

Do Redemption Gates Discipline or Destabilize? Evidence from the 2026 Run on Semi-Liquid Private Credit Funds

Jihwan Woo

jihwan.woo@gmail.com

Amazon Web Services Korea

Research Article

Keywords: redemption gates, runs, private credit, business development companies, stale net asset values, liquidity transformation

Posted Date: August 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10704359/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Do Redemption Gates Discipline or Destabilize?

Evidence from the 2026 Run on Semi-Liquid Private Credit Funds

[Author blinded for double-anonymous review]*

August 2026

Abstract

Semi-liquid private credit funds cap quarterly redemptions to forestall runs. Hand-collecting 136 tender-offer filings for ten non-traded BDCs (2024–2026), we construct the first systematic measure of redemption pressure in this market during the 2026 private credit run: requests rose from 1.2 to 6.9 percent of shares outstanding and caps bound at seven of ten funds. Same-sponsor exchange-listed twins show no negative repricing when their siblings gate, locating the run in fund structure rather than assets. Requests co-move with the twins’ discounts to net asset value—consistent with redemption at stale NAVs—but the robust effect is a discrete jump in requests when the first prorations are announced. A test specified before the third-quarter filings is inconclusive in pooled form; deeply gated funds escalated 43–52 percent while a mildly gated fund’s queue cleared. Gates neither reflect asset distress nor extinguish redemption demand: they queue it.

Keywords: redemption gates; runs; private credit; business development companies; stale net asset values; liquidity transformation

*All data are from public sources (SEC EDGAR and exchange prices); a complete replication package, including the dated pre-specification draft referenced in Section 3.4, accompanies the submission.

1 Introduction

In the second quarter of 2026, investors in Apollo Debt Solutions BDC asked to redeem 11.1 percent of the fund’s shares in a single quarterly window—more than double the fund’s repurchase capacity—despite the fund paying distribution yields near nine percent and reporting no unusual credit losses. The fund honored 45 percent of the requests and prorated the rest. Across the ten largest non-traded business development companies (BDCs), a class of semi-liquid vehicles channeling roughly \$150 billion of investor capital into private corporate loans, the same pattern repeated: redemption requests that had averaged 1.2 percent of shares outstanding through 2024 and 2025 surged to 6.9 percent in the first half of 2026, and at seven of the ten funds the contractual repurchase caps—“gates” in market parlance—began to bind.

The episode is a puzzle for the theory of fund design, and its resolution matters beyond this market. Semi-liquid structures are the fastest-growing wrapper through which retail and high-net-worth savings reach illiquid private assets; the same design is spreading to private equity, real estate, and infrastructure vehicles, and regulators from the SEC to the Financial Stability Board have flagged nonbank liquidity transformation as a central monitoring priority. Whether repurchase caps actually deliver the stability they promise is therefore a live design and policy question, and the 2026 episode is its first test at scale.

Theory sharpens the puzzle. Open-end funds holding illiquid assets are vulnerable to self-fulfilling redemption spirals because early redeemers impose liquidation costs on those who remain (Chen et al., 2010). The semi-liquid design severs that externality: investors cannot redeem on demand, only tender into quarterly repurchase offers capped at five percent of shares, so no investor can dilute the claims of others by exiting first. If the payoff externality were the root cause of runs, capping redemptions should have eliminated them. It did not.

Understanding why requires distinguishing two mechanisms that the cap leaves intact, and this paper measures both. The first is a *valuation* motive. Semi-liquid funds repurchase at net asset value, and NAVs of private loan portfolios adjust slowly; when market-based valuations of the same asset class fall below carrying values, redeeming at NAV is profitable regardless of what other investors do—the private-market analogue of the stale-price arbitrage documented in open-end funds (Zitzewitz, 2003; Chalmers et al., 2001). The second is a *queue* externality. When aggregate requests exceed the cap, each investor’s fill rate depends on how many others tender; unfilled requests must resubmit; and because a tender is a free option, any news that raises expected future congestion raises requests today. Cipriani et al. (2014) show theoretically that anticipated gates can trigger preemptive runs; Li et al. (2021) find that money market funds subject to gating rules suffered larger outflows in March 2020. Whether the same logic governs vehicles in which the gate is not an emergency device but the permanent contractual state of affairs has been unknown—largely because these vehicles publish no flow data.

Our first contribution is to show that they do publish flow data, in disguise. Non-traded BDCs execute their quarterly repurchases as issuer tender offers under Rule 13e-4, and the final amendments to Schedule TO disclose, fund by fund and quarter by quarter, the shares investors sought to redeem, the shares repurchased, and whether rationing applied. Parsing 130 such filings for the ten largest funds from January 2024 through July 2026, plus six third-quarter results collected as they were filed, yields a revealed-preference panel of redemption demand for a market segment previously opaque to researchers and supervisors alike. The oversubscription ratio in these filings—requests over capacity—is, we argue, a run gauge that regulators could obtain in real time at negligible cost.

The background, the theoretical tension, and the new measurement together force a single question. The cap removed the payoff externality yet the run occurred; two candidate mechanisms—stale-NAV arbitrage and queue coordination—survive the cap and are separable with our data. Does a redemption gate restrain a run on a semi-liquid fund, or does it

accelerate one? This is the question the paper answers.

