

Speed Limits for Liquidity: Queues, Gates, and Runs in Semi-Liquid Funds

Sergei Sukhov*

This version: September 2026

Abstract

Semi-liquid funds—non-traded BDCs, interval and tender-offer funds—repurchase shares each quarter at net asset value up to a cap, typically 5% of shares, and ration excess requests pro rata. We develop an equilibrium model of this mechanism.

A cap works as a speed limit. If the drain it permits, $\kappa[\theta + \Lambda(1 + \theta)]$, does not exceed the fund's excess return plus costless liquidity, $\pi + \Lambda\mu$, no investor base can run the fund.

Above this floor, runs are sustained by the backlog of liquidity-constrained investors they trap in the queue. We characterise the run region in closed form, show that a backlog above a tipping level eliminates the no-run equilibrium, and show that tighter caps can be more fragile.

We then characterise all run-proof designs. Among work-conserving anonymous rules, a fund is robust to every investor base if and only if its cap satisfies the speed limit—most cheaply implemented through a fee we show to be minimal—or all liquidity needs are verified and prioritised and requests are withdrawable. Irrevocable, first-in-first-out and last-in-first-out queues are more fragile than withdrawable pro rata, and leverage lowers the floor.

Tender-offer filings for 46 funds show that proration, absent in private-credit BDCs before 2026, reached 23% of their offers in 2026 and is highly persistent.

AI disclosure: This paper was prepared with substantial assistance from a large language model (Claude, Anthropic). The author's and the AI system's roles are described in the Declarations section. The author takes full responsibility for the content.

JEL: G01, G23, G28. **Keywords:** semi-liquid funds, private credit, BDC, redemption gates, proration, runs, liquidity management tools.

1 Introduction

Semi-liquid vehicles have become the main channel through which individual investors hold private credit, private equity and real estate. Evergreen private-credit assets alone exceed half a trillion dollars. Their liquidity promise is modest and precise. Each quarter the fund offers to repurchase shares at NAV, but only up to a cap (usually 5% of shares outstanding). If requests exceed the cap, every request is filled pro rata, and the unfilled part either lapses and can be resubmitted, or is carried forward to the next window. In 2026 this promise was tested for the first time at scale. In our data, requests at the largest funds reached 10–18% of shares per quarter against 5% caps. Fills

*Market Microstructure Research Lab (MMRL), Yerevan, Armenia. E-mail: s.sukhov@mmrls.com. Replication package (code, data, and an automated test suite that reproduces every number, table and figure): [GitHub repository](#). The use of generative AI in this work is described in the Declarations section.

fell to 30–50% of requested shares, and several funds prorated in consecutive quarters (Table 10). Fang et al. (2026) document that 5% caps have not eliminated run-like dynamics in this sector.

Two policy views coexist. One holds that a quantity cap makes a run impossible: nobody can get out faster than 5% per quarter, so there is no first-mover advantage to race for. The other holds that caps are what trigger runs, because investors rush to get in line before the gate binds. Existing theory cannot settle the question because it does not model the mechanism in use. Gates in the literature are suspensions of convertibility (Engineer, 1989; Cipriani et al., 2014; Voellmy, 2021, 2024; Huang and Keister, 2025; Kiss et al., 2025). The swing-pricing literature has no quantity limit (Capponi et al., 2020). Dynamic fund-run models concern daily open-end funds (Zeng, 2017). Stress-test models of BDCs treat redemptions as exogenous (Chernenko and Scharfstein, 2025). None contains the defining feature of semi-liquid funds: *a quantity cap with pro-rata rationing and rollover of unfilled requests*. Concurrent empirical work by Woo (2026) sketches a two-period queue model with an exogenous gap between NAV and market value. In it, exits do not change the value of the claims that remain. We show that this feedback—dilution at a stale NAV and fire-sale costs—is what determines whether a cap stops a run.

What we do. Investors in our model hold claims on a fund whose illiquid assets earn an excess return π over their outside option. A fraction μ of assets repays each period at no cost; further cash is raised at a fire-sale haircut λ . Reported NAV may overstate true value by a factor θ . Each investor is hit by a liquidity shock with hazard ϕ and must then exit. A fraction δ of investors is *attentive* and chooses strategically when to request repurchase. The fund fills at most a fraction κ of shares per period, pro rata.

Results.

1. **Speed-limit floor** (Theorem 1). Let κ_1 solve $g(\kappa_1) = 0$, where $g(q)$ is the growth rate of the exit price when a fraction q is filled. If $\kappa \leq \kappa_1$ there is no run equilibrium for *any* δ . In continuous time $\kappa_1 = (\pi + \Lambda\mu)/[\theta + \Lambda(1 + \theta)]$. This test uses no information about the investor base, which regulators cannot observe. Table 1 translates it into the excess return a 5% cap requires.
2. **Backlog channel** (Theorem 2). Above the floor, rescission (unfilled requests lapse) makes strategic tendering unravel from the end: the last strategic investors in line prefer to walk away and keep the franchise. A run can therefore survive only if, when strategic tendering stops, enough liquidity-shocked investors remain in the queue that the decline while they are served makes staying worse than leaving. We solve the model in closed form. The run region is $\max_i s(i; \delta, \kappa) \geq 1 - x^*(\kappa)$, where s is the backlog the run builds and x^* the backlog it needs. The run’s duration is explicit, and a scaling law reduces the run region to three dimensionless numbers. Tighter caps turn out to be *more* fragile through almost the whole corridor above κ_1 .
3. **Lump channel** (Proposition 2, Section 5). With quarterly NAV strikes, all fills in a window share one stale price, and a single window’s drain can sustain a run without any backlog. This channel vanishes as the window shrinks. We prove the exact one-window condition for lax caps and verify numerically that the quarterly model converges to the continuous-time boundary.
4. **Queue design** (Theorem 3, Proposition 3). (i) Priority for verifiable hardship redemptions eliminates the backlog channel. (ii) Irrevocable carry-forward removes the unravelling and lowers the run threshold δ^* in every case we compute, by 11–80% (median 42%). (iii) First-in-first-out (FIFO) ordering is more fragile than pro rata. This formalises the industry’s historical move

away from FIFO (PREA, n.d.) and informs current regulatory choices, such as ESMA’s question on whether investors may cancel the unexecuted part of gated orders.

5. **Optimal robust design** (Theorem 4, Section 7). Among work-conserving, Markov, anonymous rules, a fund is run-proof for every investor base *if and only if* its cap satisfies the speed limit, or all liquidity needs are verifiable, prioritised and requests withdrawable. With imperfect verification, the fee $f^\dagger$ is the minimal fee that achieves robustness at a given cap (Corollary 2), and each additional share of verified shocks raises the run threshold (Proposition 4). History-dependent rules do not escape the frontier. Recency rules collapse to pro rata when requests can be withdrawn (Proposition 5). With irrevocable requests, LIFO is the best ordering but remains far more fragile than withdrawable pro rata (Proposition 6). Withdrawability is the first-order design choice.
6. **Tipping point and welfare** (Theorem 5, Section 8). A backlog of forced sellers above $1 - x^*$ destroys the no-run equilibrium. Below it, the backlog lowers the run threshold. The threshold is 8–21% of shares for thin franchises, which is comparable to the 9–13% unfilled requests at the most stressed funds in mid-2026. In a simulated wave of liquidity needs, a cap at the floor minimises deadweight fire-sale losses unless waiting is extremely costly; a triggered fee is the efficient complement to looser caps.
7. **Leverage** (Proposition 7, Section 9). With a credit line, the floor is set by the constrained regime, where each dollar paid out forces deleveraging. At equal franchise a 1x line lowers the floor from 4.25% to 3.5% and the run threshold at 5% from 0.82 to 0.57. Runs start while the line is still being drawn and prices are still rising: a *leverage clock*.
8. **Evidence** (Section 10). From 271 SEC tender-offer filings we build a fund-quarter panel of requests, caps and fills. Proration was absent before 2026 and hit 23% of offers in 2026. It is highly persistent: $\Pr(\text{prorate}_{t+1} \mid \text{prorate}_t) = 0.80$ against 0.03 otherwise. Requests relative to the cap rise from 0.94 in the quarter before the first proration to 1.85 and then 2.23.

Why it matters. For managers and boards the message is concrete. A cap is safe for any investor base only if it lies below κ_1 , and κ_1 falls sharply as marks become stale (θ) and secondary discounts widen (λ). Above the floor, *how* the queue is run matters as much as how wide it is. For regulators the results rank liquidity-management designs by robustness, and they identify a cheap tool—hardship priority—that removes the self-fulfilling component of a queue. Stale NAVs matter twice. They lower the floor through θ , and they create the lump channel through discrete NAV strikes.

Related literature. We contribute to four literatures.

Bank runs, suspension and payout restrictions. Starting from Diamond and Dybvig (1983), a large literature studies how contracts and restrictions on withdrawals affect run equilibria: suspension of convertibility (Engineer, 1989; Wallace, 1988), efficient allocations under sequential service (Green and Lin, 2003; Peck and Shell, 2003; Ennis and Keister, 2009, 2010), and fundamentals-driven crises (Allen and Gale, 1998; Chari and Jagannathan, 1988). More recent work designs fees and gates for funds (Cipriani et al., 2014; Voellmy, 2021, 2024; Huang and Keister, 2025) and settlement rules for failing intermediaries (Matta and Perotti, 2021); Kiss et al. (2025) test fees and gates in the laboratory. In all of these a gate is a suspension or a payout schedule. Our gate is a *rate*: it never stops redemptions, it rations them and rolls the remainder over. This produces a mechanism

that suspension models cannot generate—runs sustained by the backlog they create—and design questions they cannot pose: rescission, FIFO, hardship priority.

Fund fragility. Strategic complementarities among fund investors are documented by Chen et al. (2010) and Goldstein et al. (2017), and in MMF runs by Schmidt et al. (2016), Li et al. (2021) and Cipriani and La Spada (2021). Li et al. (2021) show that the prospect of redemption restrictions itself triggered pre-emptive MMF outflows in 2020. Liquidity management and dilution are studied by Zeng (2017), Morris et al. (2017), Chernenko and Sunderam (2016), Ma et al. (2022) and Falato et al. (2021). Swing pricing is analysed by Capponi et al. (2020, 2023); Lewrick and Schanz (2017); Jin et al. (2022). Our dilution parameter θ connects to stale NAV pricing in open-end funds (Chalmers et al., 2001; Greene and Hodges, 2002; Zitzewitz, 2003) and to return smoothing in illiquid funds (Getmansky et al., 2004). Closest in subject are Fang et al. (2026), who document fragility in semi-liquid BDCs empirically, and Woo (2026), who hand-collects tender-offer filings for ten non-traded BDCs. Woo finds a discrete jump in requests when the first proratations are announced and no repricing of the funds’ exchange-listed twins.

Woo argues that the cap severs the payoff externality of open-end funds, because no investor can exit faster than the cap. Theorem 1 makes this precise and qualifies it: the cap rate-limits the externality, and removes it only when the cap lies below the floor κ_1 .

