathcal{L} + \kappa\mathrm{RC}^{\mathrm{Basel}} + \psi)} \approx \frac{\kappa\,|\partial_{q_c}W|}{\partial^2_{q_c}\mathcal{L} + \partial^2_{q_c}\psi}. \tag{29}$$

Under 10, $\partial^2_{q_c}\psi$ is strictly positive; its magnitude is not identified in this paper because LP covenant-pricing data are not in the present calibration. A ten-fold band $\partial^2_{q_c}\psi \in [1, 100]$ around an illustrative midpoint $\partial^2_{q_c}\psi = 10$ translates into a wedge-shift order of magnitude of

$$\Delta W \approx |\partial_{q_c}W|\cdot(q_c^{FB} - q_c^*) \tag{30}$$

that spans [\$0.1, \$10] billion on the \$400 billion notional used in Section 6, with an illustrative midpoint near \$1 billion. This is 0.5% to 40% of the \$23.58 billion total missed-capital baseline of Proposition 3. The order-of-magnitude band is the honest statement; the point estimate depends on LP-welfare curvature that this paper does not identify.

5 Discussion

5.1 Extensions

Four extensions preserve the mechanism but broaden scope. First, decoupling the gate and warehouse LTV thresholds in 1 produces two kinks at $\bar{\epsilon}_1^\star < \bar{\epsilon}_2^\star$ with the jump formula of Proposition 1 applied at each threshold separately; the super-additive wedge of Proposition 3 generalizes, with the quantile range $(F(\bar{\epsilon}^\star), p_S(\rho))$ of Lemma 2 replaced by a union of analogous ranges at the two thresholds. Second, replacing the Gaussian one-factor LP structure with a Student-t_ν one-factor copula changes Proposition 2's closed form but not its qualitative content; the u-substitution that underlies (11) is Gaussian-specific, but the convex-payoff argument that $G^\star$ is increasing in the factor loading extends to any elliptical one-factor family. Third, introducing a systematic component in the direct-loan channel, $L_D(\epsilon, Z) = \gamma_D \epsilon + \kappa Z$, preserves strict positivity of the wedge (Proposition 3(a)) and adds an interaction term whose sign is positive when κ aligns with the subscription-line channel. Fourth, moving from a single aggregate shock ϵ to a multi-factor representation replaces the kink point with a kink surface; the local structure at the surface carries over pointwise.

5.2 Limitations

Eight limitations deserve explicit acknowledgement. Assumptions 8 (Basel channel-additive rule) and 9 (LP informational asymmetry) are postulated rather than microfounded from deeper primitives; 8 is the load-bearing friction for Proposition 7, motivated empirically by Basel Committee on Banking Supervision [2019] without a structural derivation of why the regulator aggregates counterparty exposures channel-by-channel rather than jointly. Assumption 10 (local strict convexity at interior solutions) is verified numerically at the Corollary 2 calibration but not globally. Assumption 1 collapses the gate and LTV triggers into a single threshold. Assumption 6 orders the gate below the collateral cushion; the ordering is empirically defensible from BDC 10-K disclosures but is not equilibrium-derived, and a complete model would endogenize bilateral contracting between the bank and the fund. The Gaussian one-factor LP structure is load-bearing for the closed form in Proposition 2; relaxing it preserves the sign of the comparative static but not the expression. I do not identify the primitives ρ_{LP}, $\bar{\epsilon}^\star$, or q_c empirically; Section 6 discusses how Form 5500 Schedule H, BDC 10-K LP disclosures, and Call Report RC-C NBFI commitment data would be combined to do so, but I do not attempt the exercise here. Conjecture 1 (stated below) on full-range monotonicity of $\rho \mapsto W(\rho; \alpha)$ is unproved; my formal results cover only the endpoint comparison, and Corollary 3 delivers only existence of an $\rho^\star$ matching any policy-target wedge on the open interval, not uniqueness. Finally, the calibration supporting the Case-1 theorem (Proposition 3(b′)) fires in a bounded range of σ_ϵ (roughly $\sigma_\epsilon \in (0.11, 0.21)$ at other baseline primitives); outside this range the calibration operates in the Case-2 regime and relies on Corollary 2.

Conjecture 1 (Full-range monotonicity)**.** *Under* ($C^\star$) *at a calibration,* $\rho \mapsto W(\rho; \alpha)$ *is strictly increasing on* $(0, 1)$.

Conjecture 1 is supported numerically over twelve Case-2 parameterizations. The policy claims in Section 6 use only the endpoint comparison $W(1^-; \alpha) > W_0(\alpha)$, not interior monotonicity. Conjecture 1 is load-bearing for any reading of the paper's results that interprets LP correlation as a continuous interior policy lever over $\rho \in (0, 1)$: the proved content covers the two endpoints, and an interior-monotonicity policy statement requires the conjecture. The channel-additive-VaR aggregation targeted by the super-additivity headline (Proposition 3(a)) is internal bank MIS and legacy channel-by-channel supervisory exercises; I am not aware of a documented instance of a regulator using channel-additive VaR aggregation under the Basel III/IV IMA or FRTB, and the paper's headline does not claim such usage.

Corollary 3 (IVT existence of a target correlation)**.** *Under* ($C^\star$) *at a calibration, for every* $\bar{W} \in (W_0(\alpha), W(1^-;\alpha))$ *there exists at least one* $\rho^\star \in (0,1)$ *with* $W(\rho^\star;\alpha) = \bar{W}$.

Proof. Continuity of $W(\cdot;\alpha)$ (Proposition 3(b)) and the intermediate-value theorem. □

5.3 Alternative microfoundations of the covenant-setting layer

Section 4 grounds the covenant externality in Assumption 8 (channel-additive Basel rule) and Assumption 9 (LPs do not observe warehouse terms). Two alternative microfoundations deliver the same direction on the covenant gap.

The first replaces 9 with a manager reputation cost of gate-break realizations: the fund manager bears a private cost proportional to the indicator $\mathbf{1}\{\epsilon > \bar{\epsilon}^\star\}$, unobserved by the bank. The cost pushes the fund toward stricter gates in the same direction as 9; a revealed-preference argument analogous to Proposition 7(b) delivers $\bar{\xi}^{\star,FB} \geq \bar{\xi}^{\star,*}$ under the same sign condition $\partial_{\bar{\xi}^\star} W \leq 0$ from Proposition 5. This alternative relaxes the assumption on LP information at the cost of adding a manager preference primitive. I do not carry out the full analysis here; it is a direct variant of Section 4.

The second variant reinterprets the gate-break event as a global-games coordination equilibrium in LP withdrawal, rather than as a scheduled covenant trigger. The gate threshold is then pinned down by the LP coordination refinement (as in Goldstein and Pauzner [2005]'s general framework) rather than by a fund covenant choice. In the private-credit calibration of interest—gates at 15 to 25 percent of NAV, q_c at 30 to 40 percent of uncalled commitments, LP base heavily concentrated in low-frequency-rebalancing endowment and pension capital—the primitive regime required for the global-games refinement to produce an interior equilibrium is not met: the range of aggregate shocks over which LP withdrawal is a strict best response is too narrow given the gate-covenant buffer, and the refinement generically pins $\bar{\xi}^\star$ to a corner. This is a genuine negative finding: the coordination-failure channel is not the operative one in the private-credit environment this paper targets, even though it is the operative one in related environments like open-end corporate bond funds [Goldstein et al., 2017]. Section 4's mechanism is the applicable one here.

5.4 Relationship to coherent risk measures

The super-additivity characterization in Proposition 3 applies to channel-additive VaR aggregation. Expected Shortfall is coherent and therefore sub-additive; for ES,

$$\mathrm{ES}_\alpha(L^\Sigma) \leq \sum_{i \in \{D,W,S\}} \mathrm{ES}_\alpha(L_i)$$

by construction, and the paper's super-additivity headline does not mechanically apply to ES aggregation. This does not defeat the informational-gap reading. Channel-additive aggregation continues to appear in internal bank MIS, in informal channel-by-channel supervisory assessments, and in legacy stress exercises, and the mechanism of Proposition 3, a subscription-line marginal with an atom at zero whose mass exceeds the VaR confidence level, produces the same directional misreading for sum-of-ES-by-channel computed at low α as for channel-additive VaR. The point of the paper is that any aggregation scheme that treats the three channels as independently measured summands misses the atom-driven joint tail contribution; ES-based coherent aggregation of the joint loss distribution avoids this by construction, but requires the joint distribution as an input, which is precisely what channel-additive aggregation tries to sidestep.