Three findings emerge. First, redemption demand escalated industry-wide—by a factor of roughly eight after normalizing by shares outstanding, an estimate whose inference survives wild cluster bootstrap with only ten fund clusters—and the caps bound at seven of ten funds, with unfilled requests reaching 7.6 times capacity at the extreme. Second, the two candidate channels are separable in the data, and the queue channel is the robust one. Tender demand co-moves with the discount to net asset value at which each fund’s exchange-listed twin trades—measured without look-ahead, at the twins’ filing dates—but the association is marginal (wild-cluster-bootstrap $p = 0.06$); conditional on the twin discount, demand jumps discretely after the first proration announcements of late March 2026 ($p = 0.03$). Third, the structure, not the assets, bears the run. Exchange-listed twins managed by the same sponsors and holding the same asset class show no negative abnormal returns on the days their semi-liquid siblings disclose proration. A test specified in a dated working draft before the third-quarter 2026 filings existed is, in its committed pooled form, inconclusive—six funds leave the permutation floor at 1/15—and we report it as such. Within those six observations, the three funds whose caps had bound deeply (oversubscription of 1.85 to 2.21 times) escalated by 43 to 52 percent with deepening rationing, while the one fund whose cap bound marginally saw its queue clear. We present this split as hypothesis-generating rather than confirmatory: it is driven by a single de-escalating fund, and several interpretations besides gate depth—the fund’s smaller size, mean reversion from an unusual second-quarter spike, a different investor base—remain open. What lends the reading independent support is that the same lifecycle appears in the only prior binding of this contract, the 2022–2024 episode of Blackstone’s non-traded REIT, where the queue jumped 72 percent in the two months after the first gate and then drained over fourteen months of continuous rationing.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 relates the paper to the literature and derives hypotheses from a two-period queue model (formalized in Appendix A). Section 4 describes the data. Section 5

presents results, and Section 6 concludes.

2 Institutional Background

2.1 Semi-liquid private credit vehicles

Non-traded perpetual BDCs emerged after 2020 as the retail-facing arm of the private credit industry. Sponsored by the largest alternative asset managers—Blackstone, Apollo, Blue Owl, Ares, HPS, Oaktree, Golub, Fidelity, and others—they raise capital continuously at net asset value, invest in senior secured loans to middle-market and sponsor-backed companies, and distribute yields historically between eight and ten percent. Shares do not trade; the sole liquidity channel is a quarterly share repurchase program, disclosed in the prospectus, under which the fund offers to repurchase up to five percent of outstanding shares at NAV. The board retains discretion to raise, reduce, or suspend the offer. When tenders exceed the cap, the fund repurchases pro rata and unfilled requests must be resubmitted the following quarter—the request queue does not carry over.

Several sponsors operate exchange-listed BDCs alongside their non-traded vehicles—Blackstone’s BXSL beside BCRED, Blue Owl’s OBDC beside OCIC and OTIC, Ares’s ARCC beside ASIF, Golub’s GBDC beside GCRED, Oaktree’s OCSL beside OSCF. The twins hold the same asset class, frequently originated on the same platforms, but their investors exit by selling shares at market prices: no queue, no gate, and a market-determined price that can and does deviate from NAV. This twin structure supplies both our control group (Section 5.4) and our measure of the market’s valuation of the asset class relative to carrying values (Section 5.3).

The repurchase offers are executed as issuer tender offers under Exchange Act Rule 13e-4. Each quarter the fund files a Schedule TO-I announcing the offer and, roughly five weeks later, a final amendment (SC TO-I/A) disclosing shares validly tendered, shares accepted, and any proration. These filings are the paper’s primary data source.

2.2 The 2026 private credit stress

Strains in the \$1.8 trillion private credit market surfaced publicly in March 2026, when mark-downs at several semi-liquid funds and rising redemption requests drew sustained press coverage; PIMCO publicly warned that investors had underpriced the liquidity risk of the asset class. No wave of realized defaults accompanied the stress during the sample window: the funds continued to report high interest coverage and low non-accruals, and their listed twins continued to pay uninterrupted distributions. The market signal that did deteriorate—well before the press coverage—was the listed twins’ price relative to NAV, which moved from premium to double-digit discount over 2025 (Section 5.3, Table 5). The stress was, on the observable record, a repricing of fund structures and of NAV reliability rather than of measured asset performance.

Why do sponsors not simply mark NAVs down to market and remove the wedge? Incentives point the other way: management fees are levied on NAV, so faster mark-downs mechanically cut fee income, and a lower NAV crystallizes underperformance relative to the listed twin that the semi-liquid wrapper is marketed to smooth away. The valuation channel examined below is thus not an accident of appraisal technology but an equilibrium outcome of fee design—a point we return to in the policy discussion.

3 Related Literature, Research Gap, and Hypotheses

3.1 Three strands of literature

The paper sits at the junction of three literatures, and the order in which we review them mirrors the logic of the research question: why runs happen, whether restricting redemptions stops them, and why the answer matters for this particular market.

The first strand is the theory of runs. In Diamond and Dybvig (1983), demandable claims on illiquid assets support a run equilibrium; Goldstein and Pauzner (2005) derive the unique

run threshold under global-games refinements. The design question—whether suspending convertibility prevents or provokes the run—is as old as the theory, and its modern form is due to Cipriani et al. (2014): if investors anticipate that a gate may bind tomorrow, they run today.

The second strand takes these mechanisms to open-end funds. Payoff complementarities make illiquid-asset funds fragile: outflows respond more sharply to bad performance where liquidation costs fall on remaining investors (Chen et al., 2010; Goldstein et al., 2017), a fragility realized system-wide in March 2020 (Falato et al., 2021) and propagated through funds’ liquidity management (Ma et al., 2022; Zeng, 2017). In money market funds, where gates exist as contingent tools, the evidence is pointed: Schmidt et al. (2016) show sophisticated investors run first, Cipriani and La Spada (2021) show investors fled the fund types made gate-eligible by the 2014 reform, and Li et al. (2021) show that in March 2020 the gate-eligible funds ran *harder*—the stabilizer became the focal point of the run. A parallel line documents a second, non-coordination motive for exit: stale NAVs create riskless redemption profits (Zitzewitz, 2003; Chalmers et al., 2001), and pricing remedies such as swing pricing dampen outflows (Jin et al., 2022). Fee incentives can keep NAVs stale (Kacperczyk and Schnabl, 2013).

The third strand explains why the semi-liquid form exists at all. Competition pushes funds toward open-ending even when collectively inefficient (Stein, 2005); closed-end structures price liquidity transformation into a premium or discount (Cherkes et al., 2009); and funds holding illiquid assets sort into closed forms (Deli and Varma, 2002). Private credit is the newest and largest asset class to face this design choice (Chernenko et al., 2022; Davydiuk et al., 2024; Block et al., 2024), and the non-traded BDC—closed-end economics with open-end marketing—is the market’s newest answer.