His appendix model is static in values: the wedge is exogenous and the current fill rate enters the tender condition. In our dynamic model an unfilled share remains invested and can be resubmitted, so the fill rate cancels (Lemma 1). Rationing then affects requests only through the expected path of values—the channel behind the backlog and tipping results.

Share restrictions. Lockups, notice periods and gates in hedge funds are studied as sources of illiquidity premia and arbitrage capacity (Aragon, 2007; Ang and Bollen, 2010; Hombert and Thesmar, 2014; Teo, 2011). These papers treat restrictions as given and study pricing or performance; we study the strategic interaction the restrictions induce.

Dynamic runs, fire sales and private credit. Our continuous-time structure, with a queue that drains at a known rate, relates to dynamic debt runs (He and Xiong, 2012; Cheng and Milbradt, 2012; He and Manela, 2016) and to predatory or pre-emptive trading (Bernardo and Welch, 2004; Brunnermeier and Pedersen, 2005). Fire-sale costs follow (Shleifer and Vishny, 1992, 2011; Coval and Stafford, 2007; Chernenko and Sunderam, 2020). Equilibrium selection by global games (Carlsson and van Damme, 1993; Morris and Shin, 1998; Frankel et al., 2003; Goldstein and Pauzner, 2005) does not apply off the shelf, and we show why (Section 11). On private credit and BDCs, we build on Chernenko et al. (2022), Block et al. (2024) and Chernenko and Scharfstein (2025); the latter’s stress tests motivate our leverage section. Roncalli (2021) describes gated redemption schedules with deferral in a stress-testing context, without an equilibrium analysis.

2 Model

Fund. Time is divided into repurchase windows $t = 0, 1, \dots$ of length Δ (a quarter in the baseline). Values are measured relative to the investor’s outside option, so cash paid out does not grow. True NAV per share is a_t . Reported NAV is $(1 + \theta)a_t$ with $\theta \geq 0$ (stale or optimistic marks). The exit price is $P_t = p a_t$ with $p = (1 - f)(1 + \theta)$, where f is a repurchase fee that stays in the fund. Illiquid assets grow at gross rate $G = 1 + \pi$ per window. A fraction μ of assets repays at par each window, and any further cash is raised by selling assets at haircut λ . Write $\Lambda = \lambda/(1 - \lambda)$ and $\zeta = p(1 + \Lambda) - 1$, the value drained from remaining holders per unit of shares filled.

Policy. At most κS_t of the S_t outstanding shares are filled per window. If requests exceed the cap, they are filled pro rata. The baseline rule is *rescission*: the unfilled part of a request lapses, and the holder decides again next window. This is economically identical to cancel-and-resubmit plans, and to carry-forward plans in which requests can be withdrawn. Section 6 studies alternatives.

Investors. There is a continuum of investors. Each is hit by a liquidity shock with hazard ϕ ; shocked investors request repurchase every window until filled. A fraction δ is attentive: while unshocked, they choose each window whether to request. The rest request only when shocked.

Law of motion. If a fraction q of shares is filled at price multiple p ,

$$a_{t+1} = \rho(q; p) a_t, \quad \rho(q; p) = \frac{G[1 - qp - \Lambda(qp - \mu)^+]}{1 - q}. \quad (1)$$

For $p \geq 1$, ρ is decreasing in q . In normal times fills equal ϕ and $\rho_n = \rho(\phi; 1 + \theta)$. An unshocked holder who never requests strategically is worth $V_n = \mathbb{E}[\rho_n^\tau] = \phi \rho_n / (1 - (1 - \phi)\rho_n)$ times the current exit price, where $\tau \sim \text{Geom}(\phi)$. We assume $\phi < \kappa$ (the cap does not bind in normal times) and $V_n \geq 1$ (staying is worthwhile in normal times; otherwise exit is fundamental rather than a run).

Equilibrium. A symmetric equilibrium is a request path $\sigma_t \in \{0, 1\}$ for attentive unshocked investors that is a best response to the price and fill path it generates. A *run* is an equilibrium with $\sigma_t = 1$ for some t . The no-run equilibrium ($\sigma \equiv 0$) always exists when $V_n \geq 1$.

Lemma 1 (Pro-rata neutrality). *Let $x_t \in (0, 1]$ be the window- t fill ratio and W_t the value of a share carried unfilled into window $t + 1$ under optimal continuation. Requesting is worth $x_t P_t + (1 - x_t)W_t$ and not requesting is worth W_t . Hence requesting is optimal iff $P_t \geq W_t$, whatever x_t . W_t is a convex combination of future exit prices.*

The cap never deters requests directly. It works only through the price path it induces.

3 The speed-limit floor

Theorem 1 (Floor). *Let κ_1 be the root of $\rho(\kappa; 1 + \theta) = 1$:*

$$\kappa_1 = \frac{G(1 + \Lambda\mu) - 1}{G(1 + \theta)(1 + \Lambda) - 1} \quad (\text{fire-sale branch}), \quad \kappa_1 \approx \frac{\pi + \Lambda\mu}{\pi + \theta + \Lambda(1 + \theta)}.$$

(i) $V_n \geq 1 \iff \phi \leq \kappa_1$. (ii) If $\phi < \kappa \leq \kappa_1$, no run equilibrium exists for any $\delta \in [0, 1]$.

Proof. The map $h(q) = 1 - qp - \Lambda(qp - \mu)^+ - (1 - q)/G$ is concave with $h(0) = 1 - 1/G > 0$. Hence $\{q \in [0, 1] : \rho(q) \geq 1\} = [0, \kappa_1]$, which gives (i) because $V_n \geq 1 \iff \rho_n \geq 1$. For (ii), under any request profile the fill rate satisfies $q_t \leq \kappa \leq \kappa_1$, so $\rho(q_t) \geq 1$, with strict inequality for $q_t < \kappa_1$. The exit price is therefore non-decreasing along every path, and strictly increasing whenever $q_t < \kappa_1$. A holder who does not request exits almost surely at a later price at least as high. So $W_t \geq P_t$, with strict inequality except in the knife-edge case $\kappa = \kappa_1$, and by Lemma 1 requesting is not strictly optimal. $\square$

Table 1: Speed limit: excess return per quarter (%) needed for a 5% quarterly cap to be run-proof for every investor base, $\pi \geq 0.05\zeta - \Lambda\mu$.

θ	$\lambda = 5\%$		$\lambda = 8\%$		$\lambda = 12\%$	
	$\mu = 0.5\%$	1.5%	0.5%	1.5%	0.5%	1.5%
0%	0.24	0.18	0.39	0.30	0.61	0.48
1%	0.29	0.24	0.45	0.36	0.67	0.53
3%	0.39	0.34	0.55	0.47	0.78	0.65
5%	0.50	0.45	0.66	0.58	0.90	0.76

A cap below κ_1 is a *speed limit*: it caps the rate at which exits can drain value below the rate at which the portfolio replenishes it. This is the only region in which run-proofness does not depend on the investor base. In continuous time (Section 4) the floor is exactly $\kappa_1 = (\pi + \Lambda\mu)/\zeta$. The run-proof condition is therefore

$$\underbrace{\kappa [\theta + \Lambda(1 + \theta)]}_{\text{drain the cap permits}} \leq \underbrace{\pi + \Lambda\mu}_{\text{replenishment}}. \quad (2)$$

Table 1 gives the excess return per quarter needed for a 5% cap to pass this test. With accurate marks and moderate discounts, 0.2–0.4% per quarter suffices. With 3% overstated marks and an 8% secondary discount it takes 0.5–0.55%, and with 5% overstatement and 12% discounts close to 0.9%. Stress raises θ and λ and lowers π and μ , which pushes funds above the floor exactly when it matters.

Fees. A fee charged when requests exceed κ_1 restores $\rho(\kappa; p) \geq 1$ if $f \geq f^\dagger(\kappa) = 1 - y(\kappa)/[\kappa(1 + \theta)]$, where $y = [1 + \Lambda\mu - (1 - \kappa)/G]/(1 + \Lambda)$. This fee is run-proof for every δ by the proof of Theorem 1. At $\kappa = 1$ it equals $1 - [1 - \lambda(1 - \mu)]/(1 + \theta)$, which is full-internalisation swing pricing, so gates and fees lie on one frontier with swing pricing at its end. The exact fee that deters runs as $\delta \rightarrow 1$ coincides with $f^\dagger$ near the floor (0.538% vs 0.539% at $\kappa = 4.5\%$; 1.572% vs 1.588% at 5%, $\pi = 0.40\%$) and lies below it at looser caps (Figure 1 in the Appendix). Near the floor $f^\dagger$ is a small difference of order- κ terms, so its level depends on timing conventions at order $\pi/[(1 + \Lambda)(1 + \theta)] \approx 0.36$ percentage points.

4 Continuous time: the backlog channel

Set-up. Let windows shrink to zero with all rates fixed. With fill rate $q_t \in [0, \kappa]$ (share of shares per quarter), the exit price grows at

$$\dot{a}_t/a_t = g(q_t), \quad g(q) = \pi + q(1 - p) - \Lambda(pq - \mu)^+,$$

which is concave, piecewise linear and decreasing, and equals $\pi + \Lambda\mu - \zeta q$ on the fire-sale branch. Then $\kappa_1 = (\pi + \Lambda\mu)/\zeta$ and $V_n = \phi/(\phi - g(\phi))$, so again $V_n \geq 1 \iff \phi \leq \kappa_1$. If the requesting mass is positive, the cap binds: shares fall at rate κ and each requesting share is filled at hazard $h_t = \kappa/r_t$, where r_t is the queue as a fraction of shares. Otherwise shocked investors are filled on arrival. For $\kappa > \kappa_1$ write $\gamma = -g(\kappa) > 0$ for the rate at which the price falls while the cap binds, and $\beta = \kappa - \phi > 0$.




Figure 1: Minimum run-detering fee at $\delta \rightarrow 1$ (solver) against $f^\dagger(\kappa)$.

HJB form of Lemma 1. Let $v = V/P$ and $u = U/P$ be the values of unshocked and shocked holders relative to the price. Then

$$\dot{v} = -\sigma h(1 - v) - \phi(u - v) + \gamma v, \quad \dot{u} = -h(1 - u) + \gamma u,$$

and requesting ($\sigma = 1$) is optimal iff $v \leq 1$, independent of h .

