When Investors Don't Trust the NAV: Valuation Opacity, Runs, and the Retailization of Private Equity

This Version: June 2026

Abstract

Information asymmetry creates coordination problems. When investors cannot assess portfolio value, uncertainty about losses triggers redemption races during stress. Among 195 semi-liquid alternative funds, those with above-median exposure to unobservable-input assets reach their first binding redemption-cap event more than twice as fast (Cox hazard ratio 2.6; four-year incidence 23% versus 10%). In open-end funds, where within-fund variation permits a difference-in-differences design around COVID-19, a one-point rise in Level 3 assets raises crisis-quarter redemptions by 0.12 points. The channel is distinct from liquidity mismatch and NAV staleness, and structural, not behavioral. Safeguards should target opacity, not structure.

JEL Codes: G11, G23, G28, M41

Keywords: Mutual Funds, Financial Fragility, Portfolio Opacity, Fair Value Accounting, Redemption Risk, ASC 820

I Introduction

The boundary between public and private capital markets is blurring. Over the past decade investors have increasingly accessed traditionally opaque asset classes, private equity, venture capital, and direct lending, through liquid mutual fund and semi-liquid structures. This "retailization" of illiquid markets accelerated after the Department of Labor's Information Letter of June 3, 2020, which clarified that 401(k) fiduciaries may prudently include private equity within diversified portfolios, and was catalyzed by Executive Order 14330 of August 7, 2025, "Democratizing Access to Alternative Assets for 401(k) Investors," extending opaque-asset exposure toward the tens of millions of households whose retirement savings sit in defined-contribution plans.[1] The premise is appealing, to democratize access to high-returning alternatives while preserving daily redemption rights, but it raises a question for financial stability: does portfolio opacity create fragility distinct from the well-documented risks of liquidity mismatch?

We answer that question using regulatory data on portfolio valuation uncertainty, the mandatory SEC disclosures that classify fund holdings by fair-value hierarchy under ASC 820, in particular the Level 3 assets marked with unobservable inputs such as appraisals, discounted-cash-flow models, and manager discretion. Two findings anchor the paper. In the semi-liquid funds that retirement menus would actually access, opacity drives sharply faster runs on the redemption gate: across 195 interval and tender-offer funds, those with above-median exposure to unobservable-input assets reach their first binding redemption-cap event more than twice as fast as transparent peers (Cox hazard ratio 2.6, $p < 0.01$; log-rank $p < 0.01$). In open-end mutual funds, a difference-in-differences design around the COVID-19 crisis shows that a one-percentage-point increase in lagged Level 3 assets raises crisis-quarter redemptions by 0.10 to 0.12 percentage points within fund, significant at the

[1] DOL Information Letter 06-03-2020 (June 3, 2020), issued to Groom Law Group on behalf of Pantheon Ventures and Partners Group; Executive Order 14330 (August 7, 2025); and the DOL/EBSA proposed rule of March 30, 2026 establishing process-based fiduciary safe harbors. Over $8.7 trillion resides in 401(k) defined-contribution plans, covering more than 60 million active participants (Vanguard Group, 2020).

5 percent level under two-way clustering. The first result shows where opacity fragility bites hardest; the second pins down its magnitude where within-fund variation lets us difference out everything else about the fund.

This opacity effect is distinct from portfolio illiquidity. When a genuine cash buffer receives its own crisis interaction, the opacity gradient is essentially unchanged, with a persistence ratio between 0.92 and 0.99 across opacity proxies, so opacity is not proxying for liquidity mismatch. Our aim is not to displace the liquidity-mismatch mechanism of Chen et al. (2010) and Morris et al. (2017) but to establish that information asymmetry operates as a separate amplifier: investors' inability to verify NAV calculations creates coordination failures alongside fire-sale externalities. The amplification is strongest where valuation uncertainty is greatest, largest in the COVID-19 shock, roughly half-sized in the transparent 2022 monetary-tightening episode, and undetectable around the late-2022 non-traded-REIT gating episode, scaling with the informational severity of the stress rather than its price volatility.

The contribution is both empirical and conceptual. Empirically, we are the first to measure portfolio-level opacity from regulatory fair-value disclosures and link it to redemption behavior and to the timing of fund-level run events.[2] Conceptually, we identify a fragility mechanism distinct from liquidity-driven runs: opacity-driven coordination failures. When NAVs rely on unobservable inputs, investors cannot independently verify fund valuations, and during crises this creates a strategic complementarity, since stale pricing inflates remaining shares' exposure to unmarketed losses if others redeem first. The theory is genuinely two-sided. One view, following Zeng (2017), holds that sticky NAVs should stabilize funds by muting the first-mover advantage, while the information-economics view, in the tradition of Diamond and Dybvig (1983), predicts that unverifiable valuations amplify runs because redemptions signal private information (Acharya and Merrouche, 2013). Our evidence supports the second: unlike traditional runs driven by fire-sale discounts on saleable-but-illiquid

[2] Prior work proxies opacity using asset-class labels (such as "emerging markets") or bid-ask spreads, which conflate illiquidity and information asymmetry.

assets, opacity runs emerge from non-verifiable but liquid-on-paper portfolios where the manager's discretion substitutes for market prices, and the negative, non-robust staleness interaction we estimate is consistent with the Zeng channel operating, in attenuated and opposite form, alongside it.

Identifying the effect of opacity on redemptions faces two challenges, that opacity is endogenous (funds may add Level 3 holdings precisely when expecting low redemptions) and that it correlates with unobserved fund characteristics, which we address with a difference-in-differences design exploiting COVID-19 as a plausibly exogenous shock, embedded in a strategy that spans open-end flows and interval-fund oversubscription. *Regime 1* estimates how the opacity-redemption relationship shifts during COVID in a panel of 15,709 open-end funds (214,828 fund-quarters, 2019Q4–2025Q2) with fund and time fixed effects and fund-clustered errors, identifying from pre-crisis heterogeneity in Level 3 exposure interacted with the crisis, and uses an additivity test of three nested models to show that opacity operates through an independent information channel rather than proxying for illiquidity. *Regime 2* turns to interval funds, whose periodic repurchase caps under Rule 23c-3 mean that demand to exit in excess of the offer surfaces as oversubscription, a "phantom run" invisible in open-end flow data, recovered from N-CSR filings for 179 CIKs. The ideal within-fund test of retail versus institutional behavior is not estimable, since a systematic audit finds zero CIK overlap between the interval and open-end universes, but the open-end heterogeneity analysis shows no differential opacity response across investor types ($\hat{\beta}_7 = -0.031$, $p = 0.725$), consistent with structural rather than behavioral fragility. With monthly N-PORT reporting beginning only in 2019Q4, a leads-based pre-trend test is unavailable, so we validate the design through the stress-episode comparison and a placebo battery, in which the COVID interaction ($+0.131$, $p = 0.034$) is largest, the 2022 tightening interaction ($+0.061$, $p = 0.028$) about half as large, the BREIT-episode interaction undetectable, and the COVID estimate sits in the upper tail of a non-crisis placebo distribution.

The economic stakes lie in extrapolation. Within the regulation-bounded open-end universe the estimated amplification is modest, since 78 percent of fund-quarters report zero Level 3 exposure and the 99th percentile is 4.4 percent, so a fund at that percentile faces roughly half a percentage point of excess crisis redemptions, about 12 percent of the mean. But the semi-liquid vehicles that retailization would bring to retirement savers carry a mean unobservable-input share of 33 percent, ten times the open-end 99th percentile, where the same per-unit response implies amplification on the order of the entire baseline redemption rate, delivered through funds whose quarterly repurchase caps are already chronically oversubscribed. This is the channel a calibrated welfare analysis evaluates. The SEC's framework, Rule 22e-4's liquidity programs and, since the July 2023 reforms, mandatory liquidity fees in place of the repealed money-market gates, conditions on liquidity but not on opacity, so a fund holding 40 percent cash and 60 percent Level 3 private equity triggers none of its safeguards. Comparing redemption gates, swing pricing, and an opacity-contingent hybrid, gates dominate swing pricing because they bind only on tail-event redemptions while fees tax all excess redemptions above a low threshold, and a hybrid that pairs the uniform gate with an opacity-scaled anti-dilution fee improves further by internalizing dilution where it is largest while leaving transparent funds unconstrained; and the channel is prospective, since target date funds currently hold no detectable opaque sleeves, so the analysis evaluates a regime regulators are now constructing.

The paper contributes to three literatures. It extends the mutual fund fragility literature (Chen et al., 2010; Goldstein et al., 2017; Falato et al., 2021) by identifying opacity as a channel additive to, not substitutive of, liquidity mismatch, and connects to the stale-pricing dilution literature (Choi et al., 2022; Chalmers et al., 2001) by supplying a regulatory opacity measure, a gated laboratory, and a within-fund investor-type test. It provides the first large-sample evidence that ASC 820 Level 3 classifications predict real economic outcomes, extending the fair-value-accounting literature (Song et al., 2010; Goh et al., 2015) from stock-price reactions to allocative consequences. And it informs the retailization debate (Parker et al., 2023): because opacity affects retail and institutional investors alike, the risk of retailization is scale rather than investor naivety, and effective regulation must address opacity directly rather than rely on liquidity buffers alone.

The remainder of the paper proceeds as follows. Section II provides institutional background on ASC 820 and fund liquidity regulation; Section III describes the N-PORT data, the interval-fund oversubscription extraction, and variable construction; Section IV presents the empirical strategy; Section V reports the main results; Section VI provides robustness and falsification tests; Section VII presents the welfare analysis; and Section VIII concludes.

II Institutional Background

This section establishes the three institutional facts the design relies on: the fair-value hierarchy that makes opacity measurable, the redemption rules that differ across fund structures, and the distinction between the opacity and fire-sale mechanisms that the additivity test separates.

Fair Value Accounting and the ASC 820 Hierarchy. Portfolio opacity in this paper is measured from the fair-value hierarchy of FASB ASC 820, which sorts the inputs used to value a holding into three levels of decreasing market observability: Level 1 (unadjusted quoted prices for identical assets, such as exchange-listed equities and Treasuries), Level 2 (other observable inputs, such as dealer quotes, yield curves, and matrix prices), and Level 3 (unobservable inputs, where appraisals, discounted-cash-flow models, and manager estimates rather than the market set the price).[3] The classification turns on *valuation methodology, not liquidity*: a weekly-traded bond valued from comparable spreads is Level 2, while a private property stake valued by quarterly appraisal is Level 3 even if it could be sold within a month. This is the lever of our identification, since it separates opacity (whether investors can verify

[3] ASC 820, originally SFAS 157 (FASB, September 2006); Level 3 is defined at §820-10-35-52. Common Level 3 holdings are private equity, venture capital, direct real estate, distressed debt, and complex structured products.

a valuation) from illiquidity (whether an asset can be sold without a fire-sale discount), a distinction that asset-class labels and bid-ask spreads cannot draw. All registered funds, including the interval and tender-offer funds, disclose security-level hierarchy classifications in Form N-PORT, which we use throughout and, for semi-liquid funds, adjust for holdings carried at NAV as a practical expedient (Section D); the aggregate hierarchy tables in semi-annual N-CSR reports provide an independent cross-check.[4]

Redemption Regulation and Retailization. Four rules govern how the structures we study honor redemptions. Section 22(e) of the Investment Company Act requires open-end funds to redeem within seven days and permits suspension only in narrow emergencies, so equity and bond mutual funds cannot gate.[5] Rule 22e-4 (2016) requires every open-end fund to run a liquidity risk management program and caps illiquid assets at 15% of net assets, but it conditions on *illiquidity, not opacity*: a fund holding 40% cash and 60% Level 3 private equity is fully compliant yet highly opaque.[6] Rule 22c-1, amended alongside it, authorized swing pricing to charge exiting investors for the dilution they impose, but U.S. adoption remains negligible.[7] Rule 23c-3 (1993) lets interval funds make periodic repurchase offers (5–25% of shares) and pro-rate when demand exceeds the offer, so redemption pressure above the cap surfaces as *oversubscription* rather than realized outflows, the feature Regime 2 exploits.[8] The retailization that makes this regulatory gap consequential has opened in three steps, the Department of Labor's June 3, 2020 Information Letter, Executive Order 14330 of August 7, 2025, and the DOL's proposed fiduciary safe-harbor rule of March 30, 2026, together extending opaque-asset exposure from institutions toward defined-contribution retirement

[4] N-PORT's filer scope was set in the 2016 reporting-modernization rule and reaffirmed in the SEC's August 28, 2024 amendments; through the sample, reports were prepared monthly but filed quarterly with a 60-day public lag, with monthly filing phased in from late 2025. The N-CSR requirement derives from Regulation S-X, Rule 6-03(d).

[5] Money market funds briefly held gate authority under the 2014 reform of Rule 2a-7, but the July 2023 money market reforms (SEC Release No. IC-34959) repealed it and adopted mandatory liquidity fees for institutional prime and tax-exempt funds; non-money-market open-end funds have never had comparable authority.

[6] SEC Release No. IC-32315, October 13, 2016.

[7] SEC Release No. IC-32316, October 13, 2016. The SEC proposed mandatory swing pricing in November 2022 (Release IC-34441) but withdrew it after industry opposition.