3.2 Literature summary and its mapping into hypotheses

Table 1 summarizes the studies most relevant to the design of our tests and records which element of the analysis each one informs.

3.3 The research gap

Reading the three strands together exposes a specific gap. The gate literature studies gates as *contingent emergency devices* attached to daily-liquidity vehicles, where their mere eligibility repels investors (Cipriani and La Spada, 2021) and their anticipated imposition accelerates runs (Li et al., 2021). It says nothing about the now-\$150-billion class of vehicles in which the gate is not a threat state but the *permanent contractual design*, operating every quarter. Second, the two exit motives that theory identifies—stale-NAV arbitrage (Zitzewitz, 2003) and coordination on future rationing (Cipriani et al., 2014)—have been studied in separate settings and never separated empirically within one market where both plausibly operate. Third, and the reason the gap has persisted: semi-liquid funds publish no flow data, so their redemption dynamics have simply not been measurable. Our tender-offer panel removes the measurement obstacle; the same-sponsor listed twins provide the instrument for separating the two motives; and the 2026 episode provides the first binding of the permanent gate.

3.4 Hypotheses

The hypotheses follow from the gap and from the model of Appendix A, in which investors hold shares repurchasable at NAV subject to a per-period cap, the market value of the assets may sit below NAV by a wedge π (observable in the listed twins’ discount, whose measurement Section 4 defines), and tender requests are free options filled pro rata. The wedge makes tendering profitable for investors with low holding values regardless of others’ actions (the valuation motive of Zitzewitz, 2003); expected future congestion makes tendering *today* optimal for investors who would otherwise wait (the preemptive logic of Cipriani et al.,

Table 1: Related literature and its role in deriving the hypotheses. H1 = escalation; H2 = structure channel; H3 = gate escalation; V = valuation (stale-NAV) channel of the discriminating test.

Study	Key finding (abridged)	Informs
Diamond and Dybvig (1983)	Demandable claims on illiquid assets support self-fulfilling runs	Foundation of H1, H3
Goldstein and Pauzner (2005)	Unique run threshold in fundamentals under global games	H1 (run trigger)
Cipriani et al. (2014)	Anticipated gates trigger preemptive runs today	H1, H3 (core)
Chen et al. (2010)	Illiquidity strengthens outflow-performance sensitivity	H1, H2
Goldstein et al. (2017)	Concave flow-performance in corporate bond funds	H1
Falato et al. (2021)	System-wide fund run in March 2020	H1 (benchmark)
Ma et al. (2022)	Liquidity management propagated the 2020 run	H2 (asset side)
Zeng (2017)	Cash rebuilding after outflows creates future run risk	H3 (dynamics)
Schmidt et al. (2016)	Sophisticated investors ran first in 2008 MMF run	H1 (heterogeneity)
Cipriani and La Spada (2021)	Investors fled fund types made gate-eligible by reform	H3 (ex ante)
Li et al. (2021)	Gate-eligible MMFs ran harder in March 2020	H3 (direct antecedent)
Zitzewitz (2003)	Stale NAVs invite dilutive arbitrage redemptions	V
Chalmers et al. (2001)	Fund NAVs embed a “wildcard option” for traders	V
Jin et al. (2022)	Swing pricing dampens outflows in UK funds	V (remedy); policy
Kacperczyk and Schnabl (2013)	Fee incentives shape money-fund risk choices	V (why NAVs stay stale)
Stein (2005)	Competition forces inefficient open-ending	Setting (why gates exist)
Cherkes et al. (2009)	Closed-end premium prices liquidity transformation	H2 (structure pricing)
Deli and Varma (2002)	Illiquid assets sort into closed structures	Setting
Chernenko et al. (2022); Davydiuk et al. (2024); Block et al. (2024)	Growth and behavior of private credit lenders and BDCs	Setting

2014, now operating in slow motion).

H1 (escalation). Tender demand rises with the NAV wedge and jumps discontinuously when proration is first observed, industry-wide rather than only at affected funds.

H2 (structure channel). Because the run reprices the fund structure and not the assets, gate-related disclosures leave exchange-listed twins' abnormal returns non-negative; the redemption wave concentrates in gated structures.


H3 (gate escalation). Conditional on a fund's cap binding, subsequent-quarter demand rises relative to funds whose caps did not bind, and oversubscription deepens rather than mean-reverts.

Pre-specified H3 test. To foreclose any suspicion of ex post specification search, we specified the exact test in a working draft dated July 28, 2026—before the first third-quarter filing existed—and preserve that draft, with file metadata, in the replication package: regress log tender demand in 2026Q3 on an indicator for whether the fund prorated in 2026Q2, controlling for log 2026Q2 tender demand, with the predicted coefficient on the proration indicator positive. Fidelity Private Credit Fund (never prorated) and any fund that lifted its cap to fulfill all requests serve as the comparison group. Section 5.6 reports this regression as specified, regardless of outcome.

Discriminating test (valuation vs. queue). The valuation channel predicts tender demand tracks the own-twin discount (measured without look-ahead, at twin filing dates); the queue channel predicts a residual jump at the first proration announcements conditional on that discount. Finding only the former would reduce the episode to stale-NAV arbitrage; only the latter, to pure coordination. Section 5.3 runs the horse race with fund-specific twin discounts.

Figure 1 summarizes the framework: the two candidate channels on the left feed the observable request queue in the center, and each hypothesis on the right maps to the estimating equation that tests it.

Figure 1: Research framework: channels, measurement, and tests.