Lemma 2 (Queue kinematics). *While the cap binds, the unshocked non-requesting fraction i_t grows as $\dot{i} = \beta i$, and the queue is $r = 1 - i$. A queue that starts with non-requesting fraction i_0 therefore clears after exactly $\ln(1/i_0)/\beta$. If all attentive investors request from $t = 0$ with no initial backlog, then $i_t = (1 - \delta)e^{\beta t}$, and the attentive still in line are*

$$a(i) = \delta \left(\frac{i}{1 - \delta} \right)^{-\phi/\beta} \left(\frac{1 - i}{\delta} \right)^{\kappa/\beta}, \quad \text{so the shocked backlog is } s(i) = 1 - i - a(i).$$

Proof. Normalising by shares, which fall at rate κ , gives $\dot{i} = -\phi i + \kappa i$. Attentive investors in line satisfy $\dot{a} = a(\kappa - \phi - \kappa/(1 - i))$. Integrate using $\int \kappa/(1 - ce^{\beta t}) dt = \kappa t - (\kappa/\beta) \ln(1 - ce^{\beta t})$ and substitute $e^{-\phi t} = (i/(1 - \delta))^{-\phi/\beta}$. $\square$

Lemma 3 (Backlog phase). *Suppose strategic requests have stopped and a fraction x of shares is unshocked and outside the queue, so only shocked investors are in line until x reaches 1. Let $A = (\kappa + \gamma)/\beta$, $B = \kappa/\beta$, $C = (\phi + \gamma)/\beta$. Then the value (relative to price) of a shocked share in line is*

$$u(x) = 1 - \frac{\gamma}{\beta} (1 - x)^{-B} \int_x^1 \left(\frac{x}{s} \right)^A (1 - s)^B \frac{ds}{s} \in (0, 1),$$

and the value of an unshocked share outside the queue is

$$W(x) = x^C V_n + \frac{\phi}{\beta} \int_x^1 u(s) \left(\frac{x}{s} \right)^C \frac{ds}{s}.$$

$W(1) = V_n > 1$, $W(0^+) \leq \phi/(\phi + \gamma) < 1$, and at any x with $W(x) = 1$, $W'(x) = [(\phi + \gamma) - \phi u(x)]/(\beta x) > 0$. Hence W crosses 1 exactly once, at a unique $x^*(\kappa) \in (0, 1)$.

Proof. Change variables from t to x ($dt = dx/(\beta x)$) in the linear ODEs of the HJB with $h = \kappa/(1-x)$ and $\sigma = 0$. For $y = 1 - u$, $y' - y[\kappa/(\beta x(1-x)) + \gamma/(\beta x)] = -\gamma/(\beta x)$. The integrating factor $x^{-A}(1-x)^B$ vanishes at $x = 1$, where $y = 0$ (the queue clears), which gives the formula; $y > 0$ because $\gamma > 0$. For W , $W' = [(\phi + \gamma)W - \phi u]/(\beta x)$ with $W(1) = V_n$; integrate with factor x^{-C} . The bound on $W(0^+)$ follows from $u \leq 1$. The slope at a crossing uses $u < 1$, and a function that can cross level 1 only upward crosses it once. $\square$

Lemma 4 (Backlog envelope). *Consider any request behaviour of attentive investors: any timing, symmetric or not, pure or mixed. Let I_t be the inattentive fraction of shares and s_t the backlog of shocked investors in line, starting from (I_0, s_0) . While the cap binds, $\dot{I} = \beta I$, and*

$$s_t \leq \bar{s}(I_t), \quad \frac{d\bar{s}}{dI} = \frac{\phi(1-\bar{s}) - \kappa \bar{s} I/(1-I)}{\beta I}, \quad \bar{s}(I_0) = s_0,$$

with equality if and only if all attentive investors are in line. With $s_0 = 0$ and $I_0 = 1 - \delta$, $\bar{s}$ is the function $s(i)$ of Lemma 2, and $\max \bar{s}$ is strictly increasing in δ .

Proof. The queue consists of the shocked backlog and the attentive investors in line, so $r_t \leq 1 - I_t$. Normalising by shares, $\dot{s} = \phi(1-s) - \kappa s(1/r - 1) \leq \phi(1-s) - \kappa s I/(1-I)$. Inattentive investors are never in line, so $\dot{I} = \beta I$ while the cap binds, and fractions do not move when it does not bind. In the variable I the right-hand side is decreasing in s , so the standard comparison argument gives $s \leq \bar{s}$: if $s - \bar{s}$ became positive, its derivative would be negative. For monotonicity, a larger δ starts the same ODE earlier (smaller I_0) at $\bar{s} = 0$, so its solution lies above the solution for a smaller δ wherever both are defined. $\square$

Theorem 2 (Backlog channel: exact run region). *Let $\phi < \kappa$ and $\kappa > \kappa_1$. There is an equilibrium in which attentive investors request on $[0, \tau)$ and hold afterwards if and only if*

$$\max_{i \in [1-\delta, 1)} s(i) \geq 1 - x^*(\kappa).$$

The switch point solves $s(i_\tau) = 1 - x^$ with $\tau = \ln(i_\tau/(1-\delta))/\beta$, and the cap binds until $T = \tau + \ln(1/x^*)/\beta$. No other kind of run equilibrium exists: the condition is necessary for every equilibrium with strategic requests, whatever its timing and whether or not it is symmetric. Generically, when δ exceeds the threshold $\delta^*(\kappa)$ there are two single-interval runs (short and long) besides the no-run equilibrium.*

Proof. Sufficiency. After τ , attentive investors are outside the queue with ratio $w(\tau) = W(x^*) = 1$. Forward in time, whenever $w = 1$ we have $\dot{w} = \phi(1-u) + \gamma > 0$, so $w \geq 1$ on $[\tau, T]$ and holding is optimal. Before τ , $v(\tau) = 1$ and $\dot{v}|_{v=1} = \phi(1-u) + \gamma > 0$, so going backward v drops below 1 and cannot return to it: requesting is optimal on $[0, \tau)$. *Necessity, for any equilibrium with strategic requests.* Let τ be the supremum of the times at which a positive mass requests strategically.

- τ is finite. Strategic requests make the cap bind, and I grows at rate β while it binds, so total binding time is at most $\ln(1/(1-\delta))/\beta$. As $t \rightarrow \infty$ the remaining binding time goes to zero, so $v_t \rightarrow V_n > 1$, which rules out requests at arbitrarily late dates.
- Indifference at τ . After τ only shocked investors are in line, so the continuation is the backlog phase from $x = 1 - s_\tau$. Continuity of the value function forces $W(1 - s_\tau) = 1$, i.e. $s_\tau = 1 - x^*$.
- Conclusion. By Lemma 4, $s_\tau \leq \max \bar{s} = \max_i s(i)$.

The same argument, started from (I_0, s_0) , covers an initial backlog (Theorem 5). $\square$




Figure 2: Policy map. By Proposition 1, δ^* depends only on the cap multiple κ/κ_1 and the franchise V_n ($\zeta = 0.12$).

Reading. In continuous time strategic requests alone cannot sustain a run. Near the end, the few strategic investors still in line would rather step out and keep the franchise, so the run unravels unless it has trapped enough *liquidity-shocked* investors behind it. Those investors keep the cap binding, and the price falling, after strategic requests stop. The run is self-fulfilling through the investors it traps. This mechanism cannot arise in models where gates are suspensions.

Corollary 1 (Shape of the run region). (i) $\delta^*(\kappa) > 0$ for every κ : continuous time has no “ceiling” above which small attentive shares suffice. (ii) $\delta^*(\kappa) \rightarrow 1$ as $\kappa \downarrow \kappa_1$, because $\gamma \rightarrow 0$ forces $x^* \rightarrow 0$. The convergence is logarithmic: at $\pi = 0.32\%$, $\delta^* = 0.97, 0.79, 0.53$ at $\kappa - \kappa_1 = 10^{-7}, 10^{-4}, 3 \cdot 10^{-3}$. (iii) (heuristic) $\delta^*(\kappa) \rightarrow 1$ as $\kappa \rightarrow \infty$. (iv) (numerical) Outside a thin layer above κ_1 , δ^* increases with κ : tighter caps are more fragile. (v) The attainable backlog is increasing in δ (Lemma 4), so runs exist exactly for $\delta \geq \delta^*(\kappa)$.

Proposition 1 (Scaling law). The continuous-time equilibrium set depends on the primitives only through $(\kappa/\phi, \gamma/\phi, V_n)$. On the fire-sale branch, with ζ fixed, it depends only on the cap multiple $m = \kappa/\kappa_1$ and the franchise V_n .

Proof. Rescaling time by ϕ turns every ODE above into one with rates $\kappa/\phi, \gamma/\phi, \beta/\phi$ and terminal value V_n . With ζ fixed, $\kappa_1/\phi = 1 + (1 - 1/V_n)/\zeta$ and $\gamma = \zeta(\kappa - \kappa_1)$. $\square$

This is why Figure 2 can present the whole run region as a function of two numbers a fund can estimate: how far its cap sits above its floor, and how valuable staying is in normal times.

5 Quarterly windows: the lump channel

Actual funds strike NAV once per window. All fills in that window are paid the same price, while the loss they cause lands on remaining holders the next window. This creates a second, purely discrete source of runs.

Proposition 2 (One-window runs). Let $q_0 = \phi + \delta(1 - \phi)$ and $\kappa \geq q_0$, so the cap does not bind when all attentive investors request at once. The profile “attentive investors request in window 0 only” is an equilibrium if and only if

$$\rho(q_0; 1 + \theta) V_n \leq \rho_n.$$

Table 2: Run threshold δ^* under pro rata with rescission: discrete windows of length Δ (quarters) against the closed form. $\theta = 3\%$, $\lambda = 8\%$, $\mu = 1.5\%$, $\phi = 3\%$ per quarter.

π	κ	$\Delta = 1$	$\Delta = 1/3$	$\Delta = 1/13$	continuous time
0.32%	5.0%	0.511	0.535	0.542	0.543
0.32%	7.5%	0.620	0.671	0.687	0.689
0.32%	10.0%	0.675	0.761	0.784	0.789
0.32%	15.0%	0.663	0.859	0.890	0.897
0.32%	20.0%	0.347	0.901	0.940	0.946
0.32%	30.0%	0.187	0.933	0.979	0.983
0.40%	5.0%	0.816	0.835	0.843	0.843
0.40%	7.5%	0.866	0.886	0.890	0.890
0.40%	10.0%	0.913	0.933	0.940	0.940
0.40%	15.0%	0.952	0.976	0.983	0.981
0.40%	20.0%	0.956	0.991	0.995	0.993
0.40%	30.0%	0.823	0.999	0.999	0.999




Figure 3: Run threshold $\delta^*(\kappa)$: closed form (Theorem 2) against the discrete solver at window lengths of one quarter, one month and one week. The gap at quarterly windows is the lump channel.

Proof. All attentive investors are filled at P_0 . From window 1 onward only inattentive holders remain, fills equal ϕ , and $P_{t+1} = P_t \rho_n$. A deviator who holds exits at the first shock $\tau \geq 1$ and receives $\mathbb{E}[P_0 \rho(q_0) \rho_n^{\tau-1}] = P_0 \rho(q_0) V_n / \rho_n$. Requesting later is never profitable because the price rises from window 1 onward. $\square$

This is the Diamond–Dybvig run of an uncapped NAV fund, and the solver reproduces it exactly (test suite). When $\kappa < q_0$, rationing spreads the run over several windows. A single window of drain can still outweigh the franchise when $\rho(\kappa) V_n < 1$, that is, when κ exceeds the Part-I ceiling $\bar{m}$ (21.5% and 33.5% of shares in our two calibrations). Figure 3 and Table 2 separate the two channels. Refining the window at fixed rates (quarter, month, week) makes δ^* converge to the closed form of Theorem 2. For example, at $\pi = 0.40\%$ and $\kappa = 10\%$, $\delta^* = 0.913, 0.933, 0.940$ against 0.940. The solver’s run timing matches the long root of Theorem 2 to within 0.2 quarters: $(\tau, T) = (52.9, 64.7)$ against $(52.9, 64.9)$. Everything the quarterly model shows beyond the continuous-time curve is the lump channel. It switches on near $\bar{m}$ and disappears as windows shrink.