5.5 Relationship to He and Li [2026] and Zhu et al. [2023]

He and Li [2026] analyze multi-layer credit chains in which households lend to an entrepreneur through a sequence of intermediary funds facing rollover risk. Their chain-length result minimizes run risk by insulating interim fundamental shocks from triggering liquidation. The mapping to this model goes as follows. A single layer of their chain corresponds to a single fund in my setup; their fund faces rollover risk from its lenders, which in my setup corresponds to the gate-break event at $\bar{\epsilon}^{\star}$. What their model does not contain is (i) the distinction between NAV-collateralized and commitment-collateralized bank debt, (ii) correlated LP capital-call defaults, and (iii) bank-side channel aggregation. Conversely, this paper's model does not contain their chain-length optimization or their dynamic rollover structure. The two papers address complementary questions: He and Li [2026] asks how many intermediary layers minimize run risk for a given contract; this paper asks how a single bank should aggregate risk across channels of a single-layer exposure.

Zhu et al. [2023] derives conditions for asymptotic VaR super-additivity under tail dependence, using regular variation of the stable tail dependence function (STDF). The marginals in Zhu are continuous with regularly varying tails; the mechanism is tail dependence; the super-additivity is asymptotic in the confidence level. Proposition 3 delivers an exact, non-asymptotic super-additivity on a bounded quantile range $(F(\bar{\epsilon}^{\star}), p_S(\rho))$, driven by an atom at zero in one marginal (the subscription-line channel) rather than by tail regular variation. The mathematical object and the economic channel are both different. Proposition 4 implies that tail-heaviness alone (Student-t_ν replacement of the Gaussian F, for ν down to 3) does not invert the super-additivity direction, which further distinguishes the two mechanisms.

6 Policy Implications

6.1 The informational gap reading

Proposition 3(a) states that channel-additive VaR strictly under-reads joint VaR on the quantile range $(F(\bar{\epsilon}^{\star}), p_S(\rho))$. I read this as an informational gap in channel-additive aggregation, not as a critique of any specific regulatory rule. The scope of the gap is channel-additive aggregation as used in internal bank MIS, in informal channel-by-channel supervisory assessments, and in legacy stress exercises that compute a total capital number by summing channel-wise contributions. Modern Basel III/IV aggregation under the SA-CCR and FRTB internal models approach (IMA), and Expected Shortfall computed on the joint loss distribution, fall outside this scope by construction, because coherent aggregation on the joint loss distribution does not partition the problem into channels and the IMA requires joint-distribution inputs rather than channel-wise summation. The gap's mechanism is the atom at zero in the subscription-line marginal: for $\alpha \in (F(\bar{\epsilon}^{\star}), p_S(\rho))$, the channel's standalone VaR is zero while the joint event already assigns positive loss. Proposition 3(c) then pins down exactly when positive LP correlation further widens this gap.

6.2 Illustrative calibration

At the baseline calibration, the channel-additive aggregation miss against a \$400 billion notional private-credit exposure book is \$23.58 billion. I emphasize that this is an illustrative midpoint, not a point estimate of any specific bank's exposure. Across the 97-point primitive grid, the total missed capital ranges from \$0.5 billion to \$76 billion, spanning two orders of magnitude. Table 2 reports the baseline and the robustness range.

The headline of the paper is not the dollar figure. It is the characterization: channel-additive aggregation has an informational gap of characterized sign on a characterized quantile range, and

Table 2: Illustrative missed-capital calibration and robustness range. Baseline: $\gamma_D = 0.3$, $B_W = B_S = 1$, $\delta_1 = 0.5$, $\beta = 2$, $\tau_0 = 1.2$, $q_c = 0.30$, $\bar{\xi}^\star = 0.25$, $\epsilon \sim N(0, 0.2^2)$, $\alpha = 0.99$, book size \$400 billion. Grid: $(q_c, \sigma_\epsilon, \bar{\xi}^\star)$ over 140 points of which 97 are valid under 6 and the Proposition 3 quantile range.

Quantity	Baseline	97-point range
Zero-correlation wedge $W_0(\alpha)$	0.0190	$[0, 0.062]$
Full-correlation wedge $W(1^-;\alpha)$	0.0589	$[0.001, 0.190]$
Incremental correlation-induced wedge	0.0400	$[0.001, 0.134]$
Missed capital at \$400B notional	\$23.58B	[\$0.5B, \$76B]
(C*) margin ratio, Case 2	3.83×	min 3.28×

the sign is pinned down by the primitive-parameter inequality ($C^\star$). Any particular dollar number is a product of the primitive choices. Corollary 2 establishes that ($C^\star$) holds at every valid point of the grid with minimum Case-2 margin 3.28×; Proposition 4 establishes that the conclusion survives scale-matched Student-t replacement of F with degrees of freedom down to $\nu = 3$ and subscription-to-direct ratios up to 200.

6.3 The policy chain

Three implications from the characterization chain into a policy statement. First, joint VaR strictly exceeds the channel-additive sum on the quantile range of Proposition 3(a). Second, at calibrations satisfying ($C^\star$) (which includes every point of the 97-point grid), the wedge at the full-correlation endpoint $\rho \to 1^-$ strictly exceeds its zero-correlation endpoint value (Proposition 3(c)); in the Case-1 regime, this endpoint strict dominance is unconditional (Proposition 3(b′)). Third, tighter fund-level gates and lower LP correlation each reduce the wedge at both endpoints, and the magnitude of the gate-strictness response grows as LP correlation increases from the low- to the high-endpoint (Proposition 5, with the magnitude ordering verified numerically across the 97-point grid). Both levers act in the same direction at the two endpoints of $\rho \in (0, 1)$, and the gate-strictness lever bites harder at the high-correlation endpoint. Interior-to-$(0, 1)$ monotonicity of $\rho \mapsto W(\rho; \alpha)$ beyond the endpoint comparison is stated as Conjecture 1; the policy claims here rely only on the proved endpoint inequality, except where the interior reading is explicitly flagged as conditional on the conjecture.

6.4 What this tells and does not tell a regulator

Three positive statements follow from the characterization. A bank with concentrated LP allocations in correlated investor types (endowments with overlapping chief investment officers, retail wealth platforms with common gating triggers) plus multi-channel PC exposure plus shallow fund-level gates is where the channel-additive gap is largest; the primitive condition ($C^\star$) provides the exact threshold. A regulator assessing whether channel-additive aggregation is adequate for a given bank can compute both sides of ($C^\star$) from observable primitives, given estimates of ρ_{LP}, $\bar{\epsilon}^\star$, and q_c. The direction of the gap under positive LP correlation is pinned down by the characterization in this paper, not an empirical guess.

Three statements I do not make, and a regulator should not infer, follow from the same characterization. I do not propose a specific capital surcharge. The primitives ρ_{LP} and $\bar{\epsilon}^\star$ are imperfectly observable from standard disclosures and require empirical work not undertaken here. Conjecture 1 on full-range monotonicity of $\rho \mapsto W(\rho;\alpha)$ is unproved, so interior statements beyond the endpoint comparison are out of scope. The dollar figure in Section 6.2 is illustrative and should not be read as a target capital reallocation for any specific institution.

6.5 Covenant floors and the counter-intuitive gate lever

Proposition 8 operationalizes the wedge reduction through supervisory covenant floors. The haircut lever is the expected one: a minimum q_c floor raises the subscription-line collateral cushion and closes one margin of the wedge. The gate lever is not. The supervisor's instrument on $\bar{\xi}^{\star}$ is a *floor* on gate laxity, not a cap. The direction is counter-intuitive relative to the conventional reading of gates under which a gate that is set too loosely lets sponsors extract liquidity at LP expense. That reading is orthogonal to the mechanism here: in this paper the fund is over-cautious on gate laxity because it does not internalize Proposition 5's channel ($\partial_{\bar{\xi}^{\star}} W \leq 0$, laxer gates reduce the wedge), which is a bank-capital externality not priced through the fund's spread. The second force runs opposite to the first, and at the covenant-setting stage the second force is the operative one for the wedge.