4 Data

The sample comprises the ten largest non-traded perpetual BDCs by assets: Blackstone Private Credit Fund (BCRED), Blue Owl Credit Income Corp. (OCIC), Blue Owl Technology Income Corp. (OTIC), Apollo Debt Solutions BDC (ADS), HPS Corporate Lending Fund (HLEND), Ares Strategic Income Fund (ASIF), Oaktree Strategic Credit Fund (OSCF), Fidelity Private Credit Fund (FPCF), Golub Capital Private Credit Fund (GCRED), and Nuveen Churchill Private Capital Income Fund (PCAP). For each fund we download every Schedule TO-I/A filed between January 2024 and July 2026—130 filings—and parse, per filing and share class, shares validly tendered, shares accepted, proration language, and board actions on the cap; the six third-quarter 2026 results were collected as they were filed between July 29 and August 11, 2026 (Section 5.6). Multi-class disclosures are summed across classes. Every extracted figure is archived with the verbatim disclosure sentence it was matched from; the replication package reproduces a random verification sample of twenty filings and the full archive. Shares outstanding come from XBRL company facts for eight

funds; OCIC and OTIC do not tag shares outstanding in XBRL, so normalized specifications use the remaining eight funds. OSCF and PCAP do not disclose accepted-share counts in their TO-I/A filings and are excluded, not imputed, where acceptance is required.

Table 2: Sample and measurement.

Item	Value
Funds (non-traded perpetual BDCs)	10
Filings parsed (SC TO-I/A, 2024-01–2026-07)	130
Third-quarter 2026 results, collected as filed	6
Fund-quarter observations with request data	100
Observations normalizable by shares outstanding	78
Listed same-sponsor twins (control group)	5
Typical repurchase cap (% of shares/quarter)	5

The key variable is *tender demand*, the revealed redemption pressure of fund i in quarter t :

$$D_{i,t} = 100 \times \frac{\text{SharesTendered}_{i,t}}{\text{SharesOutstanding}_{i,t}}, \quad \omega_{i,t} = \frac{\text{SharesTendered}_{i,t}}{\text{SharesAccepted}_{i,t}}, \quad (1)$$

where $\omega_{i,t}$ is the *oversubscription ratio*: the cap binds and requests are prorated when $\omega_{i,t} > 1$, and the realized fill rate is $1/\omega_{i,t}$. The *listed discount* of twin j , timed to its filing date so that every observation was public knowledge on its date, is

$$\pi_{j,\tau} = \frac{P_{j,\text{filed}(\tau)}}{\text{BVPS}_{j,\tau}} - 1, \quad (2)$$

the twin’s share price on the day quarter- τ book value per share became public, over that book value, minus one. Because unfilled requests do not roll over, each quarter’s tender is a fresh, revealed-preference measure of the stock of investors seeking exit.

5 Results

5.1 The escalation of redemption demand

Figure 2 plots tender demand by fund-quarter. Through 2024 and the first half of 2025 demand is low and stable—the cross-fund mean is 1.17 percent of shares outstanding per quarter, well inside the cap. Aggregate tendered shares rise from 59.6 million in 2025Q1 to 135.2 million in 2025Q3, then discontinuously in the 2026 filings: the panel mean reaches 6.86 percent, with Apollo Debt Solutions at 11.1 percent and HPS Corporate Lending at 9.1 percent in the May 2026 windows. Demand was already rising through the second half of 2025, two quarters before the March 2026 press coverage; the redemption wave predates the news cycle, so the news cannot be its cause.

Figure 2: Tender demand by fund-quarter, 2024–2026.




Table 3 reports the panel regression

$$\ln D_{i,t} = \alpha_i + \beta \text{Post}_t + \varepsilon_{i,t}, \quad (3)$$

where α_i are fund fixed effects and Post_t marks results filings on or after January 15, 2026 (the windows covering the stress quarters); e^β is the implied demand multiple. Inference is appropriate to ten clusters: p -values are from a restricted wild cluster bootstrap (Rademacher weights, 9,999 replications, clustered by fund), following Cameron et al. (2008). With 10 (8) clusters the Rademacher distribution has 1,024 (256) unique sign vectors, so the draws effectively enumerate it.

Table 3: Redemption demand around the 2026 stress. Post = results filings on or after 2026-01-15.

	(1) log(shares tendered)	(2) log(tendered, % shares out.)
Post	2.568	2.058
CR1 t -statistic	12.60	8.61
Wild cluster bootstrap p	0.0015	0.0093
Fund FE	Yes	Yes
Observations	100	78
Funds (clusters)	10	8
Pre / post mean of dep. var. (level)	—	1.17% / 6.86%

5.2 Gates bind

Table 4 collects every 2026 gate event. Seven of ten funds prorated; oversubscription reached $7.61\times$ (OTIC, July window).

Table 4: Gate incidence, 2026.

Fund	Filing date	Tendered / accepted (shares)	Oversub.	Prorated
OTIC	2026-07-24	109,888,749 / 14,438,906	$7.61\times$	Yes
OTIC	2026-04-27	119,934,135 / 17,246,881	$6.95\times$	Yes
OCIC	2026-04-27	463,820,625 / 105,851,996	$4.38\times$	Yes
OCIC	2026-07-24	395,428,504 / 105,249,082	$3.76\times$	Yes
ASIF	2026-06-26	56,891,917 / 19,767,079	$2.88\times$	Yes
ASIF	2026-03-27	45,301,737 / 19,526,058	$2.32\times$	Yes
ADS	2026-05-06	66,882,332 / 30,261,082	$2.21\times$	Yes
HLEND	2026-05-01	45,605,292 / 24,646,802	$1.85\times$	Yes
GCRED	2026-05-14	15,148,979 / 8,891,200	$1.70\times$	Yes
FPCF	2026-03-09	2.8% of shares	$1.00\times$	No
FPCF	2026-06-05	2.9% of shares	$1.00\times$	No

A net-flow check addresses the concern that gross tenders overstate exit pressure while subscriptions continue. Quarterly share growth—which nets issuance against repurchases,

and therefore includes mechanical growth from distribution reinvestment—turned negative in 2026Q1 for the first time in the sample at BCRED (−3.0 percent) and OSCF (−3.3 percent), and fell to near zero at ADS; FPCF and GCRED, the funds with the mildest tender pressure, kept growing at +5.1 and +6.0 percent. The run is visible in net terms exactly where the queue is deepest.