6 Queue design

We compare four rules actually used by funds. (PR-R) pro rata, requests lapse or can be withdrawn—the baseline. (PR-C) pro rata, unfilled requests carried forward irrevocably. (FIFO-C) first-in-first-out by request date, pro rata within a date, irrevocable. ($H\omega$) PR-R plus priority for a fraction ω of liquidity shocks that are verifiable (death, disability, required minimum distributions).

Theorem 3 (Hardship priority removes the backlog channel). (continuous time) *Under (H1) no run equilibrium exists in continuous time, for any δ and κ .*

Proof. Shocked investors are served before everyone else, and the shock flow is at most $\phi < \kappa$, so the queue never contains shocked investors. Let τ be the supremum of times with strategic requests; τ is finite because the attentive mass decays at rate at least ϕ . After τ the cap does not bind and the price grows at $g(\phi) \geq 0$, so the ratio of an unshocked holder right after τ is $V_n > 1$. By continuity $v > 1$ just before τ , and nobody requests there—a contradiction unless $\tau = 0$. $\square$

Proposition 3 (Irrevocability and FIFO increase fragility). (i) *Under FIFO-C, attentive investors who request at the onset are filled earlier, in the first-order stochastic dominance sense, than under PR-C. Since the price falls while the cap binds, their value is higher.* (ii) (numerical) $\delta_{\text{FIFO-C}}^* < \delta_{\text{PR-C}}^* < \delta_{\text{PR-R}}^*$ in all 48 continuous-time cells on the grid $\phi \in \{2, 3, 6, 10\}\%$, $V_n \in \{1.01, 1.02, 1.05, 1.10\}$, $\kappa/\kappa_1 \in \{1.2, 1.5, 2.5\}$, with no violations, and in quarterly and monthly windows (Table 3).

Proof of (i). The cohort requesting at $t = 0$ is served ahead of all later cohorts under FIFO-C, and receives only a pro-rata share of capacity under PR-C, so its fill-time distribution under FIFO-C is dominated. $\square$

Under irrevocable carry-forward, a strategic investor cannot step out of line when the run is about to end. The unravelling that disciplines runs under rescission disappears, and the committed attentive mass acts as its own backlog. Table 3 and Figure 4 show magnitudes. At a 5% cap and $\pi = 0.32\%$, the run threshold falls from 0.52 (PR-R) to 0.30 (PR-C) and 0.27 (FIFO-C) with quarterly windows. Priority for half of all shocks raises it to 0.64, and priority for all shocks eliminates runs until the lump region. Figure 2 generalises the comparison over $(\kappa/\kappa_1, V_n)$ using Proposition 1.

Enlarging the offer. Several managers raised caps in 2026 (e.g. BCRED to 7% of shares). In the corridor, backlog-channel fragility falls as κ rises (Corollary 1(iv)): at $\pi = 0.32\%$, δ^* moves from 0.52 at 5% to 0.62 at 7.5% with quarterly windows. The gain reverses once κ approaches the lump region near $\bar{m}$.

7 Optimal robust design

A board or regulator cannot observe how many investors are attentive. The natural design objective is therefore robustness: no run equilibrium for *any* $\delta \in (0, 1)$. We characterise all designs with this property.

Table 3: Run threshold δ^* by queue rule. “none”: no run for any $\delta < 1$.

Queue rule	quarterly windows				monthly windows			
	5.0%	7.5%	10.0%	20.0%	5.0%	7.5%	10.0%	20.0%
$\pi = 0.32\%$ per quarter								
Pro-rata, requests lapse	0.52	0.62	0.68	0.35	0.54	0.67	0.76	0.91
Pro-rata, carried forward	0.30	0.29	0.28	0.22	0.31	0.31	0.31	0.30
FIFO, carried forward	0.27	0.27	0.26	0.22	0.28	0.28	0.29	0.29
Pro-rata + priority for 50% of shocks	0.64	0.77	0.83	0.49	0.66	0.81	0.89	0.98
Pro-rata + priority for all shocks	none	none	none	none	none	none	none	none
$\pi = 0.40\%$ per quarter								
Pro-rata, requests lapse	0.82	0.87	0.91	0.96	0.84	0.89	0.94	0.99
Pro-rata, carried forward	0.56	0.50	0.49	0.44	0.59	0.52	0.52	0.52
FIFO, carried forward	0.52	0.46	0.46	0.43	0.55	0.48	0.48	0.49
Pro-rata + priority for 50% of shocks	0.89	0.94	0.98	0.99	0.90	0.95	0.98	1.00
Pro-rata + priority for all shocks	none	none	none	none	none	none	none	none




Figure 4: Queue design. Lines: continuous-time thresholds; markers: discrete solver (o monthly, x quarterly). Blue: pro rata with lapse; orange: pro rata carried forward; red: FIFO; green: priority for 50% of shocks.

Design class. A design is a cap κ , a fee f charged when the cap binds, and an allocation rule. A liquidity shock is *verifiable* with probability ω (death, disability, required distributions), and the rule may give verifiable requests priority. Requests are either *withdrawable* or *irrevocable*. The rule is *work-conserving* (it fills κS whenever requests reach κS) and *Markov and anonymous*: the fill hazard of a request depends only on the current aggregate state and the requester’s observable type.

Lemma 5 (Anonymity implies pro rata). *Under a Markov anonymous work-conserving rule, all unverifiable requests receive the same fill hazard $h = (\kappa - F_v)/r$, where F_v is the capacity used by verifiable requests and r the unverifiable queue.*

Proof. An unverifiable liquidity-shocked investor and a strategic requester have the same observable type, so they receive the same hazard. Work conservation pins down its level. $\square$

With full priority for verifiable requests, the analysis of Section 4 carries over with three changes:




Figure 5: Quarterly model: run equilibria in (κ, δ) for $\pi = 0.32\%$ and 0.40% (1,480 equilibrium solves).

- verifiable shocks are served on arrival, so unverifiable shocked investors in line are filled at hazard $h_\omega = (\kappa - \omega\phi x)/(1 - x)$;
- their value becomes $u_\omega(x) = 1 - \frac{\gamma}{\beta}(1 - x)^{-B_\omega} \int_x^1 (x/s)^A (1 - s)^{B_\omega} ds/s$ with $B_\omega = (\kappa - \omega\phi)/\beta$;
- the holding value becomes $W_\omega(x) = x^C V_n + \frac{\phi}{\beta} \int_x^1 [\omega + (1 - \omega)u_\omega(s)](x/s)^C ds/s$, and the envelope drift becomes $F_\omega(s, I) = (1 - \omega)\phi(1 - s) - (\kappa - \omega\phi(1 - s))s/(1 - I) + \kappa s$.

W_ω crosses 1 once, at x_ω^* . At the crossing $W'_\omega = [(\phi + \gamma) - \phi(\omega + (1 - \omega)u_\omega)]/(\beta x) > 0$ whenever $\omega < 1$, and $W_\omega(0^+) \leq \phi/(\phi + \gamma) < 1$. Theorem 2 holds verbatim with $(\bar{s}_\omega, x_\omega^*)$ in place of $(\bar{s}, x^*)$.

Theorem 4 (Robust design frontier). (continuous time) *Let $\phi < \kappa$. A design is run-proof for every investor base $\delta \in (0, 1)$ if and only if one of the following holds:*

- speed limit: $\kappa \leq \kappa_1(f) = (\pi + \Lambda\mu)/[(1 - f)(1 + \theta)(1 + \Lambda) - 1]$; or
- complete verification: $\omega = 1$, verifiable requests have full priority, and requests are withdrawable.

In particular, if some liquidity needs cannot be verified ($\omega < 1$), the speed limit is the only design robust to every investor base.

Proof. Sufficiency: (a) is Theorem 1; (b) is Theorem 3.

Necessity. Suppose $\kappa > \kappa_1(f)$, so $\gamma > 0$.

Case 1: $\omega < 1$, or verifiable requests do not receive full priority. Pooled verifiable requests are indistinguishable in the queue from unverifiable ones, so this case is covered by the parameter $\omega' < 1$, the share of shocks that is actually prioritised. By Lemma 5 the unverifiable queue is served pro rata. As $\delta \rightarrow 1$, $I_0 \rightarrow 0$ and the inattentive fraction stays near zero for a time of order $\ln(1/I_0)/\beta \rightarrow \infty$. Along the envelope, $\dot{s} \rightarrow \phi(1 - s)(1 - \omega' + \omega's) > 0$, so $\sup \bar{s}_{\omega'} \rightarrow 1$. Since $x_{\omega'}^* > 0$, the condition of Theorem 2 holds for δ close to 1, and a run exists.

Case 2: $\omega = 1$ with full priority, but irrevocable requests. Consider all attentive investors committing at $t = 0$. As $\delta \rightarrow 1$, the committed fraction stays close to 1 and the cap binds for

Table 4: Run threshold δ^* with a share ω of liquidity shocks verified and given priority (continuous time, pro rata with withdrawable requests). “none”: no run for any δ ; “> 0.999”: runs only for δ above 0.999.

ω / κ	5.0%	7.5%	10.0%	15.0%	20.0%
$\pi = 0.32\%$ per quarter					
0.00	0.543	0.689	0.789	0.897	0.946
0.25	0.593	0.749	0.843	0.934	0.970
0.50	0.664	0.823	0.904	0.969	0.989
0.75	0.773	0.915	0.966	0.993	0.999
0.90	0.875	0.973	0.994	0.999	0.999
1.00 (Thm. 3)	none	none	none	none	none
$\pi = 0.40\%$ per quarter					
0.00	0.843	0.890	0.940	0.981	0.993
0.25	0.871	0.920	0.961	0.990	0.997
0.50	0.903	0.951	0.980	0.996	0.999
0.75	0.941	0.980	0.995	0.999	0.999
0.90	0.969	0.995	0.999	0.999	0.999
1.00 (Thm. 3)	none	none	none	none	none

a time that diverges. A committed share is filled at total hazard κ (a shock gives priority, and pro-rata fills add the rest), so its value tends to $\kappa/(\kappa + \gamma)$. A deviator who stays out is filled only when shocked, at hazard ϕ , so her value tends to $\phi/(\phi + \gamma) < \kappa/(\kappa + \gamma)$. Committing later is worse, because the price is falling. By continuity, committing at $t = 0$ is a strict best response for δ near 1. $\square$

Corollary 2 (Minimal robust fee). *If $\omega < 1$, the smallest fee that makes a cap $\kappa > \kappa_1(0)$ robust to every investor base is*

$$f^\dagger(\kappa) = 1 - \frac{\kappa + \pi + \Lambda\mu}{(1 + \Lambda)(1 + \theta)\kappa}.$$

Any lower fee admits a run for some $\delta < 1$.