A regulator implementing the Proposition 8 policy has three instrument options. A minimum q_c (case (i)) operates on the haircut alone and leaves the gate free. A minimum $\bar{\xi}^{\star}$ (case (ii)) operates on gate laxity alone and leaves the haircut free. A joint floor (case (iii)) restores the full first-best. The welfare cost of a single-instrument policy is second-order in the distance between the unconstrained gate and the first-best gate (equation (28)); when the two policy levers are substitutes (Remark 2's numerical observation at the baseline), a single-instrument policy captures most of the first-best gain.

The quantitative size of the equilibrium-first-best shift is governed by the LP-welfare curvature $\partial^2_{q_c}\psi$ as in Section 4.7; under a plausible band for that parameter, the covenant-floor policy delivers a wedge-reduction order of magnitude in the \$0.1–\$10 billion range on the \$400 billion notional basis of Section 6.2, an illustrative \$1 billion midpoint. This does not change the policy claim: the existence and direction of the gap are Proposition 7's content, and the quantitative magnitude is conditional on primitives the present paper does not identify.

6.6 Empirical implementation agenda

A full empirical implementation of the characterization would draw on three data sources. Form 5500 Schedule H pension-type allocation data at the plan-sponsor level would support estimation of LP-base correlation ρ_{LP} via the common-allocation structure of large endowment and pension commitments. BDC 10-K LP disclosures, where available, would directly identify top-LP concentrations and investor type. Call Report Schedule RC-C, as augmented by the post-2024 NBFI commitments field, would identify bank-level subscription-line and warehouse-facility exposures to private-credit counterparties. The mapping from these raw fields to $(\rho_{LP}, \bar{\epsilon}^{\star}, q_c)$ in the model requires structural assumptions not undertaken here. I flag this as the natural policy-implementation agenda; the present paper provides the characterization that gives the implementation a target.

7 Conclusion

A bank exposed to a single gated private-credit fund through direct loans, a warehouse NAV facility, and a subscription line holds three channels of risk that channel-additive VaR aggregation measures incorrectly. Expected bank loss has a kink at the gate-break threshold whose derivative jump is strictly increasing in LP correlation; channel-additive VaR strictly under-reads joint VaR on the quantile range where the subscription-line marginal has its gate-induced atom at zero; when the zero-correlation baseline wedge is strictly positive, the wedge at the full-correlation endpoint strictly exceeds its zero-correlation endpoint value if and only if an explicit integral inequality holds, and a direct numerical certificate establishes the inequality at a policy calibration with a $3.83\times$ margin that survives a 97-point primitive grid and scale-matched Student-t heavy-tail replacement. The sign of the endpoint comparison is governed by global tail integration, not by local geometry at the kink. Gate strictness and LP diversification move the wedge in the same direction at both endpoints of $\rho \in (0, 1)$; full-range interior monotonicity is left as a conjecture. The dollar-figure headline is illustrative of a two-orders-of-magnitude robustness range; the paper's contribution is the mechanism and the exact endpoint iff characterization, not the number.

References

Viral V. Acharya, Manasa Gopal, Maximilian Jager, and Sascha Steffen. Shadow always touches the feet: Implications of bank credit lines to non-bank financial intermediaries. *Working Paper*, 2025. NBER WP 33590.

Basel Committee on Banking Supervision. MAR — calculation of RWA for market risk, and standardised approach for counterparty credit risk. Technical report, Bank for International Settlements, Standard D457, 2019.

Basel Committee on Banking Supervision. Banks' interconnections with non-bank financial intermediaries. Technical report, Bank for International Settlements, Discussion Paper D598, 2025.

Sergey Chernenko, Robert Ialenti, and David Scharfstein. Bank capital and the growth of private credit. *Working Paper*, 2025. SSRN 5097437.

Jon Daníelsson, Bjorn N. Jorgensen, Gennady Samorodnitsky, Mandira Sarma, and Casper G. de Vries. Fat tails, VaR and subadditivity. *Journal of Econometrics*, 172(2):283–291, 2013.

Itay Goldstein and Ady Pauzner. Demand-deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327, 2005.

Itay Goldstein, Hao Jiang, and David T. Ng. Investor flows and fragility in corporate bond funds. *Journal of Financial Economics*, 126(3):592–613, 2017.

Michael B. Gordy. A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12(3):199–232, 2003.

Sharjil Haque, Simon Jang, and Yiming Wang. Indirect credit supply. Technical report, Federal Reserve Board, FEDS 2025-059, 2025.

Zhiguo He and Jian Li. Intermediation via credit chains. *Journal of Finance*, 2026. forthcoming; NBER WP 29632.

International Monetary Fund. Global financial stability report: Shifting ground beneath the calm. Technical report, IMF, 2025.

Anil K. Kashyap and Jeremy C. Stein. Cyclical implications of the Basel II capital standards. *Economic Perspectives, Federal Reserve Bank of Chicago*, 28(1):18–31, 2004.

Rafael Matta and Enrico Perotti. Pay, stay, or delay? how to settle a run. *Review of Financial Studies*, 37(4):1368–1407, 2024.

Office of Financial Research. Measuring counterparty exposures to private credit. Technical report, OFR Brief 26-02, 2026.

Oldrich Vasicek. The distribution of loan portfolio value. *Risk*, 15(12):160–162, 2002.

Chuancun Zhu et al. Asymptotic subadditivity/superadditivity of value-at-risk under tail dependence. *Mathematical Finance*, 33(4), 2023. doi: 10.1111/mafi.12393.

A Proofs

A.1 Proof of Proposition 1

$L_D(\epsilon) = \gamma_D\epsilon$ is C^∞. $L_W(\epsilon) = B_W\delta_1(\epsilon - \bar{\epsilon}^\star)^+$ is C^∞ on $\mathbb{R}\setminus\{\bar{\epsilon}^\star\}$. The function $\epsilon \mapsto (\epsilon - \bar{\epsilon}^\star)^+\, G(\epsilon;\rho)$ is C^∞ on $(\bar{\epsilon}^\star, \infty)$ and identically zero on $(-\infty, \bar{\epsilon}^\star)$, so $\bar{L}$ is C^1 on $\mathbb{R}\setminus\{\bar{\epsilon}^\star\}$.

For $\epsilon < \bar{\epsilon}^\star$, $\bar{L}'(\epsilon;\rho) = \gamma_D$. For $\epsilon > \bar{\epsilon}^\star$,

$$\bar{L}'(\epsilon;\rho) = \gamma_D + B_W\delta_1 + B_S G(\epsilon;\rho) + B_S(\epsilon - \bar{\epsilon}^\star)\,\partial_\epsilon G(\epsilon;\rho).$$

The partial $\partial_\epsilon G(\epsilon;\rho)$ is finite by dominated convergence with the envelope $\|\partial_\epsilon(\Phi(W) - q_c)^+\|_\infty \le \|W_\epsilon\|_\infty < \infty$ on compacts. Taking $\epsilon \downarrow \bar{\epsilon}^\star$, the last term vanishes and $G(\epsilon;\rho) \to G^\star(\rho)$. Subtracting the left-derivative γ_D yields

$$J(\rho) = B_W\delta_1 + B_S G^\star(\rho) > 0$$

under 5 and the non-negativity of $G^\star$. □

A.2 Proof of Proposition 2

At $\epsilon = \bar{\epsilon}^\star$, $\tau(\bar{\epsilon}^\star) = \tau^\star$. Change variables in the Z-integral to $u = (\sqrt{\rho}\,Z - \tau^\star)/\sqrt{1-\rho}$; then $Z = (\sqrt{1-\rho}\,u + \tau^\star)/\sqrt{\rho}$, $dZ = \sqrt{(1-\rho)/\rho}\,du$, and

$$G^\star(\rho) = \int_{\Phi^{-1}(q_c)}^{\infty} (\Phi(u) - q_c)\,\phi(Z(u;\rho))\,\sqrt{(1-\rho)/\rho}\,du.$$

Differentiating under the integral and collecting exponents yields (11). The exponent

$$-\frac{(\tau^\star)^2}{2} - \frac{u_c(\rho)^2}{2(1-\rho)}$$

is finite for $\rho \in (0,1)$, so $dG^\star/d\rho > 0$ strictly.