[8] SEC Release No. IC-19399, March 23, 1993.

savers.[9]

Fund Structures. The three structures impose different constraints on redemptions, generating the variation the regime design uses. *Open-end funds* (Regime 1) redeem daily at NAV with no gate, so they face both the fire-sale externality of Chen et al. (2010) and Goldstein et al. (2017), when meeting redemptions forces asset sales that dilute remaining investors, and, when holdings are opaque, a distinct friction: investors cannot verify whether the NAV reflects conservative or aggressive marks. *Interval funds* (Regime 2) cap repurchases at a board-set percentage, which mechanically prevents fire-sale-driven outflows and makes oversubscription, the ratio of tendered to accepted shares, a direct measure of the latent demand a "phantom run" represents. *Dual-access funds* (Regime 3) would offer retail and institutional share classes on an identical portfolio, isolating investor type from portfolio characteristics, but a systematic audit finds zero CIK overlap between the interval and open-end universes, so the within-fund triple-difference is not estimable from current filings and the open-end heterogeneity test (no differential retail response, $\hat{\beta}_7 = -0.031$, $p = 0.725$) provides the available evidence.

Why Opacity and Illiquidity Are Distinct. The fund-run literature emphasizes one mechanism, the fire-sale externality from liquidity mismatch (Chen et al., 2010; Goldstein et al., 2017; Morris et al., 2017): selling illiquid assets to meet redemptions moves prices, dilutes those who stay, and creates a first-mover advantage that can trigger self-fulfilling runs. We argue that opacity adds a second, distinct mechanism, a coordination failure from information asymmetry. The two act on different margins. Illiquidity constrains redemption *supply* and operates through realized asset sales, so a cash buffer directly mitigates it; opacity amplifies redemption *demand* and operates through beliefs, since investors who cannot tell whether a manager has marked Level 3 assets conservatively or aggressively may redeem to avoid bearing a future write-down even when the fund holds ample cash, the coordination

[9] DOL Information Letter 06-03-2020 (to Groom Law Group on behalf of Pantheon Ventures and Partners Group); EO 14330; DOL/EBSA proposed rule, March 30, 2026, whose evaluation factors include performance, fees, liquidity, valuation, benchmarks, and complexity.

structure of Diamond and Dybvig (1983) amplified by unverifiable valuations. Because the two channels work on independent margins, a specification including both should leave each with predictive power; if opacity merely proxied for illiquidity, controlling for liquidity would drive its coefficient toward zero. The additivity test (Section A) confirms separability with a persistence ratio near unity, and the opacity gradient is itself invariant to the fund's cash level, so a larger cash buffer neither resolves nor enables the opacity run; Section VII draws the policy implication that liquidity buffers alone cannot neutralize opacity-driven fragility. A third channel runs the other way: NAV staleness, when managers smooth reported values, attenuates the first-mover advantage (Zeng, 2017), so higher staleness predicts *fewer* crisis redemptions ($\hat{\gamma}_{\text{staleness}\times\text{COVID}} = -0.375$, $p < 0.05$), distinguishing it from the amplifying opacity channel. Table II.1 summarizes how a cash buffer bears on each.

Table II.1: Cash Buffer Effectiveness by Fragility Channel

Channel	Mechanism	Effect of a 20% cash buffer
Illiquidity (Level 1 $\downarrow$)	Fire-sale dilution	*Full coverage*: absorbs 20% redemptions with no asset sales
Opacity (Level 3 $\uparrow$)	Coordination failure	*Irrelevant*: the uncertainty is informational, not about solvency
Staleness (ξ $\uparrow$)	Sticky-NAV dampening	*Partial*: slows dilution but does not resolve the valuation ambiguity

Notes: Level 1 share proxies liquidity (higher $\Rightarrow$ more liquid); Level 3 share proxies opacity; ξ is the NAV staleness index (Getmansky et al., 2004). A fund holding 40% cash and 60% opaque Level 3 assets can meet all redemptions at face value, yet rational investors may still redeem if they believe the NAV is overstated.

III Data and Variable Construction

This section describes the data sources, sample construction, and variable definitions, with particular attention to the most consequential measurement issue: the mechanical correlation between portfolio valuation changes and flow-based redemption measures, which a lagged-NAV denominator severs so that opacity does not predict redemptions through accounting mechanics.

Primary Data Sources. We combine two SEC regulatory filings to construct a panel of fund portfolio characteristics and investor flows. Form N-PORT, filed monthly with quarterly public disclosure under the 2016 Investment Company Reporting Modernization rule, reports complete security-level holdings including the ASC 820 fair-value hierarchy level (Item C.9: Level 1 quoted prices, Level 2 observable inputs, Level 3 unobservable inputs), the asset's market value (Item C.3), and the Rule 22e-4 liquidity classification (Item C.7).[10] We aggregate these holdings to the fund-quarter level into portfolio-weighted measures of opacity, liquidity, and asset composition over 2019Q1 to 2025Q4.[11] Form N-CEN, filed annually, supplies fund-level operating data, total net assets, expense ratios, management fees, and share-class structure, which we merge to N-PORT on the SEC Central Index Key and series identifier at a 94.8% match rate for open-end funds.[12]

From the 1.29 million fund-quarter universe in N-PORT we apply three restrictions. We retain only open-end mutual funds offering daily redemptions, excluding closed-end funds, ETFs, and interval and tender-offer funds (analyzed separately in Regime 2), which reduces the sample to 335,718 observations.[13] We then drop observations with missing or invalid redemption rates, negative net assets, or Level 3 shares above 100%, removing a further 74,105 observations, and we drop singleton funds with a single quarter of data (Cameron and Miller, 2015). The full open-end panel is 214,828 fund-quarters for 15,709 series, of which roughly 191,000 to 194,000 enter the regression specifications depending on which controls each requires (the additivity test, which needs Level 1 data, has $N = 193{,}792$; the headline difference-in-differences with staleness controls has $N \approx 193{,}700$; Table A.2). Appendix Table A.1 and the data appendix (Section A) detail the attrition at each stage, which leaves

[10] 17 CFR §270.30b1-9. The rule became effective May 2018, with full compliance by April 2019. See U.S. Securities and Exchange Commission (2016).

[11] Early filings (2019Q1–Q2) have incomplete coverage as funds transitioned from N-Q reporting; we retain them and verify robustness to excluding the phase-in period.

[12] 17 CFR §270.30a-1. Unmatched observations are chiefly money market funds (excluded) and recently liquidated series.

[13] Interval and tender-offer funds register on Form N-2 but also file N-PORT under Rule 30b1-9. Their fair-value-hierarchy data require a measurement adjustment for holdings carried at NAV as a practical expedient (Section D), which is why they form a separate regime rather than being pooled with open-end funds.

roughly 70 to 75 percent of the initial open-end universe.

Mechanical Correlation in Flow Measures. Constructing redemption rates from portfolio valuations raises a mechanical-correlation problem. The standard flow measure divides net flows by current-period NAV,

$$\text{Flow}_t^{\text{Naive}} = \frac{\text{TNA}_t - \text{TNA}_{t-1} \times (1 + R_t)}{\text{TNA}_t}, \tag{1}$$

where R_t is the fund's total return: if the fund marks down Level 3 assets, TNA_t falls, the denominator shrinks, and the flow measure inflates even when no redemptions occur. Because our opacity measure works directly through discretionary NAV, equation (1) would link opacity to redemptions through accounting rather than behavior. We sever the link with a lagged-denominator measure,

$$\text{Redemption Rate}_t = \frac{\text{TNA}_t - \text{TNA}_{t-1} \times (1 + R_t)}{\text{TNA}_{t-1}}, \tag{2}$$

which scales flows by predetermined prior-period NAV, so a quarter-t markdown affects TNA_t but not the denominator. The measure passes a placebo check: in non-crisis periods (2019Q1–2020Q1) the opacity-redemption slope is flat and insignificant (-0.002, $p = 0.84$), whereas the naive measure produces a spurious negative slope (-0.18, $p < 0.001$) driven by valuation-induced denominator movement. Where N-PORT reports flows directly (Item C.7.a, available for 12% of observations), our lagged-NAV measure correlates with them at 0.96 against 0.71 for the naive measure, confirming that it captures investor behavior rather than accounting artifacts.

Key Variable Definitions. *Portfolio opacity*, our primary explanatory variable, is the share of fund assets valued with unobservable inputs,

$$\text{Opacity}_{i,t} = \frac{\sum_{j \in i} \text{Value}_{j,t} \times \mathbb{1}(\text{Level 3})_j}{\text{Total Net Assets}_{i,t}} \times 100, \tag{3}$$

the portfolio fraction marked using manager discretion (appraisals, discounted cash flow, broker quotes) rather than observable prices. ASC 820 defines Level 3 inputs as those for which there is little or no market activity (§820-10-35-52), and classification turns on valuation methodology rather than liquidity: a liquid security traded in an inactive market may be Level 3, while an illiquid security valued from observable comparables may be Level 2.[14] That distinction is what lets the additivity test separate opacity from illiquidity. The distribution is extremely right-skewed: 78.4 percent of fund-quarters report zero Level 3 exposure, the 90th percentile is 0.12 percent, and the 99th percentile is 4.4 percent, so opacity is concentrated in a thin tail of specialized vehicles, and its high within-fund persistence ($\rho_{t,t-1} = 0.94$ among funds with positive Level 3) makes it a stable fund characteristic rather than transitory noise. We complement this primary measure with several alternative opacity proxies that recur in the robustness and mechanism analyses: NAV staleness (the serial correlation in reported returns), the Level 2 share (a partial-opacity placebo), the private-asset composition (private credit, private equity, and private real estate shares), and construction variants of the Level 3 measure (forward-filled, frozen at pre-crisis values, and orthogonalized to portfolio liquidity).

To distinguish opacity from liquidity mismatch we construct two liquidity proxies, the cash buffer

$$\text{Liquidity}_{i,t}^{\text{Cash}} = \frac{\text{Cash \& Equivalents}_{i,t}}{\text{Total Net Assets}_{i,t}} \times 100 \tag{4}$$

and the quoted-price share

$$\text{Liquidity}_{i,t}^{\text{Level1}} = \frac{\sum_{j \in i} \text{Value}_{j,t} \times \mathbb{1}(\text{Level 1})_j}{\text{Total Net Assets}_{i,t}} \times 100, \tag{5}$$

both of which proxy for the ability to meet redemptions without fire-sale discounts, the friction emphasized by Chen et al. (2010) and Morris et al. (2017). These measures are orthogonal to opacity by construction: a real-estate fund holding 20 percent cash and 40

[14] A corporate bond trading twice monthly might be Level 2 (valued from spreads on similar bonds), while a frequently traded derivative on an illiquid underlying could be Level 3 (model-based).

percent private property (Level 3) is liquid yet opaque, while a fund in illiquid corporate bonds valued from dealer quotes (Level 2) is illiquid yet transparent. The same hierarchy yields an analogous Level 2 share,

$$\text{Level 2 } (\%)_{i,t} = \frac{\sum_{j \in i} \text{Value}_{j,t} \times \mathbb{1}(\text{Level 2})_j}{\text{Total Net Assets}_{i,t}} \times 100, \tag{6}$$

for assets priced from observable inputs other than quoted prices (dealer quotes, yield curves, comparable-instrument prices), a partial-opacity tier between transparent Level 1 and opaque Level 3. We enter each level as a standalone treatment in the three-level regression rather than jointly, because the shares satisfy L1%+L2%+L3% $\approx$ 100% and simultaneous inclusion is near-singular.[15]

NAV staleness captures how smoothed reported values are relative to true performance, measured by the Getmansky et al. (2004) serial-correlation index

$$\xi_i = 1 - \sum_{k=0}^{K} \hat{\theta}_k^2, \tag{7}$$

where the $\hat{\theta}_k$ are the moving-average weights fitted to fund i's NAV returns, so $\xi = 0$ is contemporaneous revelation and $\xi \to 1$ is extreme smoothing. Staleness is distinct from opacity, both conceptually, since a fund can mark Level 3 assets promptly (opaque but fresh) or smooth a transparent bond portfolio (stale but not opaque), and empirically, with a correlation of about 0.06 between staleness and the Level 3 share in our sample, weak enough to identify the two mechanisms separately. We compute the index from N-PORT monthly returns (Item C.6) using an MA(2) specification, and it is available for 99.6 percent of fund-quarters; the semi-liquid panel uses an analogous return-smoothing measure.

For interval and tender-offer funds we add an asset-type opacity proxy, the *private*

[15] The sum departs from 100% only through uninvested cash, zero-value derivatives, and unclassified securities, but the constraint is tight enough ($R^2 \approx 0.98$ from regressing the sum on a constant) that joint inclusion is near-singular; joint specifications are reported in the appendix with a singularity disclaimer.

market share, capturing the portfolio fraction in illiquid alternatives,

$$\text{Priv. Mkt. Share}_{i,t} = \frac{\text{PE} + \text{Priv. Credit} + \text{Real Estate} + \text{VC} + \text{Priv. Placements} + \text{Other Illiq. Alts.}}{\text{Total Net Assets}_{i,t}} \times 10 \tag{8}$$

summing the market value of holdings identified as private-market instruments *regardless of ASC 820 level* by ownership and asset type: private equity, private credit (direct loans and below-investment-grade debt), private real estate, venture capital, and restricted private placements.[16] This applies only to semi-liquid vehicles, since Rule 22e-4 caps open-end illiquid assets at 15 percent and the median open-end private market share is negligible. Private market share and the unobservable-input share are positively but imperfectly correlated for semi-liquid funds (about 0.42), consistent with their being alternative, imperfect proxies for the same underlying exposure to hard-to-value private holdings rather than perfect substitutes: the Level 3 share keys on valuation methodology and the private market share on asset ownership, so the two need not coincide holding by holding, yet both track a fund's concentration in opaque private assets. We therefore retain the Level 3 unobservable-input share as our primary opacity measure and use the private market share as a corroborating asset-type proxy.