Proof. $\kappa_1(f)$ is increasing in f , and $\kappa_1(f^\dagger(\kappa)) = \kappa$. Apply Theorem 4. $\square$

Proposition 4 (Second best: partial verification). *Let $\omega < 1$ and $\kappa > \kappa_1(f)$, so some investor base can run the fund. Then: (i) the attainable backlog $\bar{s}_\omega$ is decreasing in ω , because $\partial F_\omega / \partial \omega = \phi(1 - s)[s/(1 - I) - 1] \leq 0$ on the feasible set $s \leq 1 - I$; (ii) (numerical) the tipping backlog $1 - x_\omega^*$ is increasing in ω (no violations on the test grid); (iii) hence δ^* rises with the share of shocks that is verified and prioritised (Table 4). The continuous-time thresholds match the monthly-window solver to within 0.02.*

History-dependent rules. Theorem 4 covers Markov rules. Two history-dependent families matter in practice: seniority (first-in-first-out) and recency (last-in-first-out). Recency is the natural candidate for doing better. Newly shocked investors are always the most recent requesters, so a rule that serves the newest requests first gives them priority *without verification*.

Proposition 5 (Recency collapses under withdrawal). *If requests can be withdrawn and resubmitted at no cost, any rule that weakly favours more recent submissions is outcome-equivalent to pro rata with withdrawable requests. Theorem 4 therefore applies to it.*

Table 5: Run threshold δ^* under history-dependent rules (quarterly windows). The mixed-entry solver and the pure-strategy solver agree exactly on FIFO-C.

Queue rule / κ	5.0%	7.5%	10.0%	15.0%	20.0%
$\pi = 0.32\%$ per quarter					
Pro rata, withdrawable (PR-R)	0.52	0.62	0.68	0.66	0.35
LIFO, irrevocable (mixed entry)	0.33	0.30	0.29	0.25	0.22
Pro rata, irrevocable (PR-C)	0.30	0.29	0.28	0.25	0.22
FIFO, irrevocable (mixed entry)	0.27	0.27	0.26	0.24	0.22
FIFO, irrevocable (pure solver)	0.27	0.27	0.26	0.24	0.22
$\pi = 0.40\%$ per quarter					
Pro rata, withdrawable (PR-R)	0.82	0.87	0.91	0.95	0.96
LIFO, irrevocable (mixed entry)	0.62	0.54	0.52	0.48	0.44
Pro rata, irrevocable (PR-C)	0.56	0.50	0.49	0.47	0.44
FIFO, irrevocable (mixed entry)	0.52	0.46	0.46	0.44	0.43
FIFO, irrevocable (pure solver)	0.52	0.46	0.46	0.44	0.43

Proof. Resubmitting weakly improves a request’s rank and costs nothing, so every live request is resubmitted in every window. All live requests then carry the newest timestamp, and the rule allocates among them as among identical requests: pro rata by Lemma 5. $\square$

Recency can therefore bite only if requests are irrevocable (LIFO-C).

Proposition 6 (LIFO with irrevocable requests). *(i) (continuous time) No equilibrium has a mass of strategic requests at any instant. An investor who requests an instant later is served ahead of the mass, at essentially the same price, whereas members of the mass are rationed while the price falls. Runs can therefore only be gradual and pre-emptive. (ii) (numerical) Such runs exist. In quarterly windows they are equilibria with fractional entry, which we compute with a projected fixed-point solver on the entry share. Validation: the same solver reproduces the pure-strategy FIFO-C thresholds exactly in all ten cells (Table 5). (iii) (numerical) In every cell, $\delta_{FIFO-C}^* < \delta_{PR-C}^* \leq \delta_{LIFO-C}^* < \delta_{PR-R}^*$.*

Two conclusions follow. First, among irrevocable rules, favouring recency helps: LIFO-C is the most robust of them, because it serves newly shocked investors first whenever strategic entry is moderate. Second, no ordering compensates for irrevocability. The gap between withdrawable pro rata and the best irrevocable rule (0.52 against 0.33 at $\kappa = 5\%$, $\pi = 0.32\%$) is an order of magnitude larger than the gaps among irrevocable orderings. The first-order design choice is whether unfilled requests can be withdrawn; the order of service is second order. LIFO also carries a fairness cost that pro rata does not. In a wave in which forced requests exceed capacity, the earliest forced sellers are served last and can wait indefinitely.

Remark 1. For US issuer tender offers, Rule 13e-4(f)(3) requires oversubscribed offers to be accepted pro rata, with limited exceptions for odd lots and conditional tenders. Each quarterly offer is a separate offer, and unfilled shares are returned. Semi-liquid BDCs operating through tender offers therefore already use the most robust rule that does not require verification (withdrawable pro rata). Irrevocable carry-forward and seniority queues arise in share-repurchase programmes of other vehicles and in European open-ended funds, which is where our design results bite.

What the theorem says for policy. Robustness to every investor base can be bought in exactly two ways: a speed limit, or complete verification of liquidity needs with priority and withdrawable

Table 6: Tipping backlog $1 - x^*$ (% of shares): the unfilled queue of forced sellers above which the no-run equilibrium ceases to exist.

κ/κ_1	$V_n = 1.01$	$V_n = 1.02$	$V_n = 1.03$	$V_n = 1.05$	$V_n = 1.08$
1.1	10.0	21.0	31.5	49.2	68.2
1.2	9.1	18.5	27.5	43.3	61.4
1.5	8.4	16.6	24.3	38.0	54.4
2.0	8.2	15.9	23.1	35.8	51.1
3.0	8.1	15.6	22.5	34.6	49.3

requests. Complete verification is not available in practice, because reallocations and adviser-driven exits cannot be verified. The speed limit is therefore the only complete tool, and $f^\dagger$ is its cheapest implementation. Everything else—partial hardship programmes, better investor communication, larger offers, irrevocable queues—moves the threshold δ^* (Proposition 4, Section 6) but leaves some investor base able to run the fund. These tools manage fragility; they do not remove it.

8 Fundamental waves and the tipping point

The runs of Theorem 2 are self-fulfilling: fundamentals alone do not force them. In 2026 the trigger was a wave of genuine liquidity needs and reallocations that pushed requests past the cap and left unfilled requests in line. The model has a sharp answer to what such a wave does.

Theorem 5 (Tipping point). (continuous time) *Suppose that at $t = 0$ a backlog s_0 of investors who must exit is already in line, and $\phi < \kappa$, $\kappa > \kappa_1$. (i) If $s_0 > 1 - x^*(\kappa)$, the no-run equilibrium does not exist, and a run equilibrium exists for every $\delta > 0$. (ii) If $s_0 < 1 - x^*(\kappa)$, the no-run equilibrium exists. A run exists if and only if the envelope of Lemma 4, started at $(I_0, s_0) = ((1 - \delta)(1 - s_0), s_0)$, reaches $1 - x^*$. The threshold $\delta^*(\kappa, s_0)$ is decreasing in s_0 and falls to 0 at $s_0 = 1 - x^*$.*

Proof. (i) If nobody requests strategically, an attentive holder’s value ratio is $W(1 - s_0) < 1$ (Lemma 3), so requesting at $t = 0$ is strictly better, and the no-run profile is not an equilibrium. Along the all-request path the envelope starts above $1 - x^*$ and tends to 0 as $I \rightarrow 1$. By continuity it crosses $1 - x^*$ at some $\tau > 0$, and the sufficiency argument of Theorem 2 applies. (ii) $W(1 - s_0) > 1$, and $\dot{w} > 0$ whenever $w \geq 1$, so holding is optimal at every date along the no-request path. The remaining statements follow from Lemma 4 and the necessity argument of Theorem 2 with initial condition (I_0, s_0) . Monotonicity in s_0 : a higher initial point of the same ODE yields a higher envelope. $\square$

$1 - x^*$ is therefore a *tipping backlog*. Below it, a queue of forced sellers only makes runs easier: at $\kappa = 5\%$ and $\pi = 0.32\%$, δ^* falls from 0.54 to 0.45, 0.39 and 0.31 as s_0 rises to 10%, 15% and 20% of shares. Above it, the run is the only outcome, whatever the investor base. By Proposition 1, the tipping backlog depends only on the cap multiple and the franchise (Table 6). For thin franchises ($V_n = 1.01$ – 1.02) it is 8–21% of shares; for $V_n \geq 1.05$ it exceeds a third.

Mapping to data. The last column of Table 10 reports the unfilled part of requests after proration. In mid-2026 it was 12.4–12.6% of shares at the two funds with the longest proration spells, and 5–9% elsewhere. If those requests came from investors who must exit, and franchises are thin, these funds are at or near the tipping backlog. Observed unfilled requests *overstate* forced demand,

Table 7: Wave of liquidity needs (12% per quarter for four quarters), quarterly model, $\pi = 0.32\%$. Entries: run in the greatest equilibrium / deadweight fire-sale loss (% of initial equity) / waiting of shocked investors (initial-equity-quarters).

Rule / cap	3.5%	5.0%	7.5%	10.0%	15.0%
<i>attentive share $\delta = 0.15$: run? / fire-sale loss (% equity) / waiting (equity-quarters)</i>					
PR-R	no / 5.17 / 4.97	run / 5.80 / 2.12	run / 5.99 / 0.46	run / 6.04 / 0.14	run / 6.06 / 0.00
PR-C	no / 5.17 / 4.97	run / 5.80 / 2.08	run / 5.99 / 0.45	run / 6.04 / 0.14	run / 6.06 / 0.00
Hardship priority	no / 5.17 / 4.97	no / 5.72 / 1.43	run / 5.99 / 0.41	no / 6.04 / 0.12	no / 6.06 / 0.00
PR-R + fee $f^\dagger$	no / 5.17 / 4.97	no / 5.57 / 1.56	no / 5.70 / 0.48	no / 5.78 / 0.15	no / 5.92 / 0.00
<i>attentive share $\delta = 0.4$: run? / fire-sale loss (% equity) / waiting (equity-quarters)</i>					
PR-R	no / 5.17 / 4.97	run / 6.00 / 4.19	run / 6.34 / 1.46	run / 6.48 / 0.60	run / 6.60 / 0.06
PR-C	no / 5.17 / 4.97	run / 6.00 / 3.31	run / 6.36 / 1.22	run / 6.50 / 0.54	run / 6.60 / 0.05
Hardship priority	no / 5.17 / 4.97	no / 5.72 / 1.43	no / 6.04 / 0.40	no / 6.27 / 0.10	run / 6.60 / 0.00
PR-R + fee $f^\dagger$	no / 5.17 / 4.97	no / 5.58 / 1.61	no / 5.72 / 0.49	no / 5.79 / 0.16	no / 5.90 / 0.02

however. Under pro rata, an investor who needs to sell part of her position has an incentive to tender more, up to her whole position, to secure the fill she needs. When needs exceed capacity, reported requests are therefore an upper bound on the forced backlog. The implied distance to tipping is conservative, which is the right direction for a supervisory indicator. (heuristic)

Wave simulations and welfare. Table 7 feeds a four-quarter wave of liquidity needs (hazard 12% per quarter, four times normal) into the quarterly model at $\pi = 0.32\%$ ($\kappa_1 = 3.66\%$). We record, for the greatest equilibrium, whether a run occurs, the deadweight fire-sale loss, and the waiting of liquidity-shocked investors (in initial-equity-quarters). Four findings stand out. (numerical)

1. A cap below the floor (3.5%) never runs, even in the wave, as Theorem 1 implies.
2. At 5% and above, the wave tips the fund even with a small attentive share. At $\delta = 0.15$, where no run exists without the wave, the no-run equilibrium disappears (best-response iteration from “nobody requests” ends in a run).
3. Hardship priority does not help when the wave itself exceeds the cap ($\phi_H > \kappa$): the priority queue then builds its own backlog, and Theorem 3 requires $\phi < \kappa$.
4. The triggered fee $f^\dagger$ prevents runs at every cap.