For the endpoints: as $\rho \to 0^+$, $W(\bar{\epsilon}^\star, Z) \to -\tau^\star$ for every Z, so $\Phi(W) \to 1 - \bar{\xi}^\star < q_c$ under 6, giving $G^\star(0^+) = 0$. As $\rho \to 1^-$, conditional on Z the distribution of q concentrates on $\{0,1\}$ with $\mathbb{P}(q = 1 \mid Z) = \mathbb{P}(\sqrt{\rho}Z > \tau^\star) \to \mathbb{P}(Z > \tau^\star) = \bar{\xi}^\star$, yielding $G^\star(1^-) = (1 - q_c)\bar{\xi}^\star$. □

A.3 Proof of Lemma 2

$L_D = \gamma_D \epsilon$ has the same continuous CDF as ϵ up to affine rescaling. For L_W, $L_W = 0$ iff $\epsilon \le \bar{\epsilon}^\star$, contributing a point mass $F(\bar{\epsilon}^\star)$; for $\epsilon > \bar{\epsilon}^\star$, L_W is strictly increasing in ϵ, yielding $\mathrm{VaR}_\alpha(L_W) = B_W \delta_1(F^{-1}(\alpha) - \bar{\epsilon}^\star)$ for $\alpha > F(\bar{\epsilon}^\star)$.

For L_S, $L_S = 0$ iff $\epsilon \le \bar{\epsilon}^\star$ or $q(\epsilon, Z; \rho) \le q_c$. The probability of the second event conditional on $\epsilon > \bar{\epsilon}^\star$ is $\Phi(Z^*(\epsilon; \rho))$ where Z^* solves $q(\epsilon, Z^*; \rho) = q_c$. Computing explicitly,

$$q(\epsilon, Z^*; \rho) = q_c \iff \frac{\sqrt{\rho}\, Z^* - \tau(\epsilon)}{\sqrt{1-\rho}} = \Phi^{-1}(q_c),$$

which gives $Z^*(\epsilon; \rho) = (\sqrt{1-\rho}\, \Phi^{-1}(q_c) + \tau(\epsilon))/\sqrt{\rho}$. Integrating against f over $(\bar{\epsilon}^\star, \infty)$ yields (13). The strictness $p_S(\rho) > F(\bar{\epsilon}^\star)$ follows directly: Φ is strictly positive everywhere on $\mathbb{R}$, so $\Phi(Z^*(\epsilon; \rho)) > 0$ for every (ϵ, ρ); by 2, $f > 0$ on $(\bar{\epsilon}^\star, \infty)$; hence the integrand $f(\epsilon)\, \Phi(Z^*(\epsilon; \rho))$ is strictly positive on a positive-measure set, so the integral in (13) is strictly positive. No sign condition on Z^* is needed.

Continuity of $\rho \mapsto p_S(\rho)$ on $(0, 1)$ follows from joint continuity of Z^* and dominated convergence with envelope $f(\epsilon)$, which is F-integrable by 2. For the endpoint limit: as $\rho \to 1^-$, $\sqrt{1-\rho}\, \Phi^{-1}(q_c) \to 0$ and $\sqrt{\rho} \to 1$, so $Z^*(\epsilon; \rho) \to \tau(\epsilon)$ pointwise, and $\Phi(Z^*(\epsilon; \rho)) \to \Phi(\tau(\epsilon)) = 1 - p(\epsilon)$ with the integrable envelope $f(\epsilon)$. Dominated convergence yields

$$p_S(1^-) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\infty} f(\epsilon)(1 - p(\epsilon))\, d\epsilon.$$

For $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$, the CDF of L_S at zero exceeds α, so the α-quantile is zero. □

A.4 Proof of Proposition 3

Lemma 3 (VaR continuity in ρ)**.** *Under 1–6 and $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$, the map $\rho \mapsto \mathrm{VaR}_\alpha(L^\Sigma(\cdot; \rho))$ is continuous on $(0, 1)$.*

Proof. Let $G_\rho(\ell) := \mathbb{P}(L^\Sigma \le \ell)$. The joint CDF $G_\rho(\ell)$ is jointly continuous in (ℓ, ρ) on $\mathbb{R} \times (0, 1)$: the integrand defining the CDF is bounded by 1 in absolute value, and the conditional default fraction $q(\epsilon, Z; \rho)$ is jointly continuous in (ϵ, Z, ρ) on $\mathbb{R} \times \mathbb{R} \times (0, 1)$, so dominated convergence applies with envelope equal to $f(\epsilon)$ times the standard normal density of Z. Under 2, F has a positive density on a neighborhood of $F^{-1}(\alpha)$; under 3, the LP factor Z has a continuous density; together these imply G_ρ is strictly increasing in ℓ on a neighborhood of $\mathrm{VaR}_\alpha(L^\Sigma(\cdot; \rho))$. The implicit equation $G_\rho(\ell) = \alpha$ then defines $\ell = \mathrm{VaR}_\alpha(L^\Sigma(\cdot; \rho))$ uniquely on this neighborhood, and continuity of the quantile in ρ follows from the standard quantile continuity lemma for jointly continuous, strictly-monotone-in-ℓ CDFs. □

Part (a): strict positivity. Fix $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$. Then $\mathrm{VaR}_\alpha(L_S) = 0$, $\mathrm{VaR}_\alpha(L_W) = B_W \delta_1(F^{-1}(\alpha) - \bar{\epsilon}^\star)$, and $\mathrm{VaR}_\alpha(L_D) = \gamma_D F^{-1}(\alpha)$. Sum of channel VaRs:

$$\sum_i \mathrm{VaR}_\alpha(L_i) = \gamma_D F^{-1}(\alpha) + B_W \delta_1(F^{-1}(\alpha) - \bar{\epsilon}^\star) = h(F^{-1}(\alpha)).$$

We show $\mathbb{P}(L^\Sigma \le h(F^{-1}(\alpha))) < \alpha$, which implies $\mathrm{VaR}_\alpha(L^\Sigma) > h(F^{-1}(\alpha))$ and hence $W(\rho; \alpha) > 0$. Bookkeep the mass of L^Σ below $h(F^{-1}(\alpha))$ by partitioning on ϵ.

On $\{\epsilon \le \bar{\epsilon}^\star\}$: $L_W = L_S = 0$, so $L^\Sigma = \gamma_D \epsilon \le \gamma_D \bar{\epsilon}^\star < h(F^{-1}(\alpha))$ strictly (since $F^{-1}(\alpha) > \bar{\epsilon}^\star$ under $\alpha > F(\bar{\epsilon}^\star)$ and 2). This partition contributes exactly mass $F(\bar{\epsilon}^\star)$.

On $\{\bar{\epsilon}^\star < \epsilon < F^{-1}(\alpha)\}$: $L^\Sigma = h(\epsilon) + L_S(\epsilon, Z; \rho) \geq h(\epsilon)$, with $h(\epsilon) < h(F^{-1}(\alpha))$ strictly. When $L_S(\epsilon, Z; \rho) = 0$ (i.e., $q(\epsilon, Z; \rho) \leq q_c$), $L^\Sigma = h(\epsilon) < h(F^{-1}(\alpha))$; when $L_S > 0$, $L^\Sigma = h(\epsilon) + L_S$ may or may not exceed $h(F^{-1}(\alpha))$. An upper bound on the contributed mass replaces the indicator with one, giving

$$\mathbb{P}\big(\{\bar{\epsilon}^\star < \epsilon < F^{-1}(\alpha)\} \cap \{L^\Sigma \leq h(F^{-1}(\alpha))\}\big) \leq (\alpha - F(\bar{\epsilon}^\star)) \cdot p_0,$$

where $p_0 := \mathbb{P}(L_S = 0 \mid \bar{\epsilon}^\star < \epsilon < F^{-1}(\alpha)) + \mathbb{P}(L_S > 0, L^\Sigma \leq h(F^{-1}(\alpha)) \mid \bar{\epsilon}^\star < \epsilon < F^{-1}(\alpha)) \leq 1$. We show $p_0 < 1$ strictly. For $\epsilon \in (\bar{\epsilon}^\star, F^{-1}(\alpha))$,

$$\mathbb{P}(L_S > 0 \mid \epsilon) = \mathbb{P}(q(\epsilon, Z; \rho) > q_c \mid \epsilon) = 1 - \Phi(Z^*(\epsilon; \rho)),$$

with $Z^*(\epsilon; \rho) = (\sqrt{1-\rho}\,\Phi^{-1}(q_c) + \tau(\epsilon))/\sqrt{\rho}$ by Lemma 2. For any $\rho \in (0, 1)$ and any finite ϵ, $Z^*(\epsilon; \rho)$ is finite, so $\Phi(Z^*(\epsilon; \rho)) < 1$ strictly. Hence $\mathbb{P}(L_S > 0 \mid \epsilon) > 0$ pointwise on $(\bar{\epsilon}^\star, F^{-1}(\alpha))$. Under 2, F has a positive density on this interval, so it has positive Lebesgue measure. Therefore the conditional probability that $L_S = 0$ on $\{\bar{\epsilon}^\star < \epsilon < F^{-1}(\alpha)\}$ is strictly less than one, and the same strict bound transfers to p_0, giving $p_0 < 1$.