Our outcome variable is the quarterly redemption rate from equation (2), winsorized at the 1st and 99th percentiles following Sirri and Tufano (1998) and Spiegel and Zhang (2013), with a post-winsorization mean of 4.2 percent (median 3.9, SD 17.5). A central feature of the data is that there is no aggregate redemption spike during COVID: the 2020Q1–Q2 mean of 3.8 percent is slightly below the 4.2 percent non-crisis mean. The fragility we document is therefore a cross-sectional reallocation of redemption pressure toward opaque

[16] We construct this share from N-PORT security-level data: a holding is private if it carries no public security identifier (no CUSIP, ISIN, or ticker), or if its asset or issuer type is private (a direct loan, a private-fund issuer, or an issuer whose name matches private-market keywords), excluding short-term investment vehicles, government issuers, and registered funds. Because the rule keys on ownership and asset type rather than the fair-value tier, the measure spans Levels 1 through 3. Lacking the external public-asset databases (Compustat, TRACE, CRSP) that would resolve identifier-bearing-but-private holdings, our share is a conservative lower bound on the full private-holdings share; narrower asset-class cuts (private equity, private credit) are available as robustness checks.

portfolios, identified within fund from the interaction of predetermined opacity with the crisis indicator, not an economy-wide run. We define $\text{COVID}_t = 1$ for 2020Q1–Q2, the acute pandemic phase when the VIX exceeded 80 and corporate spreads widened 400 basis points (Falato et al., 2021), and a placebo $\text{FedTight}_t = 1$ for 2022Q1–Q3, when the funds rate rose 300 basis points and equities fell 20 percent but without Knightian uncertainty about how to value assets. Five fund-level controls absorb confounding variation: lagged return (performance chasing; Sirri and Tufano, 1998), log fund size (Ferreira et al., 2013), the expense ratio (Pástor and Stambaugh, 2012), the two liquidity proxies, and asset-class indicators, the last subsumed by fund fixed effects in the saturated specifications.

Interval Fund Data (Regime 2). Regime 2 extends the analysis to interval funds, which offer periodic share repurchases under Rule 23c-3 rather than daily liquidity. We identify 179 interval-fund CIKs from the closed-end fund master file, cross-referenced with EDGAR full-text search for N-23C3A repurchase notifications, spanning 2019 to 2025 and covering private credit (42 percent of funds), real estate (23 percent), multi-strategy alternatives (18 percent), and private equity funds-of-funds (17 percent). From Form N-CSR and N-CSRS shareholder reports we extract the repurchase offer date and percentage, shares tendered, shares repurchased after proration, and proration percentages, and construct the oversubscription ratio

$$\text{Oversubscription}_{i,t} = \frac{\text{Shares Tendered}_{i,t}}{\text{Shares Offered for Repurchase}_{i,t}}, \tag{9}$$

a value above one marking a phantom run in which investors demand more liquidity than the gate provides. Filing formats vary across families and require careful hand-collection (Appendix A), and the extractions span a wide range, from the BlackRock Enhanced Government Fund's 12.5× oversubscription in its 2023 annual offer (71 percent of shares tendered against a 10 percent offer) and Pioneer ILS Interval Fund's recurring 2 to 6×, down to the Invesco Senior Loan Fund's undersubscribed 0.15×, consistent with transparent, liquid holdings attracting less redemption pressure.

Applying the open-end Level 3 measure unchanged would understate semi-liquid opacity, because many of these funds' positions, especially the funds-of-funds among the tender-offer funds, are investments in underlying private funds carried at NAV as a practical expedient (ASC 820-10-35-54B) and therefore excluded from the Level 1/2/3 hierarchy and reported with no level tag, despite being the most opaque positions in the portfolio. Because such holdings are the most opaque positions in the portfolio yet carry no fair-value tag, we treat them as Level 3 and define the canonical semi-liquid opacity measure as the *unobservable-input share*,

$$\text{Opacity}_{i,t}^{\text{semi-liquid}} = \frac{\text{Level 3 Assets}_{i,t} + \text{NAV-Practical-Expedient Assets}_{i,t}}{\text{Total Portfolio Value}_{i,t}} \times 100. \tag{10}$$

The untagged holdings are genuine private-fund stakes rather than data artifacts, the largest being vehicles such as Ironwood Institutional Multi-Strategy, Millennium International, AMG Pantheon Master Fund, and Barings Private Credit, with issuer-type codes overwhelmingly "private fund." The adjustment raises mean tender-offer-fund opacity from 9 percent (hierarchy-tagged Level 3 only) to 33 percent, while interval-fund and open-end measures are essentially unchanged, and we report all semi-liquid results under both conventions so that no result depends on the classification. As an external cross-check we extract aggregate Level 1/2/3 dollar amounts from the fair-value-measurements notes of N-CSR filings; the two sources agree on levels for the largest opaque vehicles, and the structured, security-level N-PORT measure is the primary one.

Dual-Access Fund Data (Regime 3). Regime 3 is designed to exploit dual-access interval funds, those offering retail and institutional share classes on the same portfolio, to test whether investor type affects opacity-driven fragility. A data audit finds zero CIK overlap between the interval-fund universe (70 CIKs with N-CEN `IS_INTERVAL = Y`) and the open-end universe (15,709 series), because interval funds register under Section 23(c)(3) as legally separate entities, so fund-level matching cannot identify dual-access structures and share-class-level cross-referencing would be required. The within-interval triple-difference is therefore beyond the scope of this paper; we retain the design to document the data limitation and rely on the open-end retail heterogeneity evidence (Table A.9), which shows no significant differential opacity response across investor types ($\hat{\beta}_7 = -0.031$, $p = 0.725$), consistent with structural rather than behavioral fragility.

Summary Statistics. Table A.1 reports summary statistics: Panel A for open-end funds (Regime 1), Panel B for interval funds (Regime 2), and Panel C for tender-offer funds; statistics for business development companies, analyzed separately because of their near-constant Level 3 exposure (median 96.9 percent), appear in Appendix Table B.6. The open-end panel contains 214,828 fund-quarters for 15,709 series over 2019Q4–2025Q2, of which roughly 191,000 to 194,000 enter the main specifications. The median open-end fund holds zero Level 3, 94 percent Level 1, and 5 percent Level 2 assets, with a mean quarterly redemption rate of 4.2 percent (SD 17.5), and mean redemptions do not spike in aggregate during COVID (3.8 percent against 4.2 percent outside the crisis) or in the 2022 tightening (2.2 percent). High-opacity open-end funds differ from transparent peers in ways that cut against a naive retail-fragility story: they are larger (median TNA \$571M against \$141M) and have slightly lower baseline redemptions (3.9 against 4.2 percent, $p = 0.005$), and raw group means show no univariate divergence during COVID (-0.21pp, $p = 0.61$). The amplification documented in Section V emerges within fund once fixed effects absorb these compositional differences, which is exactly why the difference-in-differences design rather than a comparison of group means is required. Among interval funds, opacity is far higher (median 28 percent, 75th percentile 52 percent), the mean COVID-quarter oversubscription ratio exceeds 2.0× for the top opacity tercile against 1.1× for the bottom, and 78 percent of funds offer the 5 percent regulatory minimum repurchase. Unconditionally, opacity correlates negatively with the Level 1 share (-0.09) and near-zero with size (0.09), lagged returns (0.00), and crisis-period redemptions (0.01): the opacity-redemption relationship is conditional, appearing within fund in the interaction of predetermined opacity with the crisis indicator.

Data Quality and External Validity. ASC 820 classifications are self-reported, raising the possibility that managers facing redemption pressure reclassify Level 3 assets as Level 2 to obscure opacity, but three features limit this concern. Fair-value disclosures derive from externally audited financial statements and are an area of active SEC examination, raising the cost of misclassification; opacity is highly persistent for funds with private holdings ($\rho_{t,t-1} = 0.94$), inconsistent with opportunistic reclassification; and the N-PORT security-level measure aligns with the aggregate fair-value tables in the same funds' N-CSR reports for the largest opaque vehicles. Because N-PORT filing is mandatory for registered investment companies, the panel is not subject to the voluntary-reporting selection that affects commercial databases, and attrition relative to the raw filing universe is concentrated among money market funds (excluded by design), recently launched funds, and observations failing the data-quality filters. Together, the lagged-NAV redemption measure removes mechanical opacity-flow correlation, opacity is distinct from illiquidity by construction, the open-end panel provides identifying variation in a thin opacity tail, and the interval-fund oversubscription data recover latent demand invisible in open-end flows, establishing the data infrastructure for the regime design.

IV Empirical Strategy

This section presents our difference-in-differences design, which exploits the COVID-19 crisis as a quasi-experiment to identify the effect of portfolio opacity on redemption fragility. We set out the baseline specification, state the identifying assumption and the threats to it, and describe the additivity, hierarchy, staleness, cross-vehicle, and placebo specifications that the results section takes up.

Baseline Difference-in-Differences Specification. The primary specification estimates how the relationship between opacity and redemptions shifts during the crisis,

$$\text{Redemption}_{i,t} = \beta_0+\beta_1\text{Opacity}_{i,t-1}+\beta_2\text{COVID}_t+\beta_3(\text{Opacity}_{i,t-1}\times\text{COVID}_t)+\mathbf{X}'_{i,t}\boldsymbol{\gamma}+\alpha_i+\delta_t+\varepsilon_{i,t}, \tag{11}$$

where $\text{Redemption}_{i,t}$ is fund i's winsorized quarterly redemption rate (Section III), $\text{Opacity}_{i,t-1}$ is the one-quarter-lagged share of assets classified as ASC 820 Level 3, COVID_t equals one in 2020Q1–Q2, $\mathbf{X}_{i,t}$ collects time-varying controls (cash and Level 1 shares, log fund size, lagged returns, and the expense ratio), α_i and δ_t are fund and quarter fixed effects, and standard errors are clustered by fund. The coefficient of interest is β_3, the differential effect of opacity during the crisis: under the null of no opacity-driven fragility $\beta_3 = 0$, while the information-asymmetry hypothesis predicts $\beta_3 > 0$. We lag opacity one quarter both to mitigate reverse causality, since current redemptions cannot move prior-quarter portfolio composition, and to align with investor information sets, since investors observe opacity from the prior quarterly N-PORT filing. Fund fixed effects absorb time-invariant differences across funds (manager skill, investor base, fee structure) and quarter fixed effects absorb common shocks (monetary policy, market returns), so the identifying variation is the within-fund change in the opacity-redemption relationship between COVID and non-COVID quarters.

Because COVID is a common shock rather than a staggered rollout, the heterogeneity-and-timing bias that motivates modern difference-in-differences estimators (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021) does not arise in its usual form, but treatment intensity does vary continuously across funds through Level 3 exposure. We therefore report the two-way fixed-effects estimate (Correia, 2016) as the main specification; because COVID is a single common shock with no staggered adoption, the negative-weighting and forbidden-comparison concerns that motivate the recent heterogeneity-robust estimators do not arise in their usual form.

Identification. The crisis struck every fund simultaneously, with no staggered rollout and no never-treated control group, so the design is not a canonical treated-versus-untreated difference-in-differences but a continuous-treatment design that identifies differential response across the cross-sectional distribution of Level 3 exposure. The identifying assumption is that, absent COVID, quarterly redemption rates would have evolved in parallel across that distribution, so that conditional on fund and quarter fixed effects the differential crisis response is attributable to heterogeneous opacity. The assumption would fail if high-opacity funds systematically attracted investors with a pre-existing propensity to redeem on any negative signal, independent of the information-asymmetry mechanism, or if high-opacity funds selectively attrited during the crisis. A standard leads-based pre-trend test is not available here, because monthly N-PORT reporting begins only in 2019Q4 and the open-end panel therefore contains a single pre-COVID quarter; we validate the design instead through the stress-episode comparison and the placebo battery of Section VI, which show that the opacity interaction is largest under COVID-era valuation uncertainty, attenuates by half in the transparent 2022 tightening, and lies in the upper tail of a non-crisis placebo distribution.

The headline estimate is the standalone Level 3 $\times$ COVID specification. The reason is mechanical: the three hierarchy shares satisfy L1% + L2% + L3% $\approx$ 100%, so entering all three simultaneously induces near-exact collinearity. The joint specification is therefore a diagnostic rather than the headline, and a Level 3 coefficient that keeps its sign but loses precision in the joint specification reflects that collinearity rather than a weakening of the result.

Disentangling Opacity from Liquidity: The Additivity Test. The most salient confound is portfolio illiquidity: if Level 3 assets are systematically more illiquid, the opacity coefficient could capture fire-sale risk rather than information asymmetry. Our aim is not to displace the liquidity-mismatch mechanism of Chen et al. (2010) and Goldstein et al. (2017) but to establish that opacity is a distinct, additive channel, which we test through nested specifications: an opacity-only model

$$\text{Redemption}_{i,t} = \beta_3^A \left(\text{Opacity}_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t}, \tag{12}$$

a liquidity-only model

$$\text{Redemption}_{i,t} = \beta_3^B \left(\text{Illiquid}_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t}, \tag{13}$$

where $\text{Illiquid}_{i,t-1}$ is one minus the Level 1 share, and a joint model

$$\text{Redemption}_{i,t} = \beta_3^C \left(\text{Opacity}_{i,t-1} \times \text{COVID}_t\right) + \beta_4^C \left(\text{Illiquid}_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t}. \tag{14}$$

Additivity predicts that the persistence ratio $\rho \equiv \beta_3^C / \beta_3^A$ stays near one; if opacity merely proxied for illiquidity, adding the illiquidity interaction would drive β_3^C toward zero. As a further refinement we orthogonalize Level 3 on illiquidity,

$$\text{Opacity}^{\perp}_{i,t} = \text{Opacity}_{i,t} - \hat{\pi}\,\text{Illiquid}_{i,t}, \tag{15}$$

with $\hat{\pi}$ from a first-stage regression of opacity on illiquidity, so that any remaining opacity effect is mechanically uncorrelated with the illiquidity channel.