On welfare, the cap at the floor has the lowest fire-sale loss (5.17% of equity against 5.8–6.7% with runs) but the longest waiting. Relative to a 5% cap with fee $f^\dagger$ (loss 5.57%), it is preferable unless waiting costs exceed about 12% per quarter. Relative to a 5% cap without a fee, which runs, it is preferable for any waiting cost below about 100% per quarter. The efficient menu is a cap at or below κ_1 ; if a looser cap is required for marketing reasons, it should come with a triggered fee.

9 Leverage

Semi-liquid BDCs borrow on credit lines subject to the 1940-Act asset-coverage limit: debt up to 1–2 times equity, typically about 1 in practice. Repurchases are then funded by repayments first, then by drawing the line up to the limit, and only then by sales. Once the limit binds, every dollar paid out also forces debt repayment, and fire sales are amplified.

Table 8: Leverage and fragility (quarterly model, $\pi = 0.40\%$, target leverage $0.85\bar{\ell}$).

Max debt/equity	Debt cost (%/q)	V_n	Floor, CT (%)	Floor, quarterly (%)	δ^* at 5%	δ^* at 7.5%
0.00	0.00	1.058	4.44	4.25	0.82	0.87
0.50	0.00	1.118	4.88	4.75	0.96	0.94
0.50	0.25	1.060	4.11	3.75	0.67	0.77
1.00	0.00	1.191	5.14	5.00	none	0.98
1.00	0.25	1.063	3.93	3.50	0.57	0.68
1.25	0.00	1.234	5.23	5.00	none	0.98
1.25	0.25	1.064	3.86	3.50	0.52	0.63

Proposition 7 (Levered floor). (continuous time; transitions between regimes numerical) *Let $\bar{\ell}$ be the maximum debt-to-equity ratio, $\ell^* \leq \bar{\ell}$ the target, and c the cost of debt over the investor’s outside option. Define $k(\ell) = \pi(1 + \ell) - c\ell$ and $\Lambda_\ell = \lambda(1 + \bar{\ell})/[1 - \lambda(1 + \bar{\ell})]$. If*

$$\kappa \leq \kappa_1^{lev} = \min \left\{ \frac{k(\ell^*)}{p - 1}, \frac{k(\bar{\ell}) [1 + \Lambda_\ell \bar{\ell}/(1 + \bar{\ell})] + \Lambda_\ell \mu}{p(1 + \Lambda_\ell) - 1} \right\},$$

the fund is run-proof for every δ .

Proof. With headroom on the line, fills are debt-financed and the exit-price drift is $k(\ell) + q(1 - p) \geq k(\ell^*) + q(1 - p)$ (for $\pi \geq c$). At the limit, keeping $B = \bar{\ell}E$ requires sales $y = [(1 + \bar{\ell})(qp - \mu) - \bar{\ell}k]/[1 - \lambda(1 + \bar{\ell})]$, and the drift becomes $k[1 + \Lambda_\ell \bar{\ell}/(1 + \bar{\ell})] + \Lambda_\ell \mu - q[p(1 + \Lambda_\ell) - 1]$. Both drifts are non-negative for $q \leq \kappa_1^{lev}$, and the argument of Theorem 1 applies. $\square$

Leverage raises the carry k but multiplies the fire-sale cost from Λ to Λ_ℓ (from 0.087 to 0.19 at $\bar{\ell} = 1$, $\lambda = 8\%$). Table 8 compares funds with the same franchise; a debt cost of 0.25% per quarter keeps $V_n \approx 1.06$ across leverage levels. At equal franchise, a 1x line lowers the floor from 4.25% to 3.5% of shares per quarter, and the run threshold at a 5% cap from 0.82 to 0.57 (quarterly model). Leverage that is free relative to the investor’s alternative raises the franchise so much that it stabilises; the comparison must hold the franchise fixed.

The leverage clock. Figure 6 (right) shows a run equilibrium with a 1x line. For two quarters exits are financed by the line and the exit price *rises*. Strategic requests nevertheless start at once: by Lemma 1, the decision compares today’s price with the queue value, which already reflects the forced deleveraging to come. Runs on levered semi-liquid funds therefore begin before any price decline is visible, and headroom on the credit line is not a measure of safety. (numerical)

10 Evidence from tender-offer filings

Data. Semi-liquid BDCs repurchase shares through issuer tender offers. For each offer they file a Schedule TO-I stating shares outstanding and the offer amount, and a final amendment reporting the shares tendered and accepted. We download every Schedule TO-I filed from 2021 to September 2026 by 36 perpetual non-traded BDCs and by 113 further tender-offer funds identified through EDGAR full-text search—about 3,900 documents—and parse them. Requests are split by share class where necessary, and each final amendment is matched to its offer by expiry date. The result is a panel of 362 fund-quarters for 46 funds: 271 for 29 private-credit BDCs, and the rest for private-equity, real-estate and multi-asset tender-offer funds. Hedge-fund-of-funds vehicles tender




Figure 6: Left: welfare frontier in a wave of liquidity needs (fire-sale deadweight against waiting), by queue rule, with caps from 2.5% to 100% (the 2.5% cap lies below the normal shock rate and never clears); circled points are runs. Right: leverage clock—a run equilibrium with a 1x credit line; strategic requests (shaded) start while the exit price is still rising.

in dollar amounts and are excluded. Parsing rules and a manual check of randomly drawn filings are documented in the replication package. Coverage of some large sponsors is incomplete: filings with non-standard wording are dropped rather than guessed.

Facts. (1) Until 2025 proration never occurred, and at most 17% of offers in any year had requests above the offer. In 2026, 31% of offers were oversubscribed and 23% were prorated. (2) Proration is persistent: $\Pr(\text{prorate}_{t+1} \mid \text{prorate}_t) = 0.80$ against 0.03 otherwise, and six of seven proration spells are ongoing at the end of the sample. (3) Around the first proration ($N = 7$ funds), mean requests relative to the offer go from 0.28 four quarters before to 0.48, 0.94, 1.85 at the event and 2.23 the quarter after, while the fill ratio falls to 57–56% (Figure 7). (4) At the largest funds, requests rose over consecutive windows while the cap stayed fixed (Table 10).

Other semi-liquid funds. Table 9 shows that proration is not new to the sector as a whole. Private-equity and real-estate tender-offer funds were oversubscribed in 24–50% of offers and prorated in 10–22% during 2023–2025, when private-credit BDCs had none. Private credit’s 2026 episode is the arrival of an existing phenomenon in the largest segment.

Repayments. The floor depends on costless liquidity μ . From XBRL 10-K data (99 fund-years, 30 BDCs) the annual proceeds from dispositions and repayments, per quarter and relative to the portfolio, have a median of 2.4% in 2022 and 4.8% in 2024–2025 (interquartile range 2.6–5.9%). These figures include sales and prepayments, so they are an upper bound on μ . With $\pi = 0.3\%$, $\theta = 3\%$ and $\lambda = 8\%$, the floor ranges from about 2.9% (contractual amortisation only, $\mu \approx 0.5\%$) to about 6% (at the 2024–2025 refinancing pace). A slowdown in refinancing alone can move a fund from below to above its floor with an unchanged 5% cap.

Relation to concurrent evidence. Woo (2026) measures requests for the ten largest non-traded BDCs over 2024–2026 from the same filings. The two independently parsed datasets agree share

Table 9: Tender-offer outcomes by fund type and year (share of offers with requests above the offer, and share prorated).

Year	N	Private credit / BDC			N	Private equity, real estate, multi-asset		
		Median req.	Oversubscr.	Prorated		Median req.	Oversubscr.	Prorated
2021	6	0.9%	17%	0%	9	2.8%	56%	0%
2022	18	0.8%	6%	0%	10	0.8%	0%	0%
2023	36	1.4%	6%	0%	10	4.0%	50%	10%
2024	71	0.8%	0%	0%	18	3.3%	44%	22%
2025	92	1.2%	4%	0%	25	1.2%	24%	16%
2026	48	3.7%	31%	23%	19	3.0%	32%	16%

for share wherever they overlap: tendered and accepted counts for ADS, HLEND, GCRED, ASIF, FPCF and BCRED’s prorated offer. Our panel is broader (29 private-credit BDCs over 2021–2026, plus non-credit tender-offer funds), and our event study reproduces his jump at the first prorations.

Two points of reading differ. First, in our data BCRED’s offer that expired in March 2026 was accepted in full after the board raised the cap to 7%. BCRED’s subsequent 43% increase in requests therefore followed accommodation, not proration.

Second, Woo reports that the deeply prorated funds escalated in the next window while the mildly prorated GCRED cleared. Theorem 5 gives this a precise but weak reading. A backlog above the tipping level forces escalation, whereas below it both outcomes are equilibria. GCRED’s queue clearing therefore implies only that its unfilled backlog (3.9% of remaining shares) was below its tipping level. In the model this holds for any franchise $V_n \geq 1.005$ at every cap multiple we compute. The escalation of HLEND (4.5%) and ADS (6.4%) is consistent with the model but does not identify its parameters. Coverage of some large sponsors is incomplete in our panel, notably the Blue Owl funds, whose filings our parser does not read; reported requests at the largest of them reached about 22% of shares in early 2026.

Interpretation. The pattern of rising requests after proration is what the backlog channel implies: unfilled requests come back, and new shocks join a longer line. The four-quarter build-up before the first proration points to deteriorating fundamentals—lower π and higher θ and λ —moving funds towards and across κ_1 , rather than to a pure sunspot. Requests of 10–18% of shares are still far below the 30%-plus of shares that a self-fulfilling run requires in the model at calibrated parameters (Table 3), so the 2026 episodes look like fundamental outflows under rationing that have not (yet) tipped into self-fulfilling runs. Theorem 5 says where the tipping point lies. The unfilled backlogs of 9–13% of shares at the most stressed funds are below the tipping backlog for $V_n \geq 1.03$ but at or above it for $V_n \leq 1.02$. Whether these funds are safe therefore turns on how valuable staying is, which is exactly what a stress in private-credit returns erodes.

11 Uniqueness and global games

Under common knowledge the model has up to three equilibria (Theorem 2). A global-games refinement would pin down one. Two ingredients are present. There are dominance regions: holding is dominant if $\kappa \leq \kappa_1$, and requesting is dominant if $\phi > \kappa_1$, since then $g < 0$ on every path. And there is monotonicity in the state: values fall pointwise in the path of the queue fraction $r(\cdot)$. But the game is *not* supermodular in strategy profiles. (by example) Let σ request on $[0, \infty)$ and σ'

Table 10: Oversubscribed offers since 2025Q4 (parsed from SEC Schedule TO-I/A filings). Fund abbreviations are listed in Table 11.