5.3 Valuation motive versus queue amplification

If non-traded NAVs lag market values, redeeming at NAV is profitable irrespective of coordination—the stale-price arbitrage of Zitzewitz (2003) transplanted to private credit. Each fund’s listed twin prices the same sponsor’s asset class continuously, so the twin’s discount to book value proxies the NAV wedge. Two measurement choices matter. First, to avoid look-ahead, each discount observation is timed to the date the twin’s 10-Q or 10-K was *filed*. Second, the horse race uses each fund’s *own* twin, so the valuation regressor varies across funds as well as over time, separating it from any common time effect. Table 5 reports the reconstructed series: the wedge is heterogeneous—Ares and Blackstone twins held premiums into late 2025 while Oaktree’s OCSL sat at a 20 percent discount from mid-2025—which is precisely the cross-sectional variation the horse race exploits.

Table 5: Listed-twin discount to book value, timed to filing dates (%). Price on the 10-Q/10-K filing date over book value per share reported in that filing.

Quarter	ARCC	BXSL	GBDC	OBDC	OCSL
2025Q2	+3.5	+8.6	-6.9	-9.5	-20.8
2025Q3	+13.3	+10.5	-1.9	-6.1	-20.8
2025Q4	+2.8	-1.9	-10.4	-14.2	-19.9
2026Q1	-1.9	-8.3	-12.1	-22.0	-23.1
2026Q2	-3.4	-7.1	-4.9	-18.4	-22.5

Table 6 reports the discriminating regression

$$\ln \text{SharesTendered}_{i,t} = \alpha_i + \gamma_1 \pi_{j(i),t^-} + \gamma_2 \text{ProrationNews}_t + u_{i,t}, \quad (4)$$

where $\pi_{j(i),t-}$ is the discount (equation 2) of fund i 's own twin $j(i)$ most recently *filed* at least 30 days before the results filing (i.e., available when the tender window was open), and ProrationNews_t equals one for windows after the first proration disclosures of March 27, 2026. The valuation channel predicts $\gamma_1 < 0$ (deeper discounts raise demand); the queue channel predicts $\gamma_2 > 0$ conditional on the discount. Fund fixed effects absorb levels, and wild-cluster-bootstrap p -values are computed by full Rademacher enumeration over the 2^6 sign assignments.

Table 6: Discriminating the channels (fund-specific twin discounts).

Dependent variable: log(shares tendered)	Coef.	CR t	WCB p (enum.)
Own-twin discount (pp)	-0.033	-1.81	0.063
Post first-proration dummy	2.471	5.31	0.031
Fund FE	Yes		
Observations / funds	61 / 6		

Both coefficients carry the predicted sign, but their evidentiary weight differs. The queue channel is robust: conditional on the own-twin discount, tender demand is roughly an order of magnitude higher after the first proration disclosures, significant at 3 percent under enumeration bootstrap with six clusters. The valuation channel is present but marginal ($p = 0.06$). We conclude that the stale NAV plausibly supplied a motive, and the queue demonstrably supplied the amplifier.

A rare direct observation of the wedge arrived as the paper was completed. On July 27, 2026, HPS Corporate Lending Fund disclosed an unsolicited third-party tender offer from Cox Capital Retail Secondaries Fund to purchase HLEND Class I shares at \$18.40—an approximately 25 percent discount to the fund's contemporaneous NAV of \$24.42—which the board recommended rejecting (HLEND Form 8-K, July 27, 2026). The offer was capped at 550,000 shares, under one percent of HLEND's contemporaneous quarterly tender queue, so it provided no material exit capacity; its value to us is informational. Where our twin-based proxy infers the wedge from listed relatives, the Cox bid prices the *gated shares themselves*: a secondary buyer stands ready to monetize queued investors' impatience at a quarter off

NAV. The emergence of predatory secondary bids once gates bind—a pattern familiar from the non-traded REIT episode—both validates the wedge measure and quantifies the outside option of an unfilled tender request in the model of Appendix A.

5.4 The structure channel: same-sponsor twins as controls

As a first pass, we compared gated funds with an equal-weight portfolio of fifteen listed BDCs around three March news dates and found no negative repricing (Table 7); cumulatively, the listed portfolio returned +10.2 percent over March 1–April 15, 2026, against +2.3 percent for the S&P 500. *disclosed proration*. Abnormal returns are from the market model,

$$AR_{j,t} = r_{j,t} - (\hat{a}_j + \hat{b}_j r_{m,t}), \quad CAR_j[t_1, t_2] = \sum_{t=t_1}^{t_2} AR_{j,t}, \quad (5)$$

with $(\hat{a}_j, \hat{b}_j)$ estimated over a 120-trading-day window ending 10 days before each event and $r_{m,t}$ the S&P 500 return; the t -statistic scales CAR by $\hat{\sigma}_j \sqrt{t_2 - t_1 + 1}$, where $\hat{\sigma}_j$ is the estimation-window residual standard deviation. If gating conveyed bad news about the sponsor’s assets, the twin should fall.

Table 7: Listed BDC portfolio abnormal returns around gate-news events, March 2026. Market model versus SPY, 120-day estimation window ending 10 days before each event; equal-weight portfolio of 15 listed BDCs.

Event (2026)	CAR[0]	CAR[0,+1]	CAR[0,+5]
Mar 11 (private credit marked down)	-0.19% ($t=-0.18$)	-0.44% ($t=-0.30$)	3.56% ($t=1.42$)
Mar 18 (PIMCO liquidity warning)	1.13% ($t=1.03$)	1.67% ($t=1.08$)	2.11% ($t=0.79$)
Mar 26 (redemptions / gates reported)	0.13% ($t=0.12$)	-0.54% ($t=-0.36$)	2.19% ($t=0.83$)


No twin declines; the mean two-day abnormal return across the four unique events is positive (Table 8). We caution against reading the positive sign as substitution—the confidence intervals are wide—but the absence of any negative repricing at matched twins is difficult to reconcile with an asset-quality interpretation of the gates. On power: with four events the design cannot detect twin CARs smaller than roughly two percentage points in

Table 8: Listed twins’ abnormal returns on siblings’ proration disclosure days (unique events). OBDC enters once although two siblings disclosed on the same day.