The Three-Level Specification. A sharper test asks whether crisis amplification attaches to the opacity of the inputs specifically rather than to less-liquid pricing generally. We estimate

$$\text{Redemption}_{i,t} = \beta_1 \left(\text{Lk}_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t}, \quad k \in \{1, 2, 3\}, \tag{16}$$

using each ASC 820 level in turn as the sole treatment, estimated separately because of the near-100% summation constraint. Using one-quarter-lagged shares also contains a reclassification threat: many funds moved assets from Level 2 to Level 3 during 2020Q1–Q2 as observable inputs degraded, and a contemporaneous Level 2 share would capture that reclassification rather than pre-crisis composition. The full joint specification is reported only as a supplementary diagnostic, since the near-100% summation leaves the individual level coefficients not separately identified; the drop-one form, Level 1 and Level 3 with Level 2 as the omitted reference, is the interpretable version.

NAV Staleness as a Stabilizing Channel. The framework accommodates a third channel that works in the opposite direction from opacity. Following Zeng (2017), NAV staleness should reduce crisis-period redemptions by shrinking the payoff difference between early and late redeemers: while both staleness and high Level 3 exposure reduce the informativeness of reported NAVs, opacity creates coordination failures by making the NAV hard to verify, whereas staleness dampens runs by making the NAV slow to move. We test this by adding a staleness-COVID interaction to the headline specification,

$$\text{Redemption}_{i,t} = \beta_3 \left(\text{L3}_{i,t-1} \times \text{COVID}_t\right) + \gamma \left(\xi_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t}, \quad (17)$$

where $\xi_{i,t-1}$ is the one-quarter-lagged Getmansky-Lo-Makarov staleness index, and the prediction $\gamma < 0$ contrasts with the opacity prediction $\beta_3 > 0$. The results section reports $\hat{\gamma} = -0.375$ ($p < 0.05$, not robust to two-way clustering) and uses a lead-the-outcome timing test to separate a genuine sticky-NAV mechanism from mechanical lagged recording of crisis redemptions.

Cross-Vehicle Specification. The three-regime design assumes the opacity-fragility mechanism operates across fund structures, which we test by estimating the baseline difference-in-differences separately for open-end, interval, and tender-offer funds,

$$\text{Redemption}_{i,t} = \beta_{3,s} \left(\text{L3}_{i,t-1} \times \text{COVID}_t\right) + \mathbf{X}'_{i,t}\boldsymbol{\gamma}_s + \alpha_i + \delta_t + \varepsilon_{i,t}, \quad s \in \{\text{OE, IV, TO}\}. \quad (18)$$

The dependent variable measures different objects across structures, unconstrained submitted redemptions for open-end funds, gate-fulfilled redemptions for interval funds, and accepted tenders for tender-offer funds, so coefficient magnitudes are not comparable across columns and only sign and significance are. Because realized redemptions in the capped vehicles are supply-constrained, the demand-side evidence for interval and tender-offer funds is the oversubscription and survival analysis of Section B rather than the realized-flow coefficient.

Placebo and Stress-Episode Design. To rule out that high-opacity funds simply redeem more in any market stress, we benchmark the COVID interaction against two other episodes (Section F). The 2022 Federal Reserve tightening raised the funds rate sharply and produced a 20 percent equity decline and a 400 basis-point widening in investment-grade spreads, market stress comparable to COVID in price terms but observable and anticipated, with no deep uncertainty about how to value assets; the BREIT gating episode concentrated stress in non-traded real estate. If opacity amplifies redemptions through valuation uncertainty, the amplification should scale with the informational severity of the episode rather than its price volatility. The open-end panel ($N \approx 193{,}700$) has ample power for the estimated effect; Regime 2 ($N = 4{,}107$, binary outcome) detects its reported effect at conventional levels; and Regime 3 is not estimable in the current data (Section C).

Regime 2: Oversubscription in Interval Funds. To test whether the opacity-fragility link is an artifact of the open-end structure, Regime 2 turns to interval funds, closed-end vehicles that offer periodic share repurchases at NAV (typically 5 to 10 percent of shares per quarter) under Rule 23c-3. When investor demand to exit exceeds the offer, the fund pro-rates redemptions, producing an observable oversubscription ratio,

$$\text{Oversubscription}_{i,t} = \frac{\text{Shares Tendered}_{i,t}}{\text{Shares Offered for Repurchase}_{i,t}}, \tag{19}$$

a value above one signalling a phantom run in which investors collectively demand more liquidity than the fund provides while the gate caps actual redemptions, so that the ratio reveals latent pressure invisible in open-end flow data. We extract repurchase results from Form N-CSR and N-CSRS filings for 179 interval-fund CIKs identified through the closed-end fund master file, supplemented by N-23C3A repurchase notifications, and estimate

$$\begin{aligned}\text{Oversubscription}_{i,t} = \gamma_0 + \gamma_1\,\text{Opacity}_{i,t-1} + \gamma_2\,\text{COVID}_t + \gamma_3\,(\text{Opacity}_{i,t-1} \times \text{COVID}_t) \\ + \mathbf{Z}'_{i,t}\boldsymbol{\lambda} + \eta_i + \theta_t + u_{i,t},\end{aligned} \tag{20}$$

where $\mathbf{Z}_{i,t}$ includes fund size, offer frequency, and the repurchase offer percentage, and γ_3 tests whether opacity intensifies the phantom run during the crisis.

Regime 3: Investor Heterogeneity and the Within-Fund Test. The retailization debate turns on whether extending opaque-fund access to retail investors creates incremental fragility beyond what institutional investors already generate. The cleanest test would be a triple-difference comparing retail and institutional share classes within the same interval fund,

$$\begin{aligned}\text{Redemption}_{i,c,t} = \phi_1\,(\text{Opacity}_{i,t-1} \times \text{COVID}_t) + \phi_2\,(\text{Opacity}_{i,t-1} \times \text{COVID}_t \times \text{Retail}_c) \\ + \boldsymbol{W}'_{i,c,t}\boldsymbol{\psi} + \alpha_i + \delta_t + \mu_c + \varepsilon_{i,c,t},\end{aligned} \tag{21}$$

where c indexes share class and Retail_c marks retail classes, so that because the two classes hold an identical portfolio any differential response in ϕ_2 must come from investor type rather than portfolio composition: $\phi_2 = 0$ would indicate structural fragility from the coordination problem, $\phi_2 > 0$ behavioral fragility from retail naivety. Assembling the share-class-level redemption panel this design requires is not feasible in the current EDGAR universe (Section C), so we estimate the same triple-difference on the retail-versus-institutional dimension of the open-end sample, where it is identified (Section G); the within-interval comparison would require a share-class-level redemption panel that current filings do not support.

V Main Results

This section presents the central empirical findings across the paper's regime design. We begin with an additivity test that estimates the opacity and liquidity fragility channels jointly, then turn to the interval-fund evidence (Regime 2), where opacity predicts both chronic oversubscription and a sharply accelerated time to a fund's first binding redemption cap. We then report the open-end difference-in-differences estimates (Regime 1) showing that opacity amplifies redemptions during COVID-19, quantify the economic magnitudes, and present the validation evidence: a stress-episode comparison, a placebo battery, and a heterogeneity analysis of whether opacity effects differ for retail-oriented funds.

Opacity and Liquidity as Distinct Channels. Table A.4 and Figure A.4 ask whether portfolio opacity and portfolio liquidity are distinct fragility channels. For each opacity proxy we compare its standardized oversubscription gradient estimated alone, $\hat{\beta}^A$, with the gradient $\hat{\beta}^C$ when a genuine cash buffer, economically independent of the fair-value hierarchy, receives its own term; additivity predicts that the persistence ratio $\rho \equiv \hat{\beta}^C/\hat{\beta}^A$ is near one, while substitution predicts that it collapses toward zero. Across the Level 3 unobservable-input share and three private-holdings proxies based on ownership and asset type, the opacity gradient is essentially unmoved by the cash control: ρ ranges from 0.92 to 0.99 (Level 3 0.95, private credit 0.99, no-identifier private 0.94, broad private-holdings 0.92), all clustered just below the fully-additive benchmark of one, so opacity is not standing in for portfolio liquidity. Using a cash buffer rather than the Level 1 share matters because cash is independent of the $(1 - \text{Level 1})$ decomposition, so the separability is not a mechanical artifact of the fair-value hierarchy. A larger cash buffer neither substitutes for nor amplifies the opacity channel: interacting opacity with the cash level leaves a near-zero, insignificant interaction (+0.6 percentage points, $p = 0.80$) even as the opacity gradient (+9.9 per standard deviation) and cash's own gradient (+5.5) are each precisely estimated, so cash buffers leave opacity-driven fragility untouched rather than easing it. The significance evidence for the opacity channel is the full-panel difference-in-differences of Section D.

Regime 2: Interval Fund Oversubscription and Survival. The additivity test establishes that opacity amplifies realized redemptions in open-end funds with daily liquidity. A natural concern is that this reflects something specific to the open-end structure, daily NAV pricing, first-mover advantages, fire-sale externalities, rather than a general property of opacity, and Regime 2 addresses it by asking whether opacity predicts latent redemption pressure in interval funds, which redeem on entirely different mechanics. Interval funds offer periodic share repurchases (typically 5 to 25 percent of NAV quarterly) under SEC Rule 23c-3, and when investor demand to exit exceeds the offer, the fund pro-rates redemptions; the gap between requested and offered amounts reveals a "phantom run," latent fragility invisible in realized flow data. If opacity creates coordination failures through information asymmetry rather than through fire-sale externalities, high-opacity interval funds should show higher oversubscription during crises. We extract oversubscription from semi-annual N-CSR filings for 179 interval-fund CIKs identified through N-CEN structure flags, measuring opacity primarily by the N-PORT security-level unobservable-input share (equation (10)) and cross-checking against the aggregate fair-value tables in N-CSR (Section D).

The sharpest evidence is a survival analysis of the time to a fund's first binding redemption-cap event, with one spell per fund: entry at the first observed quarter, failure at the first oversubscribed repurchase offer, right-censored otherwise. Opacity accelerates the run on the gate. Splitting funds at the median unobservable-input share, four-year cumulative first-cap incidence is 23 percent for high-opacity funds against 10 percent for transparent peers (sixteen-quarter gate-free survival of 0.77 versus 0.90; Table A.11). A Cox proportional-hazards model puts the hazard of a first cap event at 2.61 for high-opacity funds ($p = 0.005$) and 1.10 per ten-percentage-point increase in opacity ($p = 0.011$), and the log-rank test agrees ($p = 0.004$): high-opacity funds reach their first redemption-cap event more than twice as fast as transparent ones. The same force appears offer by offer. Among repurchase offers with extractable tender and proration data, high-opacity funds were oversubscribed in 50 percent of offers against 22 percent for transparent peers, an odds ratio of 3.6: the phantom run made visible, as investors who cannot verify the NAV tender for the exit and the gate rations the excess away. The contrast is descriptive, since the 94 offers cluster within 19 funds and the 50-versus-22 gap is not significant under fund clustering ($p \approx 0.22$); identification rests on the 193,000 fund-quarter open-end panel (Section D) and the cross-sectional gradient below.

Because opacity in semi-liquid funds is essentially time-invariant, a private-credit interval fund permanently opaque and an equity-income tender fund permanently transparent, the identifying variation is between funds rather than within, and fund fixed effects would absorb exactly the variation of interest. Explanatory power for oversubscription rises from $R^2 = 0.043$ without fund fixed effects to $R^2 = 0.933$ with them (Appendix Table B.2), confirming that oversubscription is overwhelmingly a persistent fund attribute; Appendix Table B.3 reports the formal between-within variance decomposition. We therefore estimate the relationship cross-sectionally, clustering by fund, and find a positive and significant gradient, $\hat{\beta}_{\mathrm{L3}} = +0.095$ ($p < 0.05$): high-opacity funds are unconditionally far more likely to be oversubscribed. The within-fund difference-in-differences interaction is, by contrast, near zero, $\hat{\gamma}_3 = +0.0002$ ($N = 4{,}107$), which is the mechanical consequence of differencing out a between-fund characteristic rather than evidence against the channel, since once the permanent opacity level is absorbed only marginal crisis-quarter movement remains and the offer cap was already binding before COVID. Read this way, the headline Regime 2 finding is not that COVID marginally elevated oversubscription but that high-opacity semi-liquid funds are chronically oversubscribed as a structural equilibrium: opacity prevents investors from verifying whether the quarterly cap is adequate, generating persistent excess demand to exit, of which the crisis is a marginal intensification rather than a regime change. This distinction matters for policy, because opacity-contingent gate design must address the structural constraint on exit, not merely its crisis-period amplification.