Fund	Offer expiry	Tendered (% shares)	Offer (% shares)	Filled (% tendered)	Unfilled backlog (% shares)
ADS	2026-03-16	11.05	5.00	45.2	6.4
ADS	2026-06-15	16.80	5.00	29.8	12.4
ASIF	2025-12-19	7.01	6.20	100.0	—
ASIF	2026-03-20	14.54	6.26	43.1	8.8
ASIF	2026-06-18	18.10	6.29	34.7	12.6
BCRED	2026-03-02	6.96	5.00	100.0	—
BCRED	2026-05-29	10.30	5.00	48.5	5.6
BDEBT	2026-06-08	5.25	5.00	95.2	0.3
GCRED	2026-05-01	8.95	5.25	58.7	3.9
HLEND	2026-03-04	9.25	5.00	54.0	4.5
HLEND	2026-06-08	13.30	5.00	37.6	8.7
MCIP	2026-02-27	5.29	5.00	n/a	—
MCIP	2026-05-29	8.89	5.00	56.3	4.1
MCIP	2026-08-31	6.43	5.00	77.7	1.5
NHPIF	2025-12-05	5.25	5.00	n/a	—
NHPIF A	2025-12-05	5.75	5.00	n/a	—
OSCF	2026-03-16	6.82	5.00	100.0	—
VISTA	2026-04-21	11.33	5.62	n/a	—

only from $t_1 > 0$, so $\sigma \geq \sigma'$. Under σ the queue at t_1 is $1 - (1 - \delta)e^{\beta t_1} < \delta$, while under σ' it is δ . Requesting early drains the pool of future requesters. The standard theorems of Goldstein and Pauzner (2005) therefore do not apply as they stand. A formulation with staggered review opportunities in the spirit of He and Xiong (2012), with state (fundamental, queue, attentive stock), is the natural route and is left open. Our policy results do not depend on selection: they concern whether runs *can* occur (Theorems 1, 2, 3) and how the run region shifts across designs.

12 Policy implications

1. **Compute the floor.** Boards can estimate $\kappa_1 = (\pi + \Lambda\mu)/[\theta + \Lambda(1 + \theta)]$ from observable inputs: excess yield net of fees and expected losses (π), repayment rates (μ), secondary-market discounts (λ), and NAV premia over traded twins (θ). A cap below κ_1 is robust to any investor base. Above it, robustness depends on how many investors are attentive (platform- and adviser-intermediated money), which is not disclosed.
2. **Marks matter twice.** θ lowers the floor, and quarterly strikes at stale NAV create the lump channel. Charging an exit fee equal to the estimated overstatement sets $p = 1$ and raises κ_1 to $(\pi + \Lambda\mu)/\Lambda$.
3. **Let unfilled requests lapse or be withdrawn; avoid irrevocable carry-forward, FIFO and LIFO queues.** At identical caps, these choices lower the attentive share needed for a run by 11–80% (median 42%).
4. **Give priority to verifiable hardship.** It removes the self-fulfilling backlog at little cost to other investors, because hardship flows are small relative to the cap.
5. **Enlarging offers can stabilise** within the corridor, but not near the lump region.




Figure 7: Left: tendered shares as % of outstanding across semi-liquid BDCs by quarter, with the median offer. Right: requests relative to the offer around each fund’s first prorated tender.

6. **Monitor the unfilled backlog against the tipping threshold.** Unfilled requests after proration, as a share of shares outstanding, are reported in tender filings and can be compared with $1 - x^*$ (Table 6). The comparison is conservative because requests overstate forced demand.
7. **Treat credit-line headroom as a clock, not a buffer.** Runs on levered funds start while the line is being drawn; the relevant floor is the constrained one (Proposition 7).

13 Robustness, verification and limitations

Every number in this paper is produced by the replication package, which includes 30 automated tests. The tests cover: (i) the floor as an exact root; (ii) no runs below the floor on 1,480 grid cells; (iii) agreement of the reference and accelerated solvers; (iv) convergence of discrete δ^* to the closed form within 0.01 at weekly windows; (v) run timing against Theorem 2; (vi) the one-window benchmark; (vii) the scaling law; (viii) the design ranking on random parameter draws; and (ix) the removal of an earlier floating-point artefact; (x) the backlog envelope and its closed form; (xi) the tipping threshold; (xii) reduction of the leverage engine to the unlevered model; and (xiii) the partial-verification extension: its reduction at $\omega = 0$, its match with the discrete solver, monotonicity of the envelope, $\sup \bar{s} \rightarrow 1$ as $\delta \rightarrow 1$, and minimality of $f^\dagger$; (xiv) the mixed-entry solver: it reproduces pure-strategy FIFO-C, finds LIFO-C runs, and confirms that the no-run equilibrium exists. That artefact produced spurious long runs when shares outstanding fell below 10^{-15} at $\phi = 10\%$; the state is now renormalised every window.

Limitations. The model is in reduced form: π and θ are exogenous, and π —the value of staying relative to the investor’s actual alternative—is hard to measure. At $V_n = 1.05$ the backlog-channel runs last 23–86 quarters. Durations comparable to observed episodes (about 7–16 quarters) arise only when the franchise is thin ($V_n \approx 1.02$) or on the short root. Leverage is modelled through a target and a maximum debt-to-equity ratio; covenants beyond asset coverage, secured-facility borrowing bases and lender behaviour are not. Welfare is measured by deadweight fire-sale losses and waiting, not by a utility function. Inflows, time-varying marks and an explicit model of partial-position requests are left for future work. The empirical panel is descriptive, and funds whose filings use non-standard wording are missing.

A Fund identifiers

Table 11: Private-credit BDCs in the tender-offer panel: abbreviations used in the tables, SEC registrant names and Central Index Keys.

Abbreviation	SEC registrant name	CIK
AB	AB Private Credit Investors Corp	1634452
ADS	Apollo Debt Solutions BDC	1837532
APCF	Antares Private Credit Fund	1976336
ASIF	ARES STRATEGIC INCOME FUND	1918712
BCPC	Bain Capital Private Credit	1899017
BCRED	Blackstone Private Credit Fund	1803498
BDEBT	BlackRock Private Credit Fund	1902649
BPCC	Barings Private Credit Corp	1859919
CPCIC	Crescent Private Credit Income Corp	1954360
FEPCF	First Eagle Private Credit Fund	1890107
FPCF	Fidelity Private Credit Fund	1920453
GCRED	Golub Capital Private Credit Fund	1930087
GSCR	Goldman Sachs Private Credit Corp.	1920145
HLEND	HPS Corporate Lending Fund	1838126
HPSCS	HPS Corporate Capital Solutions Fund	1989817
KFIT	KKR FS Income Trust	1930679
KFITS	KKR FS Income Trust Select	1975736
MCIP	Monroe Capital Income Plus Corp	1742313
NCPIF	Nuveen Churchill Private Capital Income Fund	1911066
NHPIF	North Haven Private Income Fund LLC	1851322
NHPIF A	North Haven Private Income Fund A LLC	1973476
NMPCF	New Mountain Private Credit Fund	2037804
OCREDIT	T. Rowe Price OHA Select Private Credit Fund	1901164
ONEX	Onex Direct Lending BDC Fund	1860424
OSCF	Oaktree Strategic Credit Fund	1872371
SPRIV	Stepstone Private Credit Fund LLC	1950803
STELLUS	Stellus Private Credit BDC	1901037
TCAP	TPG Twin Brook Capital Income Fund	1913724
VISTA	VISTA CREDIT STRATEGIC LENDING CORP.	1919369

Declarations

Funding. This research received no external funding.

Competing interests. The author declares no competing interests.

Data and code availability. All data come from public SEC EDGAR filings (Schedules TO-I and TO-I/A; XBRL company facts). The code, processed data and automated test suite that reproduce every number, table and figure are publicly available in the [GitHub repository](#).

Use of generative AI. This paper was prepared with substantial assistance from a large language model (Claude, Anthropic).

The author defined the research objective and the criteria for the project, selected the problem among candidate topics, and assembled the core literature. The author also set the verification standards used throughout: every reported number is regenerated by code, and the status of every claim is stated. Finally, the author directed the development of the model and the empirical work at each stage, decided on scope, interpretation and policy conclusions, and reviewed the manuscript.

The AI system was used to formalise the model and to derive and check the proofs, to write and test the numerical code, to automate the collection and parsing of SEC filings, and to draft the text.

The AI system is not an author. The author takes full responsibility for the content, including any errors.

References

- Allen, F., Gale, D. (1998). Optimal financial crises. *Journal of Finance* 53(4), 1245–1284. [doi:10.1111/0022-1082.00052](https://doi.org/10.1111/0022-1082.00052)
- Ang, A., Bollen, N. (2010). Locked up by a lockup: valuing liquidity as a real option. *Financial Management* 39(3), 1069–1096. [doi:10.1111/j.1755-053x.2010.01104.x](https://doi.org/10.1111/j.1755-053x.2010.01104.x)
- Aragon, G. (2007). Share restrictions and asset pricing: evidence from the hedge fund industry. *Journal of Financial Economics* 83(1), 33–58. [doi:10.1016/j.jfineco.2005.11.001](https://doi.org/10.1016/j.jfineco.2005.11.001)
- Bernardo, A., Welch, I. (2004). Liquidity and financial market runs. *Quarterly Journal of Economics* 119(1), 135–158. [doi:10.1162/003355304772839542](https://doi.org/10.1162/003355304772839542)
- Block, J., Jang, Y. S., Kaplan, S., Schulze, A. (2024). A survey of private debt funds. *Review of Corporate Finance Studies* 13(2), 335–383. [doi:10.1093/rcfs/cfae001](https://doi.org/10.1093/rcfs/cfae001)
- Brunnermeier, M., Pedersen, L. (2005). Predatory trading. *Journal of Finance* 60(4), 1825–1863. [doi:10.1111/j.1540-6261.2005.00781.x](https://doi.org/10.1111/j.1540-6261.2005.00781.x)
- Capponi, A., Glasserman, P., Weber, M. (2020). Swing pricing for mutual funds: breaking the feedback loop between fire sales and fund redemptions. *Management Science* 66(8), 3581–3602. [doi:10.1287/mnsc.2019.3353](https://doi.org/10.1287/mnsc.2019.3353)
- Capponi, A., Glasserman, P., Weber, M. (2023). Swing pricing: theory and evidence. *Annual Review of Financial Economics* 15, 617–640. [doi:10.1146/annurev-financial-110921-100843](https://doi.org/10.1146/annurev-financial-110921-100843)
- Carlsson, H., van Damme, E. (1993). Global games and equilibrium selection. *Econometrica* 61(5), 989–1018. [doi:10.2307/2951491](https://doi.org/10.2307/2951491)
- Chalmers, J., Edelen, R., Kadlec, G. (2001). On the perils of financial intermediaries setting security prices: the mutual fund wild card option. *Journal of Finance* 56(6), 2209–2236. [doi:10.1111/0022-1082.00403](https://doi.org/10.1111/0022-1082.00403)
- Chari, V. V., Jagannathan, R. (1988). Banking panics, information, and rational expectations equilibrium. *Journal of Finance* 43(3), 749–761. [doi:10.1111/j.1540-6261.1988.tb04606.x](https://doi.org/10.1111/j.1540-6261.1988.tb04606.x)
- Chen, Q., Goldstein, I., Jiang, W. (2010). Payoff complementarities and financial fragility: evidence from mutual fund outflows. *Journal of Financial Economics* 97(2), 239–262. [doi:10.1016/j.jfineco.2010.03.016](https://doi.org/10.1016/j.jfineco.2010.03.016)
- Cheng, I.-H., Milbradt, K. (2012). The hazards of debt: rollover freezes, incentives, and bailouts. *Review of Financial Studies* 25(4), 1070–1110. [doi:10.1093/rfs/hhr142](https://doi.org/10.1093/rfs/hhr142)
- Chernenko, S., Erel, I., Prilmeier, R. (2022). Why do firms borrow directly from nonbanks? *Review of Financial Studies* 35(11), 4902–4947.
- Chernenko, S., Scharfstein, D. (2025). Private credit and financial stability. Working paper, December 2025. [doi:10.2139/ssrn.5930356](https://doi.org/10.2139/ssrn.5930356)
- Chernenko, S., Sunderam, A. (2016). Liquidity transformation in asset management: evidence from the cash holdings of mutual funds. NBER Working Paper 22391. [doi:10.3386/w22391](https://doi.org/10.3386/w22391)
- Chernenko, S., Sunderam, A. (2020). Do fire sales create externalities? *Journal of Financial Economics* 135(3), 602–628.