On $\{\epsilon \geq F^{-1}(\alpha)\}$: $h(\epsilon) \geq h(F^{-1}(\alpha))$, and $L^\Sigma \geq h(\epsilon)$, so $L^\Sigma \leq h(F^{-1}(\alpha))$ only on the null event $\{\epsilon = F^{-1}(\alpha), L_S = 0\}$; under 2 this contributes mass zero.

Summing,

$$\mathbb{P}(L^\Sigma \leq h(F^{-1}(\alpha))) \leq F(\bar{\epsilon}^\star) + (\alpha - F(\bar{\epsilon}^\star)) \cdot p_0 < F(\bar{\epsilon}^\star) + (\alpha - F(\bar{\epsilon}^\star)) = \alpha,$$

the strict inequality using $p_0 < 1$. Therefore $\mathrm{VaR}_\alpha(L^\Sigma) > h(F^{-1}(\alpha))$ and $W(\rho; \alpha) > 0$.

Part (b): continuity and baseline. Continuity of $\mathrm{VaR}_\alpha(L^\Sigma)$ in ρ on $(0, 1)$ is Lemma 3. The $\rho \to 0^+$ limit uses the fact that $q(\epsilon, Z; \rho)$ becomes deterministic at $\Phi(-\tau(\epsilon)) = p(\epsilon)$, reducing L_S to $B_S(\epsilon - \bar{\epsilon}^\star)^+(p(\epsilon) - q_c)^+$. Direct computation of $\mathrm{VaR}_\alpha(L^\Sigma)$ at $\rho = 0$ yields $h(F^{-1}(\alpha)) + g(F^{-1}(\alpha))$ with $g(\epsilon) := B_S(\epsilon - \bar{\epsilon}^\star)^+(p(\epsilon) - q_c)^+$, hence $W_0(\alpha) = g(F^{-1}(\alpha))$. The Case-1/Case-2 dichotomy follows from $p(\epsilon_c) = q_c$ and the monotonicity of p.

Part (b′): Case 1 strict amplification. Suppose $W_0(\alpha) = 0$, equivalently $F^{-1}(\alpha) \leq \epsilon_c$, i.e., $p(F^{-1}(\alpha)) \leq q_c$. We show $W(1^-; \alpha) > 0$ by establishing $\mathbb{P}(L_1^\Sigma \leq h(F^{-1}(\alpha))) < \alpha$ in the $\rho \to 1^-$ limit; the argument is a branch decomposition analogous to the one used in the proof of part (c), specialized to the fact that at $\ell_0 = h(F^{-1}(\alpha))$ the no-fire branch coincides with η_N at $F^{-1}(\alpha)$ and the fire branch lies strictly above.

At $\rho \to 1^-$, conditional on Z the default fraction $q(\epsilon, Z; \rho)$ degenerates to a Bernoulli indicator with firing probability $p(\epsilon)$ unconditional on Z; integrating out Z, the limiting distribution of L_S conditional on $\epsilon > \bar{\epsilon}^\star$ places mass $1 - p(\epsilon)$ on the no-fire branch $L_S = 0$ (so $L^\Sigma = \eta_N(\epsilon)$) and mass $p(\epsilon)$ on the fire branch $L_S = B_S(\epsilon - \bar{\epsilon}^\star)(1 - q_c)$ (so $L^\Sigma = \eta_F(\epsilon)$). Thus

$$\mathbb{P}(L_1^\Sigma \leq h(F^{-1}(\alpha))) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\eta_N^{-1}(h(F^{-1}(\alpha)))} f(\epsilon)(1 - p(\epsilon))\, d\epsilon \tag{31}$$

$$+ \int_{\bar{\epsilon}^\star}^{\eta_F^{-1}(h(F^{-1}(\alpha)))} f(\epsilon)\, p(\epsilon)\, d\epsilon. \tag{32}$$

Since $h = \eta_N$ on $[\bar{\epsilon}^\star, \infty)$, $\eta_N^{-1}(h(F^{-1}(\alpha))) = F^{-1}(\alpha)$. Since $\eta_F(\epsilon) = \eta_N(\epsilon) + B_S(1 - q_c)(\epsilon - \bar{\epsilon}^\star) > \eta_N(\epsilon)$ strictly for $\epsilon > \bar{\epsilon}^\star$, $\eta_F^{-1}(h(F^{-1}(\alpha))) < F^{-1}(\alpha)$ strictly. Let $\Delta := F^{-1}(\alpha) - \eta_F^{-1}(h(F^{-1}(\alpha))) > 0$. Substituting and using $f = f(1-p) + fp$:

$$\mathbb{P}(L_1^\Sigma \leq h(F^{-1}(\alpha))) - \alpha = -\int_{F^{-1}(\alpha) - \Delta}^{F^{-1}(\alpha)} f(\epsilon)\, p(\epsilon)\, d\epsilon. \tag{33}$$

Under 2, 6, and $\tau(\cdot)$ strictly decreasing in ϵ, $p(\epsilon) > 0$ on a positive-measure subset of $(F^{-1}(\alpha) - \Delta, F^{-1}(\alpha))$. Hence the right-hand side is strictly negative, so $\mathbb{P}(L_1^\Sigma \le h(F^{-1}(\alpha))) < \alpha$, and $W(1^-;\alpha) > 0 = W_0(\alpha)$.

Part (c): exact iff. Define $\eta_N(\epsilon) := h(\epsilon)$ (no-fire branch: $q = 0$) and $\eta_F(\epsilon) := h(\epsilon) + B_S(1 - q_c)(\epsilon - \bar{\epsilon}^\star)^+$ (fire branch: $q = 1$). Both are strictly increasing on $(\bar{\epsilon}^\star, \infty)$.

Lemma 4 (Branch displacements)**.** *Let $\ell_0^* := h(F^{-1}(\alpha)) + g(F^{-1}(\alpha))$, the $\rho = 0^+$ VaR of L^Σ. In Case 2,*

$$\Delta_N := \eta_N^{-1}(\ell_0^*) - F^{-1}(\alpha) = \frac{B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)(p(F^{-1}(\alpha)) - q_c)}{\gamma_D + B_W\delta_1} > 0,$$

$$\Delta_F := F^{-1}(\alpha) - \eta_F^{-1}(\ell_0^*) = \frac{B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))}{\gamma_D + B_W\delta_1 + B_S(1 - q_c)} > 0.$$

Proof. For Δ_N: $\eta_N(\eta_N^{-1}(\ell_0^*)) = \ell_0^* = h(F^{-1}(\alpha)) + g(F^{-1}(\alpha))$ and $\eta_N(F^{-1}(\alpha)) = h(F^{-1}(\alpha))$ since $F^{-1}(\alpha) > \bar{\epsilon}^\star$. Subtracting and using the constant slope $\gamma_D + B_W\delta_1$ of η_N on this interval yields the first identity. For Δ_F: $\eta_F(F^{-1}(\alpha)) - \ell_0^* = B_S(1 - q_c)(F^{-1}(\alpha) - \bar{\epsilon}^\star) - g(F^{-1}(\alpha)) = B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))$, which is positive under 6; dividing by the constant slope of η_F delivers Δ_F. □

At $\rho \to 1^-$, L^Σ concentrates along the two branches η_N and η_F with conditional weights $1 - p(\epsilon)$ and $p(\epsilon)$ respectively. Hence

$$\mathbb{P}(L_1^\Sigma \le \ell_0^*) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\eta_N^{-1}(\ell_0^*)} f(\epsilon)(1 - p(\epsilon))\,d\epsilon + \int_{\bar{\epsilon}^\star}^{\eta_F^{-1}(\ell_0^*)} f(\epsilon)\,p(\epsilon)\,d\epsilon. \tag{34}$$