The cross-sectional pattern is robust to how opacity is proxied. Estimated per standard deviation without fund fixed effects, the Level 3 unobservable-input share carries the sharpest oversubscription gradient, raising the probability of a binding cap by 10.5 percentage points ($p < 0.01$), and asset-type private-holdings proxies corroborate it in the same direction: the no-identifier private share (+6.4, $p < 0.05$), the broad private-holdings share (+5.4, $p < 0.05$), and the ownership-substance and restricted-securities shares more weakly (Table IA.15, Panel D; Figure A.3). That the asset-type proxies, imperfect lower-bound measures of the same private-holdings exposure, all load positively while the directly observed Level 3 share is the most precise is what one expects of alternative proxies for a single underlying construct, not of independent channels. Conditional on fund and quarter fixed effects, every COVID interaction is negative or null, consistent with the supply-cap mechanism in which binding caps mechanically reduce realized redemptions during COVID even as latent demand rises. Placing these results in cross-vehicle context (Appendix Table B.6), the informative contrast is between the semi-liquid panel, where the Level 3 share ranges from near zero to 100 percent, and non-traded business development companies, whose Level 3 share has a median of 96.9 percent and an interquartile range of 91 to 99 percent. That near-constant opacity makes BDC regressions uninformative, since with no cross-sectional variation the opacity coefficient cannot be identified ($\hat{\beta} \approx 0$, $p > 0.1$ throughout, reflecting the absence of variation rather than a true null). BDCs are structurally all-opaque vehicles for which the binding constraint is the redemption cap rather than opacity variation, and the appropriate comparison is descriptive: public BDCs trade at a median discount of 8.1 percent to NAV, a market price for illiquidity risk that is absent from non-traded structures.

Regime 3: Retail versus Institutional Investors Within Interval Funds. The retailization debate turns on whether extending opaque-fund access to retail 401(k) participants creates incremental fragility beyond what institutional investors already generate, and the cleanest test would compare retail and institutional share classes within the same interval-fund portfolio, so that any differential response must arise from investor type rather than portfolio composition. Constructing that dual-access panel, however, requires identifying funds that file simultaneously as interval funds and as open-end funds under a single registrant, and systematic matching of the 206-fund interval universe against the open-end universe yields zero CIK overlaps.[17] The within-interval triple-difference is therefore not estimable from the current EDGAR universe. Indirect evidence comes from the open-end sample, where the retail heterogeneity analysis (Section G) shows no differential opacity response between small, retail-oriented funds and large, institutional ones, which, together with the structural argument that opacity creates coordination failures regardless of investor sophistication, anchors the prediction that the within-interval result would also be null.

Main Difference-in-Differences Results. Table A.2 reports the open-end difference-in-differences estimates of equation (11), which test whether opacity amplifies redemptions specifically during COVID-19 under the identifying assumption that high- and low-opacity funds would otherwise have followed parallel redemption trends. The baseline specification with fund and quarter fixed effects gives $\hat{\beta}_3 = 0.104$ (SE $= 0.056$, $N = 193{,}668$), so each percentage point of lagged Level 3 exposure amplifies COVID-quarter redemptions by about a tenth of a percentage point, identified from within-fund variation. Adding lagged log size strengthens rather than attenuates the effect ($\hat{\beta}_3 = 0.113$), and adding the NAV staleness index with its own COVID interaction raises it further to $\hat{\beta}_3 = 0.118$: controlling for the return-smoothing channel isolates the informational component of Level 3 opacity, which is larger once staleness is held constant. Retaining that specification with two-way clustering by fund and quarter (Cameron and Miller, 2015) leaves the coefficient unchanged at $\hat{\beta}_3 = 0.118$ with $p = 0.045$. Across the ladder $\hat{\beta}_3$ ranges from 0.104 to 0.118, significant at the 10 percent level under fund clustering and at the 5 percent level under two-way clustering, and we carry $\hat{\beta}_3 = 0.118$ as the headline.

That opacity and staleness move in opposite directions during the crisis is itself informative, because the two are conceptually distinct: the Level 3 share is a portfolio-composition

[17] This is consistent with the regulatory structure: interval funds register under Section 23(c)(3) as separate legal entities from any co-sponsoring open-end family, so the same CIK cannot file both N-23C3A and N-PORT forms.

measure of how many assets are marked with unobservable inputs, while staleness is a price-behavior measure of whether reported returns are autocorrelated. A fund holding 80 percent private equity can mark its book down promptly (high Level 3, low staleness), while a bond fund priced from a matrix can smooth returns with zero Level 3 exposure. The two are correlated only weakly and non-monotonically in the cross-section, return autocorrelation across Level 3 terciles runs 0.013, -0.073, 0.178, and within funds over time changes in Level 3 holdings do not move any staleness proxy (Table B.5, Panel A): the cross-sectional association is entirely a time-invariant fund characteristic absorbed by fixed effects. During the crisis the staleness interaction is negative while the opacity interaction is positive, consistent with Zeng's sticky-NAV hypothesis that infrequently marked prices dampen first-mover incentives; a mechanical-displacement alternative, in which stale funds merely record crisis redemptions late, predicts a reversal that does not occur (the staleness interaction at $t{+}1$ is -0.642 and at $t{+}2$ is -0.223). Because the staleness interaction is not robust to two-way clustering we treat its sign as suggestive; the load-bearing point is that controlling for staleness raises the opacity coefficient rather than lowering it.

The amplification is specific to Level 3. Estimating separate COVID interactions for each hierarchy level, with each share used as the sole treatment because the three sum to roughly 100 percent, there is no detectable effect for Level 1 (+0.003, $p = 0.24$) or Level 2 (-0.003, $p = 0.23$), while Level 3 carries the entire effect (+0.113, $p = 0.055$, or +0.17 percentage points per standard deviation; Table A.5). Crisis amplification attaches to assets priced with unobservable inputs, not to less-liquid pricing tiers generally. The same force appears across fund structures on different margins: in the capped vehicles, interval and tender-offer redemptions are mechanically limited by the repurchase offer, so their realized-redemption interactions are near zero by construction (the within-fund $\hat{\gamma}_3 \approx +0.0002$ of Regime 2), and the demand-side evidence for them is the oversubscription and survival analysis of Section B rather than a realized-flow coefficient; BDCs are excluded because their near-constant Level 3 share leaves no identifying variation.

Economic Magnitudes. Opacity in the open-end universe is concentrated in a thin tail: 78.4 percent of fund-quarters report zero Level 3 exposure, the 90th percentile is 0.12 percent, and the 99th percentile is 4.4 percent. At the headline estimate, a fund at the 99th percentile experiences crisis amplification of about $0.118 \times 4.4 \approx 0.5$ percentage points of NAV per quarter, roughly 12 percent of the 4.2 percent mean redemption rate, and a one-standard-deviation increase in lagged opacity raises crisis redemptions by 0.17 to 0.27 percentage points (Table A.6), or 4 to 7 percent of the mean. Within the open-end universe the effect is thus real but modest, and modest by regulatory design, since Rule 22e-4 caps illiquid assets at 15 percent of net assets and registered open-end funds simply cannot reach destabilizing opacity levels. The economic significance lies at the extrapolation margin: the semi-liquid vehicles that retirement-plan reform would bring to retail investors carry unobservable-input shares averaging 33 percent (Section D), ten times the open-end 99th percentile, at which the same per-unit response implies amplification near 3.9 percentage points per quarter, on the order of the entire baseline redemption rate, arriving through vehicles whose quarterly caps are already chronically oversubscribed. For comparison, Goldstein et al. (2017) find that a one-standard-deviation decrease in portfolio liquidity predicts a 1.9-percentage-point redemption increase during stress, so the open-end opacity effect is an order of magnitude smaller per standard deviation while the implied semi-liquid effect is of comparable size.

Stress-Episode Heterogeneity. Whether opacity amplifies redemptions in any market stress or only in information-driven crises is itself a test of the mechanism, so we estimate the opacity interaction jointly across three episodes that differ in the nature of the shock: COVID (2020Q1–Q2, Knightian valuation uncertainty), the 2022 Federal Reserve tightening (the fastest rate increases since the 1980s and a 20 percent equity drawdown, but transparent repricing), and the BREIT gating episode (2022Q4–2023Q2, stress concentrated in non-traded real estate). Amplification is strongest in COVID (+0.131, $p = 0.034$), present at roughly half that magnitude in the 2022 tightening (+0.061, $p = 0.028$), and undetectable around the BREIT episode (+0.032, $p = 0.27$, incremental to the overlapping tightening window); a Wald test cannot formally separate the COVID and Fed coefficients ($p = 0.26$), so the ordering is a point-estimate pattern rather than a statistical ranking (Table A.3). A cross-shock horse-race on the clean panel anchors this reading: among six crisis windows, COVID is the only episode whose opacity gradient survives the headline controls, two-way fund-by-quarter clustering, and a frozen pre-crisis opacity measure, and in every episode only the Level 3 unobservable-input proxy amplifies, while the Level 2, Level 1, and NAV-staleness proxies do not (Table B.4). Opacity-linked amplification is thus a stress-state phenomenon largest where valuation is genuinely uncertain: in 2022 repricing was observable and even Level 3 holdings had public benchmarks, so the effect is roughly half-sized, whereas during COVID the value of hotels, office REITs, and airlines was unknowable in real time. The attenuation reflects the transparency of the shock, not its size, since 2022 redemptions (mean 12.1 percent) exceeded baseline (10.2 percent) by 1.9 percentage points.

Retail versus Institutional Heterogeneity. If opacity-driven redemptions reflected naive behavior, mistaking opacity for risk, they should concentrate among retail investors; if instead opacity creates genuine coordination problems, even sophisticated investors should redeem from opaque funds during crises (Benartzi and Thaler, 2001). We proxy retail orientation by fund size, classifying funds in the bottom tercile of total net assets as retail-oriented, motivated by the facts that small funds disproportionately serve individuals and charge higher expense ratios (mean 1.18 percent against 0.74 percent for large funds; Pástor and Stambaugh, 2012). Estimating a triple-difference on the canonical panel (equation (21)), with the opacity-COVID interaction and its retail triple interaction, the main interaction is positive and consistent with the headline, $\hat{\beta}_4 = +0.103$ (SE $= 0.068$, $p = 0.133$, $N = 193{,}668$), while the triple interaction is small, negative, and far from significant, $\hat{\beta}_7 = -0.031$ (SE $= 0.087$, $p = 0.725$): there is no evidence that retail-oriented funds exhibit differential opacity amplification.

The null heterogeneity result rules out the naive-investor explanation, since even large institutional funds with professional allocators redeem from opaque portfolios during crises, which strengthens the systemic-risk argument, because a coordination problem that sophisticated investors cannot solve cannot be mitigated through investor education. It also bears on the retailization debate: the Department of Labor's 2020 guidance assumed that retail investors, intermediated through fund structures, would behave like institutions (Vanguard Group, 2020), and our findings validate that assumption in one sense (the two investor types respond similarly to opacity) while challenging it in another (both exhibit fragility, so retailization raises systemic risk in proportion to the assets it brings in). Robustness to alternative retail proxies based on expense ratio and fund age is reported in Appendix Table IA.4.

Timing and the Limits of Pre-Trend Testing. A dynamic event study would ordinarily validate the parallel-trends assumption through pre-COVID leads, but monthly N-PORT reporting begins only in 2019Q4, one quarter before the COVID onset, so the open-end panel contains a single pre-treatment quarter and a leads-based pre-trend test is structurally unavailable. We therefore validate the timing and specificity of the effect through the placebo and stress-episode evidence already reported: the opacity interaction is largest under COVID-era valuation uncertainty, attenuates by half in the transparent 2022 tightening, and lies in the upper tail of a non-crisis placebo distribution (Section VI, where the COVID estimate exceeds 18 of 21 placebo onsets). The cross-sectional difference-in-differences in Table A.2 captures the average treatment effect over the 2020Q1–Q2 crisis window, and the mechanism predicts that the effect resolves as valuation uncertainty does, which the decay of the episode coefficients from COVID to the later windows is consistent with.

Summary of Main Results. Taken together, the results establish that portfolio opacity is a distinct fragility channel. The additivity test shows that the opacity gradient is essentially unmoved when a genuine cash buffer receives its own term, with a persistence ratio between 0.92 and 0.99 across opacity proxies, so opacity does not proxy for illiquidity. The open-end difference-in-differences put the opacity-COVID interaction at 0.104 to 0.118, significant at the 10 percent level under fund clustering and at the 5 percent level under two-way clustering: a ten-percentage-point increase in Level 3 exposure raises crisis-quarter redemptions by roughly one percentage point within fund. The effect is modest inside the open-end universe, where Rule 22e-4 bounds opacity, but large at the 33 percent unobservable-input shares of the semi-liquid vehicles that retailization would expose, where the opacity-attributable component of a COVID-scale shock approaches the entire baseline redemption rate. Amplification scales with the informational severity of the stress episode, largest under COVID and half-sized in the transparent 2022 tightening, and shows no differential response across retail and institutional investors, indicating structural rather than behavioral fragility. The hierarchy test localizes the effect to assets priced with unobservable inputs, with no detectable Level 1 or Level 2 interaction, and the interval-fund evidence shows the same force on entirely different redemption mechanics, where high-opacity funds reach their first binding redemption cap more than twice as fast as transparent peers. The one qualification the analysis adds is that NAV staleness stabilizes rather than amplifies during crises, so that controlling for it isolates and strengthens the informational component of opacity.

VI Robustness Checks

We report two parallel robustness exercises. The primary exercise covers the semi-liquid fund panel (interval and tender-offer funds) using oversubscription as the dependent variable, the appropriate fragility measure when redemptions are mechanically capped. The secondary exercise covers the open-end panel and confirms that the opacity-fragility relationship survives alternative clustering, fixed effects, opacity thresholds, sample restrictions, estimators, and a battery of falsification tests.