- Cipriani, M., Martin, A., McCabe, P., Parigi, B. (2014). Gates, fees, and preemptive runs. FRBNY Staff Report 670. [doi:10.2139/ssrn.2434524](https://doi.org/10.2139/ssrn.2434524)
- Cipriani, M., La Spada, G. (2021). Sophisticated and unsophisticated runs. FRBNY Staff Report 956. [doi:10.2139/ssrn.3757602](https://doi.org/10.2139/ssrn.3757602)
- Coval, J., Stafford, E. (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics* 86(2), 479–512. [doi:10.1016/j.jfineco.2006.09.007](https://doi.org/10.1016/j.jfineco.2006.09.007)
- Diamond, D., Dybvig, P. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91(3), 401–419. [doi:10.1086/261155](https://doi.org/10.1086/261155)
- Engineer, M. (1989). Bank runs and the suspension of deposit convertibility. *Journal of Monetary Economics* 24(3), 443–454. [doi:10.1016/0304-3932\(89\)90031-7](https://doi.org/10.1016/0304-3932(89)90031-7)
- Ennis, H., Keister, T. (2009). Run equilibria in the Green–Lin model of financial intermediation. *Journal of Economic Theory* 144(5), 1996–2020. [doi:10.1016/j.jet.2009.05.001](https://doi.org/10.1016/j.jet.2009.05.001)
- Ennis, H., Keister, T. (2010). Banking panics and policy responses. *Journal of Monetary Economics* 57(4), 404–419. [doi:10.1016/j.jmoneco.2010.04.005](https://doi.org/10.1016/j.jmoneco.2010.04.005)
- Falato, A., Goldstein, I., Hortaçsu, A. (2021). Financial fragility in the COVID-19 crisis: the case of investment funds in corporate bond markets. *Journal of Monetary Economics* 123, 35–52. [doi:10.1016/j.jmoneco.2021.07.001](https://doi.org/10.1016/j.jmoneco.2021.07.001)
- Fang, C., Goldstein, I., Zeng, Y. (2026). The fragility of semi-liquid private credit funds. NBER Working Paper 35385. [doi:10.3386/w35385](https://doi.org/10.3386/w35385)
- Frankel, D., Morris, S., Pauzner, A. (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory* 108(1), 1–44. [doi:10.1016/s0022-0531\(02\)00018-2](https://doi.org/10.1016/s0022-0531(02)00018-2)
- Getmansky, M., Lo, A., Makarov, I. (2004). An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics* 74(3), 529–609. [doi:10.1016/j.jfineco.2004.04.001](https://doi.org/10.1016/j.jfineco.2004.04.001)
- Goldstein, I., Jiang, H., Ng, D. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics* 126(3), 592–613. [doi:10.1016/j.jfineco.2016.11.007](https://doi.org/10.1016/j.jfineco.2016.11.007)
- Goldstein, I., Pauzner, A. (2005). Demand-deposit contracts and the probability of bank runs. *Journal of Finance* 60(3), 1293–1327. [doi:10.1111/j.1540-6261.2005.00762.x](https://doi.org/10.1111/j.1540-6261.2005.00762.x)
- Green, E., Lin, P. (2003). Implementing efficient allocations in a model of financial intermediation. *Journal of Economic Theory* 109(1), 1–23. [doi:10.1016/s0022-0531\(02\)00017-0](https://doi.org/10.1016/s0022-0531(02)00017-0)
- Greene, J., Hodges, C. (2002). The dilution impact of daily fund flows on open-end mutual funds. *Journal of Financial Economics* 65(1), 131–158. [doi:10.1016/s0304-405x\(02\)00137-x](https://doi.org/10.1016/s0304-405x(02)00137-x)
- He, Z., Manela, A. (2016). Information acquisition in rumor-based bank runs. *Journal of Finance* 71(3), 1113–1158. [doi:10.1111/jofi.12202](https://doi.org/10.1111/jofi.12202)
- He, Z., Xiong, W. (2012). Dynamic debt runs. *Review of Financial Studies* 25(6), 1799–1843. [doi:10.1093/rfs/hhs004](https://doi.org/10.1093/rfs/hhs004)
- Hombert, J., Thesmar, D. (2014). Overcoming limits of arbitrage: theory and evidence. *Journal of Financial Economics* 111(1), 26–44. [doi:10.1016/j.jfineco.2013.09.003](https://doi.org/10.1016/j.jfineco.2013.09.003)
- Huang, X., Keister, T. (2025). Can redemption fees prevent runs on funds? FRBNY Staff Report 1160.
- Jin, D., Kacperczyk, M., Kahraman, B., Suntheim, F. (2022). Swing pricing and fragility in open-end mutual funds. *Review of Financial Studies* 35(1), 1–50. [doi:10.1093/rfs/hhab022](https://doi.org/10.1093/rfs/hhab022)
- Kiss, H. J., Rosa García, A., Voellmy, L. (2025). Redemption fees and gates in the lab. SNB Working Paper 10/2025; *Journal of Money, Credit and Banking*, forthcoming. [doi:10.2139/ssrn.5380425](https://doi.org/10.2139/ssrn.5380425)
- Lewrick, U., Schanz, J. (2017). Is the price right? Swing pricing and investor redemptions. BIS Working Paper 664. [doi:10.2139/ssrn.3055690](https://doi.org/10.2139/ssrn.3055690)
- Li, L., Li, Y., Macchiavelli, M., Zhou, X. (2021). Liquidity restrictions, runs, and central bank interventions: evidence from money market funds. *Review of Financial Studies* 34(11), 5402–5437. [doi:10.1093/rfs/hhab065](https://doi.org/10.1093/rfs/hhab065)
- Ma, Y., Xiao, K., Zeng, Y. (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *Review of Financial Studies* 35(10), 4674–4711. [doi:10.1093/rfs/hhac007](https://doi.org/10.1093/rfs/hhac007)
- Matta, R., Perotti, E. (2021). Pay, stay or delay? How to settle a run. Working paper.

[doi:10.2139/ssrn.3487535](https://doi.org/10.2139/ssrn.3487535)

- Morris, S., Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3), 587–597.
- Morris, S., Shim, I., Shin, H. S. (2017). Redemption risk and cash hoarding by asset managers. *Journal of Monetary Economics* 89, 71–87. [doi:10.1016/j.jmoneco.2017.03.008](https://doi.org/10.1016/j.jmoneco.2017.03.008)
- Peck, J., Shell, K. (2003). Equilibrium bank runs. *Journal of Political Economy* 111(1), 103–123. [doi:10.1086/344803](https://doi.org/10.1086/344803)
- Pension Real Estate Association (PREA) (n.d.). Negative fundraising: redemption queues in open-end commingled funds. *PREA Quarterly*.
- Roncalli, T. (2021). Liquidity stress testing in asset management, Part 3. Amundi Working Paper.
- Schmidt, L., Timmermann, A., Wermers, R. (2016). Runs on money market mutual funds. *American Economic Review* 106(9), 2625–2657. [doi:10.1257/aer.20140678](https://doi.org/10.1257/aer.20140678)
- Shleifer, A., Vishny, R. (1992). Liquidation values and debt capacity: a market equilibrium approach. *Journal of Finance* 47(4), 1343–1366. [doi:10.1111/j.1540-6261.1992.tb04661.x](https://doi.org/10.1111/j.1540-6261.1992.tb04661.x)
- Shleifer, A., Vishny, R. (2011). Fire sales in finance and macroeconomics. *Journal of Economic Perspectives* 25(1), 29–48. [doi:10.1257/jep.25.1.29](https://doi.org/10.1257/jep.25.1.29)
- Teo, M. (2011). The liquidity risk of liquid hedge funds. *Journal of Financial Economics* 100(1), 24–44. [doi:10.1016/j.jfineco.2010.11.003](https://doi.org/10.1016/j.jfineco.2010.11.003)
- Voellmy, L. (2021). Preventing runs with fees and gates. *Journal of Banking & Finance* 125, 106065. [doi:10.1016/j.jbankfin.2021.106065](https://doi.org/10.1016/j.jbankfin.2021.106065)
- Voellmy, L. (2024). Preventing runs under sequential revelation of liquidity needs. *Journal of Economic Dynamics and Control* 158, 104789. [doi:10.1016/j.jedc.2023.104789](https://doi.org/10.1016/j.jedc.2023.104789)
- Wallace, N. (1988). Another attempt to explain an illiquid banking system: the Diamond and Dybvig model with sequential service taken seriously. *FRB Minneapolis Quarterly Review* 12(4), 3–16. [doi:10.21034/qv.1241](https://doi.org/10.21034/qv.1241)
- Woo, J. (2026). Do redemption gates discipline or destabilize? Evidence from the 2026 run on semi-liquid private credit funds. Research Square preprint. [doi:10.21203/rs.3.rs-10704359/v1](https://doi.org/10.21203/rs.3.rs-10704359/v1)
- Zeng, Y. (2017). A dynamic theory of mutual fund runs and liquidity management. Working paper. [doi:10.2139/ssrn.3723389](https://doi.org/10.2139/ssrn.3723389)
- Zitzewitz, E. (2003). Who cares about shareholders? Arbitrage-proofing mutual funds. *Journal of Law, Economics, and Organization* 19(2), 245–280. [doi:10.1093/jleo/19.2.245](https://doi.org/10.1093/jleo/19.2.245)