By Lemma 4, $\eta_F^{-1}(\ell_0^*) < F^{-1}(\alpha) < \eta_N^{-1}(\ell_0^*)$. Subtract $\alpha = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{F^{-1}(\alpha)} f\,d\epsilon$ and use $f = f(1-p) + fp$:

$$\begin{aligned}
\mathbb{P}(L_1^\Sigma \le \ell_0^*) - \alpha &= \int_{F^{-1}(\alpha)}^{\eta_N^{-1}(\ell_0^*)} f(\epsilon)(1 - p(\epsilon))\,d\epsilon - \int_{\eta_F^{-1}(\ell_0^*)}^{F^{-1}(\alpha)} f(\epsilon)\,p(\epsilon)\,d\epsilon &(35)\\
&= \int_{F^{-1}(\alpha)}^{F^{-1}(\alpha)+\Delta_N} f(1-p)\,d\epsilon - \int_{F^{-1}(\alpha)-\Delta_F}^{F^{-1}(\alpha)} fp\,d\epsilon. &(36)
\end{aligned}$$

Under 2, L_1^Σ has a continuous CDF on a neighborhood of ℓ_0^*, so $\mathbb{P}(L_1^\Sigma \le \ell_0^*) < \alpha$ iff $\mathrm{VaR}_\alpha(L_1^\Sigma) > \ell_0^*$ iff $W(1^-;\alpha) > W_0(\alpha)$. This is equivalent to ($C^\star$). The reverse inequality and equality cases follow identically. □

A.5 Proof of Proposition 5

The gate threshold $\bar{\xi}^\star$ enters the wedge only through $\bar{\epsilon}^\star(\bar{\xi}^\star) = (\tau_0 - \tau^\star)/\beta$ with $\tau^\star = \Phi^{-1}(1 - \bar{\xi}^\star)$; $\partial\bar{\epsilon}^\star/\partial\bar{\xi}^\star = 1/(\beta\phi(\tau^\star)) > 0$.

Part (a). At $\rho \to 0^+$, $W_0(\alpha) = B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star)^+(p(F^{-1}(\alpha)) - q_c)^+$. The map $p(F^{-1}(\alpha)) = \Phi(-\tau(F^{-1}(\alpha)))$ does not depend on $\bar{\xi}^\star$ (it depends on primitive $\tau_0, \beta, F^{-1}(\alpha)$ only). Differentiating,

$$\frac{\partial W_0}{\partial \bar{\xi}^\star} = -B_S(p(F^{-1}(\alpha)) - q_c)^+ \cdot \frac{\partial \bar{\epsilon}^\star}{\partial \bar{\xi}^\star} = -\frac{B_S(p(F^{-1}(\alpha)) - q_c)^+}{\beta\phi(\tau^\star)},$$

which is non-positive and strict iff $F^{-1}(\alpha) > \epsilon_c$.

Part (b). At $\rho \to 1^-$, the joint loss concentrates along branches η_N and η_F with conditional weights $1 - p(\epsilon)$ and $p(\epsilon)$. Let $\ell_1^\Sigma(\bar{\xi}^\star)$ be the solution of $\mathcal{P}(\ell; \bar{\xi}^\star) := \mathbb{P}(L_1^\Sigma \le \ell) = \alpha$:

$$\mathcal{P}(\ell; \bar{\xi}^\star) = F(\bar{\epsilon}^\star) + \int_{\bar{\epsilon}^\star}^{\eta_N^{-1}(\ell)} f(\epsilon)(1 - p(\epsilon))\, d\epsilon + \int_{\bar{\epsilon}^\star}^{\eta_F^{-1}(\ell)} f(\epsilon)\, p(\epsilon)\, d\epsilon.$$

Implicit differentiation in $\bar{\xi}^\star$ (recall $\bar{\epsilon}^\star = \bar{\epsilon}^\star(\bar{\xi}^\star)$; η_N, η_F depend on $\bar{\xi}^\star$ only through $\bar{\epsilon}^\star$ via h; on each branch $\partial_{\bar{\xi}^\star} \eta_N^{-1}(\ell) = -B_W \delta_1 \cdot (\partial \bar{\epsilon}^\star / \partial \bar{\xi}^\star)/\eta_N'$ because $\eta_N(\epsilon; \bar{\xi}^\star) = \gamma_D \epsilon + B_W \delta_1 (\epsilon - \bar{\epsilon}^\star)$, and similarly for η_F^{-1}) yields

$$E_N \cdot \ell_1^{\Sigma\prime}(\bar{\xi}^\star) + E_F \cdot \ell_1^{\Sigma\prime}(\bar{\xi}^\star) = -\frac{\partial \mathcal{P}}{\partial \bar{\xi}^\star}\bigg|_{\ell = \ell_1^\Sigma},$$

where

$$E_N := \frac{f(\eta_N^{-1}(\ell_1^\Sigma))(1 - p(\eta_N^{-1}(\ell_1^\Sigma)))}{\gamma_D + B_W \delta_1}, \qquad E_F := \frac{f(\eta_F^{-1}(\ell_1^\Sigma))\, p(\eta_F^{-1}(\ell_1^\Sigma))}{\gamma_D + B_W \delta_1 + B_S(1 - q_c)},$$

both strictly positive under 2 and 6. The $\partial \mathcal{P}/\partial \bar{\xi}^\star$ term collects the lower-limit contribution at $\bar{\epsilon}^\star$ (mass $f(\bar{\epsilon}^\star) \cdot \partial \bar{\epsilon}^\star / \partial \bar{\xi}^\star$, which cancels across the three integrals up to net zero because $p + (1-p) = 1$) and the branch-slope contribution $(\partial \bar{\epsilon}^\star / \partial \bar{\xi}^\star) \cdot B_W \delta_1 \cdot [E_N + E_F]$ from the η_N^{-1}, η_F^{-1} dependence on $\bar{\epsilon}^\star$. Subtracting the channel-additive sensitivity $\partial_{\bar{\xi}^\star}[\gamma_D F^{-1}(\alpha) + B_W \delta_1 (F^{-1}(\alpha) - \bar{\epsilon}^\star)] = -B_W \delta_1 \cdot \partial \bar{\epsilon}^\star / \partial \bar{\xi}^\star$ leaves the residual fire-branch contribution

$$\lim_{\rho \to 1^-} \frac{\partial W}{\partial \bar{\xi}^\star} = -\frac{\partial \bar{\epsilon}^\star}{\partial \bar{\xi}^\star} \cdot \frac{B_S(1 - q_c)\, E_F}{E_N + E_F} = -\frac{1}{\beta \phi(\tau^\star)} \cdot \frac{B_S(1 - q_c)\, E_F}{E_N + E_F} < 0.$$

Part (c). Both endpoint derivatives are non-positive. The magnitude ordering $|\partial W/\partial \bar{\xi}^\star|_{\rho \to 1^-} > |\partial W/\partial \bar{\xi}^\star|_{\rho \to 0^+}$ is verified numerically across the 97-point grid; at the baseline calibration the ratio is 3–5×. A fully analytic dominance argument is not available because the E_N, E_F ratio modulates the scaling in a calibration-dependent way; we therefore report the sign analytically and the magnitude ordering numerically. □

A.6 Proof of Proposition 4

Write $\lambda := B_S$ and hold $\gamma_D, B_W, \delta_1, \beta, \tau_0, q_c, \bar{\xi}^\star, \alpha$, and F fixed at the baseline.

Part (a): branch-displacement scaling. By Lemma 4,

$$\Delta_N(\lambda) = \frac{\lambda(F^{-1}(\alpha) - \bar{\epsilon}^\star)(p(F^{-1}(\alpha)) - q_c)}{\gamma_D + B_W \delta_1}, \qquad \Delta_F(\lambda) = \frac{\lambda(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))}{\gamma_D + B_W \delta_1 + \lambda(1 - q_c)}.$$

$\Delta_N(\lambda)$ is strictly linear in λ with positive slope (under Case 2, $p(F^{-1}(\alpha)) > q_c$), so $\Delta_N(\lambda) \to \infty$ linearly as $\lambda \to \infty$. For $\Delta_F(\lambda)$, divide numerator and denominator by λ:

$$\Delta_F(\lambda) = \frac{(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))}{(\gamma_D + B_W \delta_1)/\lambda + (1 - q_c)} \nearrow \Delta_F^\infty := \frac{(F^{-1}(\alpha) - \bar{\epsilon}^\star)(1 - p(F^{-1}(\alpha)))}{1 - q_c},$$

strictly increasing in λ with finite supremum Δ_F^∞, by monotone convergence of the $1/\lambda$ term. This gives the uniform upper bound claimed in (a).