The Semi-Liquid Panel. Because redemption rates in semi-liquid funds are endogenously capped at the quarterly offer percentage (typically 5 to 20 percent of NAV), the correct fragility measure is oversubscription, a binary indicator equal to one when redemption requests exceed the offer cap, which captures the excess investor demand the fund structure cannot accommodate. For this primary dependent variable (an oversubscription indicator, mean 14.7%, $N = 6{,}592$ fund-quarters across 206 funds) we estimate

$$\text{Oversubscribed}_{it} = \alpha_i + \gamma_t + \beta_1 L3_{it-1} + \beta_2 \text{Crisis}_t + \beta_3\, L3_{it-1} \times \text{Crisis}_t + X'_{it-1}\delta + \varepsilon_{it} \quad (22)$$

across 81 specifications, with standard errors clustered by fund throughout. The specifications span four dimensions: twelve alternative shock and window combinations, pooled and by subgroup; five opacity proxies (Level 3 share carried forward, NAV staleness from a Ljung-Box Q^2 statistic, direct-lending share, private-asset share, and holdings concentration); eight sample cuts (full sample, interval-only, tender-only, excluding the top and bottom TNA deciles, excluding the first four quarters, post-2019, and high-offer-rate funds); and six estimator variants (fund-only, quarter-only, and no fixed effects, full controls, nonlinear opacity, and triple-difference). Across these 81 specifications $\hat{\beta}_3$ is positive in 65.4% of cases and positive with $p < 0.10$ in 17, for an overall robustness ratio of 21.0% (Table VI.1). The evidence is strongest among interval funds, where the estimator variants reach a 57.1% robustness ratio, and weakest among tender-offer funds, which are uniformly null. That contrast is interpretable rather than discouraging: tender-offer boards set the offer cap at their discretion, and often cut it precisely when redemption pressure rises, which creates a different binding mechanism than the fixed quarterly window of an interval fund. The 21.0% ratio reflects genuine statistical uncertainty in a hard environment, a binary outcome in a panel of 6,592 fund-quarters, rather than an absence of signal: the consistently positive sign and the concentration of significance in the interval subsample both line up with the opacity-fragility channel. A continuous secondary measure, the gate-utilization rate, converges in only two specifications because mechanical binding at the offer cap leaves too little within-fund variation to estimate.

The Open-End Panel. On the open-end panel ($N \approx 193{,}700$ fund-quarters) the headline interaction is $\hat{\beta}_3 = 0.104$ ($p = 0.061$; Table A.2), and a ladder of ten checks confirms that it is stable in sign and magnitude across the choices an empiricist must make. Clustering matters little: fund-only clustering gives $\hat{\beta}_3 = 0.086$ ($p = 0.023$) and quarter-only clustering gives the same point estimate at $p = 0.012$, with the two-way fund-by-quarter benchmark the most conservative. The result is a within-fund phenomenon: dropping fixed effects entirely attenuates the estimate to 0.054 ($p = 0.160$), fund fixed effects alone reproduce the benchmark at 0.090 ($p = 0.026$), and time fixed effects alone leave it indistinguishable from zero (0.028, $p = 0.404$). Because more than three-quarters of open-end fund-quarters report zero Level 3 holdings, an above-median opacity indicator captures any Level 3 exposure at all and enters at 1.12 ($p < 0.001$), a concentrated distribution that motivates the continuous form as the headline specification. The effect is not a size artifact, surviving exclusion of the largest ten percent of funds (0.076, $p = 0.027$) and of the smallest ten percent (0.070, $p = 0.059$); it is crisis-specific, weakening in a pre-2022 subsample with less power (0.045, $p = 0.115$) and vanishing mechanically after 2021 when there are no COVID observations; and it is concave, since adding a squared opacity term yields a positive linear coefficient (0.246, $p = 0.002$) and a small negative quadratic (-0.003, $p = 0.003$), implying diminishing marginal amplification at very high opacity.

Three further checks probe the interpretation rather than the fit. The first asks whether the result reflects unobservable inputs specifically or broad non-Level-1 exposure: substituting the Level 2 share (observable-but-not-quoted inputs such as dealer quotes and matrix prices) as the sole treatment yields $\hat{\beta}_{\mathrm{L2}} = -0.005$ ($p = 0.071$), economically zero and, if anything, of the opposite sign, exactly as the ASC 820 reading predicts, since dealer-quoted prices create far less valuation uncertainty than pure manager discretion. The second addresses anticipation: replacing lagged opacity with each fund's frozen 2019Q4 Level 3 share, immune to any COVID-period reclassification, leaves the standardized effect intact and slightly larger (a +0.27 percentage-point response per one-standard-deviation increase in opacity versus +0.19 under the time-varying measure; Table A.6), so the result is not driven by contemporaneous portfolio adjustment. The third tests the stable-unit-treatment assumption by adding a family-leave-one-out opacity measure, the average Level 3 share of the other funds in the same sponsor. Here the own-fund interaction stays positive and, at 0.126, close to the headline, but it loses precision ($p = 0.417$) once the collinear family measure enters, while family opacity itself is positive and significant (0.437, $p = 0.036$, across 669 families). We read this as enriching rather than overturning the individual-fund interpretation: own-fund opacity retains its magnitude, but the coordination mechanism also has a fund-family dimension, with investors appearing to update beliefs about a sponsor's funds collectively. Strict SUTVA therefore does not hold, and the family channel is a feature of the opacity mechanism rather than a confound to it.

Falsification. Two falsification exercises ask whether opacity, rather than something correlated with it, drives the result. The first walks the ASC 820 hierarchy. If coordination failures are informational, transparent holdings should not amplify fragility, and they do not: the Level 1 (exchange-quoted) share interacted with COVID enters at $\hat{\beta}_{\mathrm{L1}} = 0.004$, roughly one twenty-sixth of the Level 3 coefficient and economically negligible, while the Level 2 share is essentially zero. The monotone pattern across the hierarchy, $\hat{\beta}_{\mathrm{L3}} = 0.104 \gg \hat{\beta}_{\mathrm{L2}} \approx \hat{\beta}_{\mathrm{L1}} \approx 0$, is difficult to reconcile with an omitted variable that loads on illiquidity or fund risk generally, since those would not respect the valuation-observability ordering.

The second exercise is a placebo on the timing of the shock. The open-end panel begins in 2019Q4, when monthly N-PORT reporting became available, so a placebo window in 2016 to 2019 is infeasible; we instead treat each of the 21 non-COVID quarters in the sample as a counterfactual onset and re-estimate the core specification. The actual COVID estimate of 0.104 lies in the upper tail of the resulting placebo distribution, exceeding 18 of the 21 placebo onsets (empirical one-sided $p = 0.143$); it is large relative to almost all placebos but not the single largest, the appropriate characterization of a single-crisis design. A complementary battery of designated-shock placebos on the semi-liquid panel (Table VI.1 and the Internet Appendix) tells the same story: the COVID interaction is positive while genuine non-events, a 2019Q4 trade-tension window and a stable 2021Q2 quarter, are essentially zero or negative, and other stress episodes (the 2022 Fed tightening and the BREIT gate) are positive but smaller. Taken together the falsification tests support a COVID-specific, opacity-specific reading of the main result without overstating the precision a one-crisis identification can deliver.

Table VI.1: Robustness Universe: Semi-Liquid Panel ($\mathbf{1}$(Oversubscribed) as DV)

Category	**Subgroup**	**Specs**	**Pos.** $\hat{\beta}_3$	$p < 0.10$	**Rob. Ratio**
Panel A: Liquid Open-End Funds — DV: Redemption Rate (%); N ≈ 193,793					
Time-period variants	—	12	—	—	*Pending*
Opacity measures	—	6	—	—	*Pending*
Sample restrictions	—	8	—	—	*Pending*
Specification alts.	—	10	—	—	*Pending*
Panel B: Semi-Liquid Funds — DV: $\mathbf{1}$*(Oversubscribed); N = 6,592; 206 funds*					
Redemptions are capped by gate/offer pct; oversubscription captures excess demand					
Time-period variants	Pooled	12	8/12	2/12	16.7%
	Interval	12	11/12	2/12	16.7%
	Tender	12	4/12	0/12	0.0%
Opacity measures	Pooled	5	3/5	1/5	20.0%
(valuation, pricing,	Interval	5	3/5	2/5	40.0%
credit, PE, conc.)	Tender	5	3/5	0/5	0.0%
Sample restrictions	Pooled	8	7/8	2/8	25.0%
Specification alts.	Pooled	8	5/8	6/8	50.0%
	Interval	7	4/7	4/7	57.1%
	Tender	7	5/7	0/7	0.0%
All semi-liquid (pooled)	All	81	65.4%	23.5%	**21.0%**

Notes: Robustness ratio = fraction of specifications with $\hat{\beta}_3 > 0$ and $p < 0.10$. **Primary DV:** $\mathbf{1}$(Oversubscribed) = `ncsr_any_oversubscribed` (binary LPM; full coverage 6,592 obs; mean 14.7%). Actual redemptions are mechanically capped by each fund's quarterly offer percentage (typically 5–20%); oversubscription captures excess investor demand not met by the fund, which is the correct fragility measure for semi-liquid funds. Opacity proxies (Panel B, opacity row): L3% LOCF (valuation opacity), NAV staleness (pricing opacity), direct lending % (private credit), private assets % (PE/real), portfolio HHI (concentration). Sample restrictions SR1–SR8: full sample, interval only, tender only, excl. top/bottom TNA decile, excl. first 4 qtrs, post-2019, high-offer-rate funds. **Secondary DV:** gate utilization rate (continuous; 13% coverage) run in parallel across all specs; only N=2 converged (insufficient variation at 13% coverage — consistent with mechanical binding at offer cap leaving little residual variation). Specification alts. include: fund-only FE, quarter-only FE, no FE, full controls, non-linear opacity, and triple-diff. SE clustered by fund. All continuous variables winsorized at 1%/99%.
(2026-03-15 08:10:51).

Summary. Across both fund types the opacity-fragility relationship is directionally robust. On the open-end panel the interaction is positive and stable across the clustering, fixed-effects, threshold, size, and estimator variations of the ten-check ladder, and the falsification tests place the COVID estimate in the upper tail of a placebo distribution while transparent holdings show no amplification. On the semi-liquid panel the interaction is positive in 65.4% of 81 specifications. The lower significance rate in the semi-liquid universe reflects a harder identification environment, a smaller panel ($N = 6{,}592$ against 193,793), a binary dependent variable, and a gate-binding mechanism that partially censors the very demand signal oversubscription is meant to recover. The one substantive qualification the robustness analysis adds to the headline is that opacity operates partly at the fund-family level, a dimension we treat as part of the coordination mechanism rather than a threat to it.

VII Policy Implications and Welfare Analysis

The empirical findings raise an immediate policy question: if portfolio opacity amplifies redemption fragility during crises, how should regulators respond? The current framework offers no answer, because neither the SEC's liquidity rules nor its anti-dilution pricing tools condition on valuation opacity. Rule 22e-4 governs illiquidity, and since the July 2023 money market reforms the SEC has relied on mandatory liquidity fees in place of the now-repealed redemption gates, but none of these instruments distinguishes an opaque portfolio from a transparent one. The question is becoming urgent because the regulatory opening of 401(k) plans to private-market funds, set in motion by the Department of Labor's Information Letter of June 3, 2020 and accelerated by Executive Order 14330 of August 7, 2025 and the DOL's proposed fiduciary safe-harbor rule of March 30, 2026, would extend the fragility channel we document to tens of millions of retail retirement savers. This section addresses the question through a calibrated welfare framework that compares redemption gates, swing pricing, and an opacity-contingent hybrid.

The Regulatory Gap. The SEC's Rule 22e-4, finalized in October 2016 and effective December 2018, requires every registered open-end fund to operate a liquidity risk management program: funds classify holdings into liquidity buckets and limit illiquid investments to no more than 15% of net assets.[18] For money market funds, the 2014 reform of Rule 2a-7 introduced explicit redemption gates that let a board suspend redemptions for up to ten business days when weekly liquid assets fell below 30% of net assets; the July 2023 reforms repealed those gates, severed the link between liquid-asset thresholds and fees, and adopted mandatory liquidity fees for institutional prime and tax-exempt funds.[19] The regulatory trajectory has thus moved away from quantity restrictions and toward anti-dilution pricing, which sharpens rather than blunts the question this section poses, because neither the abandoned gate regime nor the current fee regime conditions on the opacity margin we identify. A census of gate-related events confirms that formal gating remains rare while near-gate stress is widespread: 30.4% of funds experienced weekly redemptions exceeding 2% of NAV during COVID, and high-opacity funds activated gates at twelve times the rate of low-opacity funds, yet no jurisdiction currently links a gate trigger to portfolio opacity.

The anti-dilution alternative, swing pricing, was authorized alongside Rule 22e-4 under amendments to Rule 22c-1. Swing pricing adjusts the fund's NAV downward when net redemptions exceed a threshold, typically 2% of assets, charging exiting investors a dilution fee of one to two percent of redemption value and so internalizing the liquidity cost of exit without suspending redemptions; it is widely used in European funds but remains rare in the United States. The common feature of all of these instruments is that they react to illiquidity or to realized flows: the repealed gate trigger keyed on weekly liquid assets, the liquidity fees that replaced it key on flows and liquidation costs, and Rule 22e-4's classification ignores opacity entirely. A fund holding 40% cash and 60% Level 3 private equity, highly liquid but highly opaque, would trigger none of these safeguards, and because opacity amplifies redemptions before any flow threshold is reached, fee-based mechanisms remain

[18]SEC Release No. IC-32315, *Investment Company Liquidity Risk Management Programs*, October 13, 2016. Codified at 17 CFR §270.22e-4.