Listed twin	Sibling (event date)	CAR[0,+1]	CAR[0,+5]
ARCC	ASIF (2026-03-27)	+0.97% ($t=0.59$)	+1.11% ($t=0.39$)
OBDC	OCIC & OTIC (2026-04-27)	+1.27% ($t=0.62$)	+4.62% ($t=1.31$)
GBDC	GCRED (2026-05-14)	+0.07% ($t=0.04$)	-0.79% ($t=-0.26$)
BXSL	BCRED (2026-06-04)	+3.00% ($t=1.55$)	+4.83% ($t=1.44$)
Mean across unique events ($n=4$)		+1.33% ($t=2.16$)	

absolute value; what it can rule out is a repricing of the size that accompanies genuine credit deterioration in this sector. Figure 3 plots the cumulative return paths.

Figure 3: Listed BDC portfolio versus benchmarks around the March 2026 stress.



5.5 A case study in accommodation

Blackstone’s BCRED faced first-quarter 2026 requests above its five percent cap, raised the cap to seven percent, fulfilled 100 percent of requests, and disclosed nearly \$2 billion of first-quarter subscriptions alongside approximately \$150 million of investment by Blackstone and its employees (SC TO-I/A, March 2, 2026). Second-quarter requests reached approximately ten percent of shares outstanding and the fund reverted to its five percent cap, prorating

at roughly half (SC TO-I/A, June 4, 2026). The communication campaign continued: a July 30 investor presentation emphasized \$2.7 billion of quarterly loan repayments, over \$17 billion of available liquidity, and new subscriptions of roughly \$1 billion—itsself a halving from the first quarter, evidence that the inflow side of the ledger was drying up even as the fund advertised it (Form 8-K, July 30, 2026). Third-quarter requests then rose 43 percent further and the fund prorated for a second consecutive quarter. We present this as a case, not a causal estimate: there is no counterfactual BCRED that declined to accommodate, and demand rose industry-wide. What the case documents is the weaker but still instructive fact that full accommodation, insider-alignment signaling, and sustained liquidity messaging did not prevent the queue from doubling and then growing again.

5.6 The pre-specified test: the complete 2026Q3 cycle

The cycle completed on August 11, 2026; Table 9 reports every observation, and we report the committed regression exactly as specified:

$$\ln \text{SharesTendered}_{i,Q3} = a + b \text{Prorated}_{i,Q2} + c \ln \text{SharesTendered}_{i,Q2} + e_i, \quad H_3: b > 0, \quad (6)$$

supplemented by an exact permutation test on demand growth $g_i = \ln \text{SharesTendered}_{i,Q3} - \ln \text{SharesTendered}_{i,Q2}$,

$$\Delta = \bar{g}_{\text{prorated}} - \bar{g}_{\text{comparison}}, \quad p = \frac{1}{\binom{6}{4}} \sum_{\sigma} \mathbf{1}\{\Delta_{\sigma} \geq \Delta\}, \quad (7)$$

where σ ranges over all $\binom{6}{4} = 15$ assignments of the treatment labels.


The pre-specified regression does not confirm H3 at conventional levels. Figure 4 shows why at a glance: three deeply gated funds escalated sharply while the mildly gated fourth de-escalated, and the two comparison funds were flat or shrinking.

The OLS coefficient on the proration indicator is negative (-1.18 , HC1 $p = 0.07$), but

Table 9: Third-quarter 2026 tender demand (complete cycle).

Fund	Q2 prorated? (depth)	Q2 tendered	Q3 tendered	Growth	Q3 oversub.
ADS	Yes (2.21 $\times$)	66,882,332	101,541,297	+51.8%	3.36 $\times$
HLEND	Yes (1.85 $\times$)	45,605,292	66,718,586	+46.3%	2.66 $\times$
BCRED	Yes (2.0 $\times$)	133,689,552	191,782,834	+43.5%	2.06 $\times$
GCRED	Yes (1.70 $\times$)	15,148,979	9,042,171	-40.3%	none
FPCF	No	1,427,358	1,552,234	+8.7%	none
PCAP	Unknown	1,752,092	1,400,421	-20.1%	none

Figure 4: Tender demand growth from Q2 to Q3 2026, by fund. Annotations show growth and, where applicable, the Q3 oversubscription ratio.



this estimate is uninterpretable: with six observations and treatment nearly collinear with fund size, the lagged-demand control absorbs the treatment contrast. The cleaner statistic is the exact permutation test on demand growth: prorated funds grew 23 log points faster than non-prorated funds, but with six observations the one-sided p -value is 0.27 and the attainable minimum is $1/15$. We report both, as committed, and draw no causal conclusion from this test alone.

What the completed cycle does suggest—and we label this exploratory and hypothesis-generating, not confirmatory—is a split that the pooled test obscures. The three funds whose second-quarter caps bound deeply (oversubscription $1.85\text{--}2.21\times$) all escalated by 43 to 52 percent and their oversubscription deepened; the one fund whose cap bound mildly ($1.70\times$) saw its queue *fall* 40 percent and clear entirely. We emphasize what this pattern is not: it is not a continuous dose-response. Within the deeply bound group there is no gradient (growth of 46, 43, and 52 percent against depths of 1.85, 2.0, and 2.21), and the split is driven by a single fund. Interpretations other than gate depth remain open: GCRED is the smallest treated fund, its second-quarter queue may have been an idiosyncratic spike that mean-reverted, and its investor base may differ. What lends the two-regime reading—escalation where rationing leaves most of the queue unfilled, drainage where fulfillment catches up— independent support is the REIT evidence of Section 5.7, where the same lifecycle plays out at monthly frequency. In the language of Appendix A, escalation requires that proration news raise expected future congestion; where cumulative fulfillment catches up with the queue, the same news resolves it. The market’s continuing quarterly cycles can test the two-regime hypothesis formally.