Part (b): boundedness of LHS and convergence of RHS. For the LHS, since the integrand $f(\epsilon)(1-p(\epsilon))$ is nonnegative on $\mathbb{R}$,

$$\mathrm{LHS}(C^\star)(\lambda) = \int_{F^{-1}(\alpha)}^{F^{-1}(\alpha)+\Delta_N(\lambda)} f(\epsilon)(1-p(\epsilon))\,d\epsilon \;\le\; \int_{F^{-1}(\alpha)}^{\infty} f(\epsilon)(1-p(\epsilon))\,d\epsilon =: M_L.$$

M_L is finite because f is integrable and $1-p \le 1$. This bound is uniform in λ, so $\sup_\lambda \mathrm{LHS}(C^\star)(\lambda) \le M_L$.

For the RHS, define $\epsilon_\infty := F^{-1}(\alpha) - \Delta_F^\infty$. By monotone convergence (the integration interval $[F^{-1}(\alpha) - \Delta_F(\lambda), F^{-1}(\alpha)]$ is increasing in λ, with nonnegative integrand),

$$\mathrm{RHS}(C^\star)(\lambda) \nearrow \int_{\epsilon_\infty}^{F^{-1}(\alpha)} f(\epsilon)\,p(\epsilon)\,d\epsilon =: M_R.$$

$M_R > 0$ under 2 and $p > 0$ on $(\epsilon_\infty, F^{-1}(\alpha))$ (nonempty because $\Delta_F^\infty > 0$ in Case 2).

We now verify $M_R > M_L$ under the Gaussian baseline calibration. Both integrals are directly computable: at the baseline primitives of Corollary 2, numerical evaluation gives $M_L \le 2.36\times10^{-3}$ and $M_R \ge 7.05\times10^{-3}$; see `output/stage3a/iff_certificate_log.txt`. Hence $M_R > M_L$ strictly, and this strict gap is preserved throughout the 97-point grid (`output/stage3a/corollary2_robustness.json`) and across scale-matched Student-t_ν replacement of F for $\nu \in \{3, 5, 10\}$ (`iff_certificate_heavytail.json`).

Part (c): non-reversal of (C*) for large λ. Combining (a) and (b): for every $\lambda \ge \lambda_0$ where λ_0 is the smallest value at which $\mathrm{RHS}(C^\star)(\lambda) > M_L$,

$$\mathrm{LHS}(C^\star)(\lambda) \le M_L < \mathrm{RHS}(C^\star)(\lambda),$$

so ($C^\star$) holds. The leading-order Taylor condition (C_L) reverses at $\lambda > \lambda_{\mathrm{crit}} = q_c(\gamma_D + B_W\delta_1)/((1-q_c)(p(F^{-1}(\alpha)) - q_c))$, which at the baseline evaluates to $\lambda_{\mathrm{crit}} \approx 0.43$. For every $\lambda > \lambda_{\mathrm{crit}}$, the Taylor prediction says ($C^\star$) reverses, while the exact integral argument says it does not. The discrepancy is driven entirely by the asymmetric scaling established in part (a): the Taylor expansion treats Δ_N and Δ_F symmetrically to leading order and misses that Δ_F saturates while Δ_N grows. □

B Proofs for Section 4

B.1 Proof of Proposition 6

The objective $\mathcal{T}$ in (19) is the sum of four terms. $\mathcal{L}$ is C^1 on the interior of $(0, q_c^{\max}) \times (0,1)$ by dominated convergence and the smoothness of f, p, Φ established in Section 3. ψ and R_F are C^2 by 7. The wedge $W = \mathrm{VaR}_\alpha(L^\Sigma) - \mathrm{RC}^{\mathrm{Basel}}$ is C^1 in $(q_c, \bar{\xi}^\star)$ on the interior of the Case-2 parameter space: the joint quantile $\mathrm{VaR}_\alpha(L^\Sigma)$ is the α-quantile of a smooth mixture integral over the (ϵ, Z) density, which is C^1 in $(q_c, \bar{\xi}^\star)$ whenever the underlying density and the loss kernels are C^1 in primitives, by standard quantile-smoothness under continuous densities (see, e.g., the arguments in the proof of Lemma 3). The Basel capital $\mathrm{RC}^{\mathrm{Basel}}$ is C^1 on the same region by Lemma 2. Hence $\mathcal{T}$ is C^1 on the interior region where $\alpha \in (F(\bar{\epsilon}^\star), p_S(\rho))$. Taking interior first-order conditions and equating to zero yields (20)–(22). Uniqueness of the interior solution follows from 10. □

B.2 Proof of Proposition 7

Part (a). Under Bertrand competition, the bank's zero-profit condition is (17), and the fund's financing cost is $\mathcal{S} = \mathcal{L} + \kappa\,\mathrm{RC}^{\mathrm{Basel}}$ by (18). Substituting into the fund's objective (15),

$$\Pi_F = R_F(B_W, B_S) - \mathcal{L}(q_c, \bar{\xi}^\star) - \kappa\,\mathrm{RC}^{\mathrm{Basel}}(q_c, \bar{\xi}^\star) - \psi(q_c, \bar{\xi}^\star) - \bar{U}.$$

The wedge W does not appear: it is absorbed against bank equity and is not passed through spreads. Taking first-order conditions and equating to zero yields (23)–(25). Uniqueness of the interior maximizer follows from 10.

Part (b): revealed-preference wedge inequality. Define

$$\mathcal{T}(q_c, \bar{\xi}^\star) := \mathcal{L}(q_c, \bar{\xi}^\star) + \kappa\,\mathrm{RC}^{\mathrm{Basel}}(q_c, \bar{\xi}^\star) + \kappa W(q_c, \bar{\xi}^\star; \rho, \alpha) + \psi(q_c, \bar{\xi}^\star)$$

and

$$\mathcal{M}(q_c, \bar{\xi}^\star) := \mathcal{L}(q_c, \bar{\xi}^\star) + \kappa\,\mathrm{RC}^{\mathrm{Basel}}(q_c, \bar{\xi}^\star) + \psi(q_c, \bar{\xi}^\star),$$

so that $\mathcal{T} = \mathcal{M} + \kappa W$ at fixed leverage. By Proposition 6 the first-best $(q_c^{FB}, \bar{\xi}^{\star,FB})$ minimizes $\mathcal{T}$; by part (a) the equilibrium $(q_c^*, \bar{\xi}^{\star,*})$ minimizes $\mathcal{M}$. Optimality of the first-best for $\mathcal{T}$ gives

$$\mathcal{M}(\mathrm{FB}) + \kappa W(\mathrm{FB}) \le \mathcal{M}(\mathrm{EQ}) + \kappa W(\mathrm{EQ}), \tag{37}$$

and optimality of the equilibrium for $\mathcal{M}$ gives $\mathcal{M}(\mathrm{EQ}) \le \mathcal{M}(\mathrm{FB})$, equivalently $\mathcal{M}(\mathrm{EQ}) - \mathcal{M}(\mathrm{FB}) \le 0$. Subtracting the second from (37),

$$\kappa W(\mathrm{FB}) - \kappa W(\mathrm{EQ}) \le \mathcal{M}(\mathrm{EQ}) - \mathcal{M}(\mathrm{FB}) \le 0,$$

so $W(\mathrm{FB}) \le W(\mathrm{EQ})$. Strict inequality holds when FB $\neq$ EQ, which holds whenever $\partial_{q_c} W \neq 0$ or $\partial_{\bar{\xi}^\star} W \neq 0$ by 10 and the non-coincidence of first-order conditions (20)–(21) versus (23)–(24).