[19]SEC Release No. IC-31166, *Money Market Fund Reform*, July 23, 2014 (gates introduced); SEC Release No. IC-34959, *Money Market Fund Reforms*, July 12, 2023 (gates removed; mandatory liquidity fees adopted). Non-money-market open-end funds have never had comparable gate authority under Section 22(e) of the Investment Company Act, which requires redemption within seven days of tender.

reactive rather than preventive.

The retailization of opaque assets makes the gap consequential. Access to private-market funds inside defined-contribution plans has opened in three steps. First, the Department of Labor's Information Letter of June 3, 2020 clarified that 401(k) fiduciaries may include private equity exposure within professionally managed asset-allocation vehicles, subject to ERISA prudence.[20] Second, Executive Order 14330 of August 7, 2025 directed regulators to facilitate access to alternative assets in DC plans, and the DOL rescinded its cautionary 2021 supplemental statement days later. Third, the DOL's proposed rule of March 30, 2026 would establish a process-based fiduciary safe harbor whose enumerated evaluation factors, performance, fees, liquidity, valuation, benchmarks, and complexity, map directly onto the fragility margins this paper documents. Our findings inform this debate in three ways. The first is a new measured fact: scanning every target date fund in N-PORT (756 funds, 12,708 fund-quarters, 2019Q4 through 2025Q2) and matching their holdings to the universe of 195 opaque interval and tender-offer funds by both name and LEI, we find no detectable opaque-sleeve exposure in any quarter: across \$35.7 trillion of holdings scanned, the largest single opaque position is 0.00% of any fund's net assets (CIT-based target date funds fall outside N-PORT and are not covered by this statement). The 401(k) transmission channel is therefore prospective rather than realized, which is precisely when ex-ante analysis is most useful. The second is that the within-fund retail heterogeneity test (Section G) finds no differential opacity response across investor types, so the fragility channel is structural rather than behavioral and retailization operates by adding exposed investor mass rather than a more flighty investor type. The third is the welfare comparison below, which quantifies the cost of alternative policy regimes for the population that retailization would actually expose.

[20]DOL Information Letter, June 3, 2020, issued to Groom Law Group on behalf of Pantheon Ventures and Partners Group. The guidance contemplates private-market exposure inside diversified vehicles such as target date funds, not direct menu options.

A Calibrated Welfare Framework. We evaluate three regimes, no intervention, redemption gates capping quarterly outflows at 5%, and swing pricing charging a 2% fee when redemptions exceed 2%, in a simulation of 10,000 funds facing a COVID-like shock. Because open-end funds hold almost no Level 3 assets (a 99th percentile of 4.4%), a simulation calibrated to the open-end opacity distribution would understate the exposure relevant to retailization; we therefore draw opacity from the semi-liquid vehicle class that retirement menus would actually access, whose mean unobservable-input share is 33% (Section D), and apply the redemption response estimated in Regime 1. Each fund's crisis redemption rate follows

$$\text{Redemption}_i = \alpha + \beta_1 \, \text{Opacity}_i + \beta_2 \, \text{COVID} + \beta_3 \, (\text{Opacity}_i \times \text{COVID}) + \gamma \, \text{Return}_i + \varepsilon_i, \quad (23)$$

calibrated from the clean-panel estimates of Table A.2: a baseline quarterly redemption rate of $\alpha = 4.2\%$, no aggregate COVID intercept ($\beta_2 \approx 0$, since the crisis-period mean of 3.8% is statistically indistinguishable from the 4.2% baseline and the crisis reallocates redemptions toward opaque funds rather than raising their aggregate level), and an opacity-COVID slope $\beta_3 \in [0.104, 0.118]$, with the return control and idiosyncratic noise set to the empirical dispersion of quarterly redemptions. Under this calibration a semi-liquid fund at the 33% class mean faces crisis redemption demand of roughly $4.2 + 0.118 \times 33 \approx 8.1\%$ of NAV per quarter, modest by open-end norms but 1.6 times the 5% repurchase cap that governs interval-fund liquidity, implying chronic oversubscription under stress. To represent retailization we scale investor mass by 15.5%, the increase implied if private funds were added to all DC menus,[21] which raises the baseline intercept to about 4.9% while preserving the opacity interaction, consistent with the null heterogeneity finding that retailization adds

[21] This parameter is calibrated from Vanguard's 2020 participant study, which shows 8.2% of participants traded during COVID versus 5.7% at baseline, a 43% relative increase; we conservatively assume 401(k) investors exhibit 60% of this sensitivity, translating to a 15.5% higher redemption rate. The assumption is conservative relative to Benartzi and Thaler (2001), who document that retail investors exhibit two to three times the flow-performance sensitivity of institutional investors.

exposed mass rather than a more flighty investor type.

We measure welfare with a quadratic loss that penalizes binding redemption constraints, dilution fees, and the dilution externality itself,

$$W_i = -\left[\lambda_1 \left(\text{Redemption}_i^{\text{desired}} - \text{Redemption}_i^{\text{actual}}\right)^2 + \lambda_2 \left(\text{Fee}_i\right)^2 + \lambda_3 \, \text{Dilution}_i\right], \tag{24}$$

where actual redemptions are capped by gates or reduced by swing fees, the dilution externality $\text{Dilution}_i = \max\Big(0,\ \kappa \left(\text{Opacity}_i/100\right) \text{Redemption}_i^{\text{actual},2} - \text{Fee}_i \, \text{Redemption}_i^{\text{actual}}\Big)$ is the value extracted from remaining investors when redemptions occur at a stale NAV (rising in opacity and internalized by any fee charged), and the weights $\lambda_1 = 1$, $\lambda_2 = 0.5$, $\lambda_3 = 1$ reflect that investors dislike binding constraints somewhat more than monetary fees.[22] The three regimes act on this loss through different margins. A uniform gate caps quarterly redemptions at 5%, the threshold typical of interval funds and proposed by the SEC for semi-liquid funds,[23] and binds only on tail-event demand: a semi-liquid fund at the 33% class mean, with 8.1% desired redemptions, has roughly 38% of its desired outflows rationed, whereas a transparent peer at the open-end 99th percentile faces demand near 4.7% and is not rationed at all, so a gate's welfare cost concentrates in the high-opacity tail. Swing pricing instead charges every fund whose redemptions exceed the low 2% threshold (a 2% fee at a redemption-demand elasticity of -1.5 trims outflows by about three percentage points),[24] so it taxes opaque and transparent funds alike and spreads its welfare cost broadly rather than targeting the funds that actually bind. The opacity-contingent hybrid leaves the uniform 5% gate in place but adds an anti-dilution fee scaled to fund opacity, charged on the redemptions of the most opaque funds so that exiting investors internalize the dilution they

[22]The weight ratio $\lambda_1/\lambda_2 = 2$ is calibrated to revealed preferences from interval fund studies (Jin and Scherbina, 2011), in which investors accept roughly 1% annual fees to avoid 2% quarterly redemption restrictions. The dilution intensity $\kappa = 1$ is a calibration parameter; the regime ranking below is robust across $\kappa \in \{0.5, 1, 2\}$.

[23]SEC Release No. IC-34432 (November 2021) proposes limiting tender-offer funds to 5% quarterly redemptions unless the fund holds 10% weekly liquid assets.

[24]The elasticity of -1.5 is calibrated from corporate bond fund studies (Goldstein et al., 2017); the swing parameters of a 2% threshold and 2% fee match European fund practice (Jin and Scherbina, 2011).

impose, while transparent funds remain fee-free. Solving the calibrated model over 10,000 funds yields a clear ranking (Table A.13): total welfare loss is 22% lower under uniform gates than under swing pricing (690,141 versus 889,739 normalized units), and a further 5% lower under the anti-dilution hybrid (658,944). The ranking is instructive about mechanism. Gates dominate swing pricing because they bind only on tail-event demand, whereas fees tax all excess redemptions above a low threshold. But tightening the gate further on opaque funds, the natural reading of an opacity-contingent rule, does not help: in the same calibration a tiered-gate hybrid raises welfare loss 11% above uniform gates, because the additional unmet-liquidity cost it imposes on opaque funds exceeds the dilution it saves. What delivers the welfare gain is not a tighter quantity cap but a price, an opacity-scaled anti-dilution fee that internalizes the externality where it is largest while leaving the fire-sale cap uniform. The ranking is robust across dilution intensities (the hybrid's advantage holds for all $\kappa \geq 0.5$ and both $\beta_3 \in \{0.104, 0.118\}$), though the precise magnitudes are calibration-dependent.

Policy Recommendations. The framework supports three reforms. First, the SEC should make redemption safeguards opacity-contingent, but through price rather than tighter quantity caps. The welfare analysis shows that tightening gates on opaque funds raises welfare loss, whereas an opacity-scaled anti-dilution fee lowers it, so the SEC should retain the uniform redemption gate as a fire-sale cap and pair it with an anti-dilution fee that rises with a fund's Level 3 share, charged on the redemptions of the most opaque funds so that exiting investors internalize the dilution they impose while transparent funds remain fee-free. This is a rule amendment rather than new legislation, because funds already report Level 3 shares in N-PORT and the post-2023 money market reforms already establish anti-dilution liquidity fees as the SEC's preferred tool, into which an opacity scaling can be inserted. Second, because N-PORT is published with a roughly sixty-day lag, investors observe stale opacity precisely when valuation uncertainty is highest; the SEC should require more frequent disclosure of Level 3 shares during stress periods, defined for example by a VIX above 30 or corporate bond spreads beyond 200 basis points, mirroring the intensified weekly reporting that already applies to money market funds under stress.[25] Faster disclosure both reduces the information asymmetry that drives coordination-based redemptions and imposes market discipline on opacity accumulation, and our event study (Section H) shows the opacity effect peaks in 2020Q2 and decays by 2021Q2, consistent with transparency accelerating the resolution of valuation uncertainty. Third, the anti-dilution fee should scale continuously with opacity rather than switching at a single threshold, set to approximate the dilution a redemption imposes as the fund's Level 3 share rises, so that it preserves a price signal on marginal redemptions that a pure gate, by pro-rating equally across investors, would otherwise suppress, and lowers the ex-ante incentive to front-run the gate.

Limitations and Open Questions. Several caveats bound these conclusions. The calibration rests on a single crisis episode, so a COVID-specific opacity response would imply a fee schedule tuned too tightly for future crises; the regime ranking is, however, robust to plausible variation in $\hat{\beta}_3$ and to the dilution intensity, and the margin that motivates the opacity-scaled fee is the elevated redemption demand of the semi-liquid vehicles rather than the open-end coefficient itself. The simulation also assumes managers do not strategically lower Level 3 holdings to qualify for a looser tier, an assumption supported by the fact that opacity largely reflects portfolio mandate rather than discretionary reporting and that interval funds facing permanent caps show no systematic under-reporting of illiquidity. Cross-border competition with European funds that use swing pricing is unlikely to undercut the hybrid, since it improves investor welfare relative to swing pricing and recent EU reviews have questioned swing pricing's effectiveness during COVID. Finally, our retailization calibration assumes gradual adoption; faster adoption would raise the welfare cost roughly in proportion to the additional exposed mass, so regulatory vigilance is warranted if retailization accelerates. The same calibration extends to interval-fund oversubscription as 401(k) participation scales, under voluntary, forced-exit (hardship and required-minimum-distribution), and macro-stress scenarios; we report this extension in the Internet Appendix,

[25] This crisis-contingent disclosure mirrors the SEC's existing framework for money market fund weekly reporting under Rule 2a-7, which intensifies when weekly liquid assets fall below 30%.

where the figures are illustrative simulations calibrated to COVID-era 401(k) withdrawal behavior (Vanguard Group, 2020) rather than econometric estimates.

Three questions remain open, each warranting its own agenda. The first is distributional: opacity-contingent gates and disclosure impose compliance costs that may fall on investors through higher expense ratios, on managers through narrower margins, or on regulators through enforcement complexity, and because our framework measures investor welfare rather than manager profits, the pass-through of compliance costs to fees, analogous to the fee-performance literature for open-end funds (Pástor and Stambaugh, 2012), must be understood before opacity-contingent rules can be fully assessed. The second is whether the illiquidity premium that justifies retailization survives the exit risk we document: private assets are argued to earn three to five percent annually above liquid benchmarks, but that premium is contingent on not exiting at the wrong time, and a lifecycle welfare analysis incorporating endogenous crisis-exit risk is needed to determine whether it remains positive net of distressed-exit costs for 401(k) participants. The third is whether real-time valuation disclosure is practical: N-PORT already collects fair-value-level data, but more frequent updates may not improve accuracy when the underlying assets are genuinely illiquid, may invite smoothing of Level 3 marks, and raise governance questions about the optimal scope of disclosure, daily versus weekly, crisis-only versus always-on, and position-level versus fund-level.

Taken together, the analysis does not support prohibiting 401(k) access to opaque funds, but it does support conditional access: retailization should proceed in tandem with opacity-contingent safeguards, so that broader retail exposure does not destabilize funds whose fragility comes from valuation opacity rather than from liquidity mismatch. The recommendations are implementable within the SEC's existing infrastructure, requiring only that opacity thresholds be added alongside the liquidity classifications Rule 22e-4 already maintains.