5.7 External validity: the non-traded REIT episode of 2022–2024

The BDC gates of 2026 are the contract’s first binding in that market, which precludes historical validation within the sample. The same contract, however, has bound before—in non-traded REITs, which repurchase at NAV subject to two-percent monthly and five-

percent quarterly caps. Blackstone Real Estate Income Trust (BREIT), the largest such vehicle, prorated continuously from November 2022; its prospectus supplements disclose, month by month, repurchases as a share of NAV and as a share of requests, from which the request queue can be backed out:

$$D_t^{\text{BREIT}} = \frac{\text{Fulfilled}_t (\% \text{ of NAV})}{\text{Fulfilled}_t (\% \text{ of requests})/100}. \quad (8)$$


Table 10 and Figure 5 report the reconstructed series. Because both disclosed inputs are rounded, the backed-out request series carries rounding error that is largest when the fulfillment rate is small (the December 2022 figure of 5.0 percent of NAV is computed from 0.2%/4% and could range roughly a percentage point in either direction); the lifecycle shape is robust to this imprecision.

Table 10: BREIT monthly repurchase requests and fulfillment, 2022–2023. Requests backed out as (fulfilled % of NAV)/(fulfilled % of requests). By February 2024 the fund returned to 100 percent fulfillment; by mid-2024 requests in dollar terms were 87 percent below the January 2023 peak (BREIT Form 8-K, August 7, 2024)—a dollar-based figure not directly comparable to the %-of-NAV series, as NAV itself declined over the period.

Month	Requests (% of NAV)	Fulfilled (% of requests)
2022-10	2.7	100
2022-11	4.7	43
2022-12	5.0	4
2023-01	8.0	25
2023-02	5.7	35
2023-03	6.7	15
2023-04	6.9	29
2023-05	6.7	30
2023-06	5.9	17
2023-07	5.9	34
2023-08	4.7	43
2023-09	3.5	29
2023-10	3.6	56
2023-11	3.0	67
2023-12	1.9	53

Two lessons carry over. First, the onset dynamics replicate: from the first binding month to the peak, BREIT’s queue grew 72 percent in two months—the same discrete escalation the deeply gated 2026 BDC funds display in their first post-proration quarter. Second, the

Figure 5: The BREIT gate episode: requests and fulfillment, 2022–2023.



episode separates gate *news* from gate *state*. Over fourteen months of continuous proration, a naive regression of next-month request growth on the gated state is negative: after the initial jump, the queue drains steadily as cumulative fulfillment works through the stock of exit demand and the NAV wedge closes. This is exactly the structure of Appendix A—the jump is a belief update at the first proration, not a flow response to being gated—and it supplies an out-of-sample prediction for the 2026 BDC runs: absent new adverse news, the queues should peak within one to two more quarters and then decay. The prediction is falsifiable on the same filings this paper is built from.

5.8 Heterogeneity and robustness

The one fund that never gated, Fidelity Private Credit Fund, is the sample’s smallest and distributes through a retail brokerage platform rather than wirehouse and private-bank channels; its requests stayed at 2.8–2.9 percent. We flag investor composition as a promising cross-sectional margin—echoing the institutional-first pattern in money funds (Schmidt et al., 2016)—but with ten funds we do not push a regression. Results in Table 3 are invari-

ant to raw versus normalized demand, to redefining Post as filings after March 1, 2026, to dropping any single fund, and to CR1 versus wild-bootstrap inference. Parsing accuracy is documented in the replication package; ambiguous cases are excluded rather than imputed.

6 Conclusion

6.1 Theoretical contribution

The preemptive-run logic of Cipriani et al. (2014), developed for gates as contingent emergency devices, extends to vehicles where the gate is the permanent contractual design—and it interacts with a second fragility the literature has treated separately. Stale NAVs plausibly give investors a reason to leave; rationing demonstrably gives them a reason to leave *now*: in our fund-specific horse race the proration jump survives every specification while the valuation association is marginal. What the fund-fragility literature now understands that it did not is that capping redemptions changes the speed and venue of a run—from daily flows to a quarterly request queue—without removing either the valuation motive or the strategic complementarity that drive it. The oversubscription ratio in tender filings makes that queue, for the first time, directly observable.

An honest scorecard for the gate, on the evidence assembled here, has three lines. As a *firewall*, it worked: through the largest redemption wave in the market’s history, no sampled fund suspended repurchases entirely, fire-sold assets, or collapsed—the failure mode the design exists to prevent, and the one that materialized in open-end funds in March 2020, did not occur. As a *drain*, the evidence is mixed but not damning: queues clear where fulfillment catches up with demand, as at GCRED and over the BREIT lifecycle. As a *signal*, however, the gate misfires: its first binding is precisely the news that accelerates the queue. The welfare question the episode poses is therefore not fund survival but *redistribution*—who exits at a stale NAV and who remains to absorb the subsequent mark-downs. Every prorated dollar redeemed at a NAV above market value transfers wealth from remaining to exiting investors;

the gate governs the speed of that transfer, not its direction.

6.2 Managerial implications

For sponsors, the analysis reframes the problem: the queue is fed by the NAV wedge, so the first-order managerial lever is valuation credibility, not episodic generosity. Marking NAVs to reflect market conditions faster—or adopting NAV adjustments analogous to swing pricing (Jin et al., 2022)—attacks the motive; publishing forward repurchase capacity attacks expected congestion. The BCRED case suggests that accommodation and insider-alignment signals, on their own, do not empty the queue. Distribution strategy is risk management: channels that concentrate advised high-net-worth money appear to inherit correlated redemption behavior.

6.3 Policy implications

Two findings are directly actionable for the SEC. First, tender-result disclosure is informative exactly where it is least standardized: funds report tendered and accepted shares in heterogeneous formats, and two of our ten funds omit accepted counts entirely. Requiring structured, XBRL-tagged disclosure of shares tendered, shares accepted, and proration ratios in Schedule TO-I/A amendments would give supervisors a real-time run gauge at negligible cost. Second, because the run’s welfare burden falls on remaining investors through stale-NAV redemptions, the swing-pricing debate for open-end funds should extend to NAV-based repurchase programs; the gate does not remove the arbitrage motive, it queues it, and the queue determines who bears the mark-down. For FSOC and the FSB’s nonbank monitoring program, the oversubscription ratio documented here is a candidate systemic indicator for the semi-liquid segment.