Part (c): direction. Evaluate $\partial_{q_c}\mathcal{T}$ at $q_c = q_c^*$: by (23), $\partial_{q_c}\mathcal{M}|_{q_c=q_c^*} = 0$, so

$$\partial_{q_c}\mathcal{T}\big|_{q_c=q_c^*} = \partial_{q_c}\mathcal{M}\big|_{q_c=q_c^*} + \kappa\,\partial_{q_c} W\big|_{q_c=q_c^*} = \kappa\,\partial_{q_c} W\big|_{q_c=q_c^*} < 0$$

in Case 2 by Proposition 3(b) ($\partial_{q_c} W_0 = -B_S(F^{-1}(\alpha) - \bar{\epsilon}^\star) < 0$ and Case-2 continuity at $q_c = q_c^*$). Since $\partial^2_{q_c}\mathcal{T} > 0$ at FB by 10, reaching $\partial_{q_c}\mathcal{T} = 0$ from a negative value requires $q_c^{FB} > q_c^*$. The same argument applied at $\bar{\xi}^\star = \bar{\xi}^{\star,*}$ with $\partial_{\bar{\xi}^\star} W \le 0$ from Proposition 5 delivers $\bar{\xi}^{\star,FB} \ge \bar{\xi}^{\star,*}$. □

B.3 Proof of Proposition 8

For case (i), adding the constraint $q_c \ge q_c^{FB}$ to the equilibrium problem and using $q_c^* < q_c^{FB}$ from Proposition 7(c) makes the constraint bind, so $q_c^{**} = q_c^{FB}$. The constrained $\bar{\xi}^\star$ satisfies the equilibrium $\bar{\xi}^\star$ FOC at $q_c = q_c^{FB}$, namely

$$\partial_{\bar{\xi}^\star}\mathcal{L}(q_c^{FB}, \bar{\xi}^\star) + \kappa\,\partial_{\bar{\xi}^\star}\mathrm{RC}^{\mathrm{Basel}}(q_c^{FB}, \bar{\xi}^\star) + \partial_{\bar{\xi}^\star}\psi(q_c^{FB}, \bar{\xi}^\star) = 0,$$

which generically differs from the first-best $\bar{\xi}^\star$ FOC (21) by the $\kappa\,\partial_{\bar{\xi}^\star} W$ term. For case (ii), the symmetric argument using Proposition 7(c)'s $\bar{\xi}^{\star,*} \le \bar{\xi}^{\star,FB}$ makes the floor on $\bar{\xi}^\star$ bind at $\bar{\xi}^{\star,FB}$. For case (iii), imposing both floors makes both FOCs degenerate to the first-best boundary conditions and $(q_c^{**}, \bar{\xi}^{\star,**}) = (q_c^{FB}, \bar{\xi}^{\star,FB})$.

For (28), Taylor-expand $\mathcal{T}(q_c^{FB}, \cdot)$ around $\bar{\xi}^{\star,FB}$ using $\partial_{\bar{\xi}^\star}\mathcal{T}|_{\bar{\xi}^{\star,FB}} = 0$ (the first-best FOC) and $\partial^2_{\bar{\xi}^\star}\mathcal{T} > 0$ at FB by 10. The leading-order term is $\frac{1}{2}\partial^2_{\bar{\xi}^\star}\mathcal{T} \cdot (\bar{\xi}^{\star,**} - \bar{\xi}^{\star,FB})^2$, with the cubic remainder bounded by the C^2 smoothness of $\mathcal{T}$. □

C Bang-bang gate variant

Replace $(\epsilon - \bar{\epsilon}^\star)^+$ in L_W and L_S with $\Lambda \cdot \mathbf{1}\{\epsilon > \bar{\epsilon}^\star\}$, where Λ is a finite jump magnitude. The loss function now has a value jump at $\bar{\epsilon}^\star$ rather than a derivative jump:

$$\bar{L}(\bar{\epsilon}^{\star+};\rho) - \bar{L}(\bar{\epsilon}^{\star-};\rho) = B_W \delta_1 \Lambda + B_S \Lambda G^\star(\rho).$$

The $G^\star(\rho)$ factor is unchanged, so Proposition 2 applies verbatim and the jump is strictly increasing in ρ. The marginal atom-at-zero structure of Lemma 2 is preserved because the atom in the subscription-line marginal comes from the probability of the gate not breaking or of the excess LP default fraction not exceeding q_c; neither depends on the continuous-onset multiplier. Proposition 3(a) therefore holds identically. The Case 2 iff argument in Lemma 4 requires strict monotonicity of η_N and η_F, which is preserved under bang-bang up to an additive constant jump; the displacements Δ_N and Δ_F change magnitude but not sign.

D Numerical verification and reproducibility

D.1 Proposition 2 verification

At $\rho \in \{0.1, 0.5, 0.9\}$ with baseline primitives, direct evaluation of (11) matches numerical differentiation of $G^\star$ to machine precision. Log: `output/stage3a/wedge_numerical_log.txt` lines 10–12.

D.2 Corollary 2 and the 97-point grid

Baseline ($C^\star$) integrals:

$$\text{LHS}(C^\star) = \int_{0.46527}^{0.48901} f(\epsilon)(1 - p(\epsilon))\, d\epsilon = 1.648 \times 10^{-3},$$

$$\text{RHS}(C^\star) = \int_{0.38343}^{0.46527} f(\epsilon)\, p(\epsilon)\, d\epsilon = 6.313 \times 10^{-3},$$

margin ratio 3.83×. Quadrature: `scipy.integrate.quad`, adaptive Gauss-Kronrod, relative tolerance 10^{-8}, subdivision limit 200.

Robustness grid: $q_c \in \{0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50\}$, $\sigma_\epsilon \in \{0.10, 0.12, 0.15, 0.18, 0.20\}$, $\bar{\xi}^\star \in \{0.15, 0.20, 0.25, 0.30\}$, all other primitives at baseline. Of 140 grid points, 97 are valid under 6 and the quantile range of Proposition 3. ($C^\star$) holds at 97 out of 97 valid points. Case-2 minimum margin 3.28× at $(q_c, \sigma_\epsilon, \bar{\xi}^\star) = (0.25, 0.20, 0.20)$. Reproducibility: `code/explore/corollary2_robustness.py`; output: `output/stage3a/corollary2_robustness.json`.

D.3 Proposition 4: heavy-tail probe

Scale-matched Student-t_ν replacement of F, variance matched to $\sigma_\epsilon^2 = 0.04$ (scale $= \sigma_\epsilon\sqrt{(\nu-2)/\nu}$). Table 3.

Extended sweep over $B_S \in [1, 200]$ at each ν: minimum margins are 1.98× at $\nu = 3, B_S \approx 9.1$; 2.48× at $\nu = 5, B_S \approx 5.9$; 2.77× at $\nu = 10, B_S \approx 5.9$. No reversal observed. Reproducibility: `code/explore/iff_certificate_heavytail.py`, `iff_heavytail_extreme.py`.

Table 3: Heavy-tail probe. All other primitives at baseline.

ν	B_S	LHS($C^\star$)	RHS($C^\star$)	Margin ratio	($C^\star$) holds
3	1	1.10×10^{-3}	2.86×10^{-3}	2.60×	Yes
3	10	3.71×10^{-3}	7.35×10^{-3}	1.98×	Yes
5	1	1.44×10^{-3}	4.11×10^{-3}	2.86×	Yes
5	10	4.30×10^{-3}	1.10×10^{-2}	2.56×	Yes
10	1	1.58×10^{-3}	5.16×10^{-3}	3.27×	Yes
10	10	4.89×10^{-3}	1.39×10^{-2}	2.84×	Yes

D.4 Proposition 5 verification

At baseline primitives, analytic $\rho \to 0^+$ derivative equals -0.1476; finite-difference estimate at $\rho = 0.001$ equals -0.1509. At $\rho = 0.9$, finite-difference equals -0.5227; at $\rho = 0.99$, -0.6240. Endpoint signs match the proved directions and the magnitude scales 3–5× as ρ moves from near-zero to near-one. Log: `output/stage3a/wedge_numerical_log.txt` lines 29–33.

D.5 Policy calibration

Baseline missed capital at \$400B notional: $W(1^-;\alpha) \cdot \$400\text{B} = 0.05894 \cdot \$400\text{B} = \$23.58\text{B}$. Incremental correlation-induced missed capital: $(W(1^-;\alpha) - W_0(\alpha)) \cdot \$400\text{B} = \$15.98\text{B}$. Range across the 97-point grid: total missed capital from \$0.5B to \$76B. Reproducibility: `code/explore/policy_calibration_upda`