6.4 Limitations and future work

The panel is ten funds and twelve quarters; inference leans on within-fund variation, the horse-race regressors vary over fund-quarters but rest on six twin-matched funds, and the matched-twin design has power only against economically large repricings. The pre-specified gate-escalation test is reported as committed and is inconclusive in the pooled specification: with six funds, treatment is nearly collinear with size, and the permutation floor is 1/15. The two-regime reading—escalation where gates bind deeply, drainage where they bind marginally—is exploratory and awaits the market’s continuing quarterly cycles (and the Blue Owl funds’ off-cycle October filings) for formal testing. Completing shares outstanding for OCIC/OTIC and accepted counts for OSCF/PCAP from filing text is mechanical but outstanding, as is the collection of fund-level non-accrual rates and subscription flows. Extending the panel to interval funds (Rule 23c-3 vehicles), and to the Starwood REIT’s 2024 cap reduction, would test whether the two-channel mechanism and the depth dependence generalize across the semi-liquid design space.

References

- Block, J., Jang, Y. S., Kaplan, S. N., and Schulze, A. (2024). A survey of private debt funds. *Review of Corporate Finance Studies*, 13(2):335–383.
- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.
- Chalmers, J. M. R., Edelen, R. M., and Kadlec, G. B. (2001). On the perils of financial intermediaries setting security prices: The mutual fund wild card option. *Journal of Finance*, 56(6):2209–2236.
- Chen, Q., Goldstein, I., and Jiang, W. (2010). Payoff complementarities and financial

- fragility: Evidence from mutual fund outflows. *Journal of Financial Economics*, 97(2):239–262.
- Cherkes, M., Sagi, J., and Stanton, R. (2009). A liquidity-based theory of closed-end funds. *Review of Financial Studies*, 22(1):257–297.
- Chernenko, S., Erel, I., and Prilmeier, R. (2022). Why do firms borrow directly from non-banks? *Review of Financial Studies*, 35(11):4902–4947.
- Cipriani, M. and La Spada, G. (2021). Investors’ appetite for money-like assets: The MMF industry after the 2014 regulatory reform. *Journal of Financial Economics*, 140(1):250–269.
- Cipriani, M., Martin, A., McCabe, P. E., and Parigi, B. M. (2014). Gates, fees, and preemptive runs. Staff Report 670, Federal Reserve Bank of New York.
- Davydiuk, T., Marchuk, T., and Rosen, S. (2024). Direct lenders in the U.S. middle market. *Journal of Financial Economics*, 162:103946.
- Deli, D. N. and Varma, R. (2002). Closed-end versus open-end: The choice of organizational form. *Journal of Corporate Finance*, 8(1):1–27.
- Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.
- Falato, A., Goldstein, I., and 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.
- Goldstein, I., Jiang, H., and Ng, D. T. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics*, 126(3):592–613.
- Goldstein, I. and Pauzner, A. (2005). Demand-deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327.

- Jin, D., Kacperczyk, M., Kahraman, B., and Suntheim, F. (2022). Swing pricing and fragility in open-end mutual funds. *Review of Financial Studies*, 35(1):1–50.
- Kacperczyk, M. and Schnabl, P. (2013). How safe are money market funds? *Quarterly Journal of Economics*, 128(3):1073–1122.
- Li, L., Li, Y., Macchiavelli, M., and Zhou, X. (2021). Liquidity restrictions, runs, and central bank interventions: Evidence from money market funds. *Review of Financial Studies*, 34(11):5402–5437.
- Ma, Y., Xiao, K., and Zeng, Y. (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *Review of Financial Studies*, 35(10):4674–4711.
- Schmidt, L., Timmermann, A., and Wermers, R. (2016). Runs on money market mutual funds. *American Economic Review*, 106(9):2625–2657.
- Stein, J. C. (2005). Why are most funds open-end? Competition and the limits of arbitrage. *Quarterly Journal of Economics*, 120(1):247–272.
- Zeng, Y. (2017). A dynamic theory of mutual fund runs and liquidity management. Working Paper 42, European Systemic Risk Board.
- Zitzewitz, E. (2003). Who cares about shareholders? Arbitrage-proofing mutual funds. *Journal of Law, Economics, and Organization*, 19(2):245–280.

A A two-period queue model

Investors of mass 1 each hold one share of a fund whose assets have market value v per share while the fund repurchases at $\text{NAV} = 1$, so the wedge is $\pi = 1 - v \geq 0$ (observable as the listed twin’s discount). Investor i derives a private convenience value $\varepsilon_i \geq 0$ from continued holding (income clientele, tax deferral, reinvestment costs), with ε_i distributed

across investors with continuous cdf F . Each period $t \in \{1, 2\}$ the fund repurchases up to $\kappa < 1$ shares pro rata among tenders T_t : the fill rate is $f_t = \min(1, \kappa/T_t)$. Tendering is costless and unfilled requests lapse.

Investor i tenders at $t = 1$ iff the expected gain from joining the queue exceeds the convenience value foregone:

$$f_1\pi + (1 - f_1)\delta E[f_2]\pi \geq \varepsilon_i. \quad (9)$$

The left side is strictly positive whenever $\pi > 0$ and decreasing in T_1 through $f_1 = \kappa/T_1$; the threshold type ε^* satisfies (9) with equality, and $T_1 = F(\varepsilon^*)$ is interior for non-degenerate F . Three comparative statics follow.

(i) ε^* rises with π , so T_1 *is increasing in the NAV wedge*—the valuation channel, active for investors with low convenience values regardless of others' behavior.

(ii) When $T_1 > \kappa$, tenders are strategic complements in beliefs: a public signal that others' tenders are high—a proration announcement—lowers $E[f_2]$, and because delaying a tender now costs a period at a reduced expected fill rate, ε^* jumps up on the news, producing a *discrete increase in the queue conditional on π* (the queue channel; this is the object measured by the proration coefficient in Table 6). Conversely, where cumulative fulfillment has worked through most of the stock of low- ε types, the same announcement raises $E[f_2]$ relative to beliefs and the queue drains—the two-regime pattern of Sections 5.6 and 5.7.

(iii) Accommodation at $t = 1$ (raising κ to clear T_1) raises f_1 but leaves π and the distribution of ε untouched, so *the period-2 queue refills*, matching Section 5.5. ■

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [M2supplementaryreplication.zip](#)