VIII Conclusion

This paper provides evidence that portfolio valuation opacity amplifies redemption fragility during systemic crises, using ASC 820 Level 3 fair-value disclosures from SEC N-PORT filings. The headline difference-in-differences estimate is $\hat{\beta}_3 = 0.118$ ($N \approx 193{,}700$): each percentage point of lagged Level 3 exposure amplifies COVID-quarter redemptions by about 0.12 percentage points within fund, significant at the 10 percent level under fund clustering and at the 5 percent level under two-way fund-by-quarter clustering. The effect is specific to Level 3, with no detectable Level 1 or Level 2 interaction; distinct from NAV staleness, controlling for which raises rather than lowers the opacity coefficient; and present in attenuated form across stress episodes, scaling with their informational severity. Within the open-end universe the magnitude is modest, bounded by Rule 22e-4's 15 percent illiquid-asset cap, but the same per-unit response evaluated at the 33 percent mean unobservable-input share of semi-liquid vehicles implies amplification on the order of the entire baseline redemption rate, which is what makes the retailization of private markets the policy-relevant frontier.

The findings span three institutional regimes and converge on a single conclusion: opacity creates structural fragility across fund types. In the open-end panel, high-opacity funds redeem differentially more during COVID ($\hat{\beta}_3 = 0.118$), an estimate robust to a frozen pre-crisis denominator and to two-way clustering, with the amplification confined to Level 3; across opacity proxies the gradient is essentially unchanged when a genuine cash buffer receives its own term, a persistence ratio between 0.92 and 0.99, so opacity is separable from liquidity. The effect scales with the informational severity of the shock, strongest in COVID (+0.131), roughly half-sized in the transparent 2022 tightening (+0.061), and undetectable around the BREIT episode, and it shows no differential response by investor type, locating fragility in fund structure rather than investor behavior. In interval funds, which offer only periodic repurchases and so mechanically mute fire-sale externalities, opacity still drives the run on the gate: high-opacity funds reach their first binding redemption cap more than twice as fast as transparent peers (Cox hazard ratio 2.61) and are oversubscribed in 50 percent of repurchase offers against 22 percent, confirming that the mechanism operates independent of daily-liquidity mechanics. The within-interval retail-versus-institutional triple-difference is not estimable in the current EDGAR universe (Section C), but the open-end heterogeneity analysis shows a null triple interaction, so opacity-driven fragility is not behaviorally concentrated among retail investors. A direct measurement of how far retailization has progressed completes the picture: scanning every target date fund in N-PORT (756 funds, 12,708 fund-quarters) and matching to the 195 opaque interval and tender-offer funds by name and LEI, we find no detectable opaque-sleeve exposure in any quarter, so the 401(k) transmission channel is a policy under construction rather than a realized exposure, which is precisely the window in which the welfare analysis is most useful.

Contribution. This is the first paper to measure portfolio-level valuation opacity from regulatory fair-value-hierarchy disclosures and trace its effect on redemption fragility. Prior work on mutual fund runs focuses on liquidity mismatch (Chen et al., 2010; Goldstein et al., 2017; Morris et al., 2017), treating opacity as absent or collinear with illiquidity; we show that opacity is a separable channel (persistence ratio between 0.92 and 0.99 under a genuine cash control, and distinct from staleness) operating through investors' inability to verify NAV, and that it is structural rather than behavioral. The contribution touches three literatures. First, it extends the fund-run literature (Chen et al., 2010; Goldstein et al., 2017; Falato et al., 2021) by identifying opacity as a channel that survives controlling for liquidity, and it connects to the stale-pricing dilution literature (Choi et al., 2022; Chalmers et al., 2001) by supplying a regulatory opacity measure, a gated laboratory in which fire-sale dilution is mechanically muted, and a within-fund investor-type test, while distinguishing the Level 3 channel from the sticky-NAV stabilization of Zeng (2017). Second, it provides the first large-sample evidence that ASC 820 Level 3 classifications predict real economic outcomes, extending the fair-value-accounting literature (Song et al., 2010; Goh et al., 2015) beyond stock-price reactions to allocative consequences. Third, it informs the 401(k) retailization debate (Parker et al., 2023) on two margins: opacity affects retail and institutional investors alike, so the risk of retailization is scale rather than investor naivety, and the channel is prospective rather than active, so the welfare analysis evaluates a prospective regime that the DOL Information Letter, Executive Order 14330, and the 2026 proposed rule are now assembling.

Policy Implications. Because opacity is a fragility channel distinct from liquidity, effective regulation must address both, yet the SEC's current framework, Rule 22e-4's liquidity programs and, since the July 2023 reforms, mandatory liquidity fees in place of the repealed money-market gates, targets only the liquidity dimension; a fund holding 40 percent cash and 60 percent Level 3 private equity would trigger none of these safeguards. The welfare analysis (Section VII) evaluates opacity-contingent redemption design. Uniform gates reduce welfare loss by 22 percent relative to swing pricing, because gates bind only on tail-event redemptions while fees tax all excess redemptions above a low threshold; tightening gates further on opaque funds does not help, but pairing the uniform gate with an opacity-scaled anti-dilution fee lowers welfare loss a further 5 percent by internalizing dilution where it is largest, and mandatory real-time Level 3 disclosure during stress periods addresses the information-asymmetry mechanism directly. Implementation is feasible within existing infrastructure, since funds already disclose Level 3 shares in N-PORT and the post-2023 money market reforms establish anti-dilution liquidity fees as the SEC's preferred tool.

Limitations and Open Questions. Several limitations bound these conclusions and mark the open questions. The primary identifying variation comes from COVID-19; the 2022 tightening shows the channel recurs at smaller magnitude in a transparent-repricing episode, but the two coefficients cannot be formally separated, and historical crises (2008, the 1998 LTCM episode) with retrospective opacity proxies would sharpen the dose-response. Our Regime 1 measure captures realized rather than latent demand, which the Regime 2 oversubscription analysis only partially recovers; full measurement would require internal redemption-request data. We document opacity's effect on investor behavior without modelling the manager's endogenous response, and modelling the optimal opacity choice under redemption fragility is an open question. The N-CSR opacity extraction for interval funds achieves partial coverage given heterogeneous filing formats, and scaling to the full universe would require machine-assisted extraction. Finally, while opacity affects retail and institutional investors alike on average, heterogeneity along manager reputation, board governance, and valuation-methodology disclosure could identify best practices for opacity risk management.

The migration of opaque assets into liquid fund structures reflects durable trends in asset management: post-2008 yield-seeking, the shift of credit intermediation from banks to funds, and the democratization of alternatives through defined-contribution savings. This opacity is not costless: it amplifies redemption fragility during crises through an information-asymmetry channel that current liquidity regulation does not address, and because the two channels are distinct, regulators cannot neutralize opacity-driven fragility by tightening liquidity requirements alone. Just as banks face capital requirements proportional to asset risk, mutual funds accessible to retail retirement savers should face redemption safeguards proportional to portfolio opacity. Opacity-contingent design, a uniform gate paired with opacity-scaled anti-dilution pricing, together with real-time disclosure, can mitigate this fragility without prohibiting opaque asset classes or restricting retailization, by internalizing the redemption risk that valuation uncertainty creates. As the boundary between public and private capital markets continues to blur, the choice is between allowing opacity to accumulate unchecked and calibrating redemption safeguards to fund-specific transparency; this paper provides the evidence and a framework for the latter.

References

Acharya, V. V. and Merrouche, O. (2013). Precautionary hoarding of liquidity and interbank markets: Evidence from the subprime crisis. *Review of Finance*, 17(1):107–160.

Benartzi, S. and Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review*, 91(1):79–98.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

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.

Choi, J., Kronlund, M., and Oh, J. Y. J. (2022). Sitting bucks: Stale pricing in fixed income funds. *Journal of Financial Economics*, 145(2):296–317.

Correia, S. (2016). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. *Working Paper*. Stata command: `reghdfe`.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

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.

Ferreira, M. A., Keswani, A., Miguel, A. F., and Ramos, S. B. (2013). The determinants of mutual fund performance: A cross-country study. *Review of Finance*, 17(2):483–525.

Getmansky, M., Lo, A. W., and Makarov, I. (2004). An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics*, 74(3):529–609.

Goh, B. W., Li, D., Ng, J., and Ow Yong, K. O. (2015). Market pricing of banks' fair value assets reported under sfas 157 since the 2008 financial crisis. *Journal of Accounting and Public Policy*, 34(2):129–145.

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.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Jin, L. and Scherbina, A. (2011). Inheriting losers. *Review of Financial Studies*, 24(3):786–820.

Morris, S., Shim, I., and Shin, H. S. (2017). Redemption risk and cash hoarding by asset managers. *Journal of Monetary Economics*, 89:71–87.

Parker, J. A., Schoar, A., and Sun, Y. (2023). Retail financial innovation and stock market dynamics: The case of target date funds. *Journal of Finance*, 78(5):2673–2723.

Pástor, Ľ. and Stambaugh, R. F. (2012). On the size of the active management industry. *Journal of Political Economy*, 120(4):740–781.

Sirri, E. R. and Tufano, P. (1998). Costly search and mutual fund flows. *Journal of Finance*, 53(5):1589–1622.

Song, C. J., Thomas, W. B., and Yi, H. (2010). Value relevance of fas no. 157 fair value hierarchy information and the impact of corporate governance mechanisms. *The Accounting Review*, 85(4):1375–1410.

Spiegel, M. and Zhang, H. (2013). Mutual fund risk and market share-adjusted fund flows. *Journal of Financial Economics*, 108(2):506–528.

U.S. Securities and Exchange Commission (2016). Investment company reporting modernization. Final Rule, Release No. IC-32936. 17 CFR Parts 230, 232, 239, 249, 270, and 274.

Vanguard Group (2020). How america saves 2020: A report on vanguard defined contribution plans. Valley Forge, PA: Vanguard.

Zeng, Y. (2017). A dynamic theory of mutual fund runs and liquidity management. *Working Paper*.

A Tables and Figures

Table A.1: Summary Statistics

Panel A: Open-End Mutual Funds (Regime 1)

	N	Mean	SD	P25	P75
Redemption rate (%, winsorized)	197,722	4.03	19.67	−0.19	7.28
Level 3 (%, winsorized)	197,722	0.14	0.61	0.00	0.00
NAV staleness (ξ)	196,986	0.05	0.26	−0.12	0.23
log(TNA), winsorized	197,722	39.18	4.20	36.52	42.02
Return, $t-1$ (%)	196,577	1.67	7.86	−1.81	5.70
Number of funds	≈10,500				
Sample period	2019Q4–2025Q2				

Panel B: Semi-Liquid Funds (Regime 2)

	N	Mean	Median	SD	P75
Interval funds (70 funds)					
Oversubscribed (binary)	2,240	0.04	0.00	0.20	0.00
Level 3 (%)	2,045	29.6	24.4	23.1	48.9
Repurchase offer (% NAV)	1,365	9.3	5.0	5.8	12.5
Tender-offer funds (136 funds)					
Level 3 (%)	3,963	15.4	8.7	18.2	22.6
Repurchase offer (% NAV)	130	3.0	3.0	0.3	3.0
Total fund-quarters	6,592				
Sample period	2018Q1–2025Q4				

Panel C: Dual-Access Interval Funds (Regime 3)

The within-interval triple-difference requires interval-fund CIKs that file under both retail and institutional share-class structures. A data audit finds zero CIK overlap between the interval-fund universe (70 CIKs) and the open-end universe (15,709 series), because interval funds register as legally separate entities under Section 23(c)(3); the design is therefore not estimable from current filings. The retail–institutional evidence is the open-end heterogeneity test in Table A.9 (no significant differential, $\hat{\beta}_7 = -0.031$, $p = 0.725$).

Notes: Panel A covers open-end mutual funds; Panel B covers interval and tender-offer funds, with Level 3 forward-filled from semi-annual disclosures and oversubscription from N-CSR repurchase results. All continuous variables are winsorized at the 1st and 99th percentiles.

Table A.2: Portfolio Opacity and Crisis Redemptions: Difference-in-Differences Estimates

	(1) Base	(2) Size	(3) Staleness	(4) Two-way
Level 3_{t-1} × COVID	0.104*	0.113*	0.118**	0.118**
	(0.056)	(0.059)	(0.060)	(0.056)
Staleness × COVID			-0.375**	-0.375
			(0.154)	(0.500)
Level 3_{t-1}	0.047	0.049	0.050	0.050
	(0.054)	(0.056)	(0.056)	(0.065)
Log TNA$_{t-1}$ control	No	Yes	Yes	Yes
Staleness control	No	No	Yes	Yes
Fund FE, Quarter FE	Yes	Yes	Yes	Yes
Clustering	Fund	Fund	Fund	Fund × Qtr
Observations	193,668	193,667	193,667	193,667
R^2	0.591	0.592	0.592	0.592

Sample: open-end mutual funds, N-PORT canonical panel, 2019Q4–2025Q2. Dependent variable: quarterly redemption rate (%, winsorized at 1/99). Level 3_{t-1} is the lagged Level 3 share of total net assets (%). COVID equals one in 2020Q1–Q2. All specifications include fund and quarter fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Staleness is the Getmansky-Lo-Makarov return-smoothing index estimated from monthly N-PORT returns. Column (4) repeats column (3) with two-way (fund and quarter) clustered standard errors.

Table A.3: Opacity-Linked Redemption Amplification Across Stress Episodes

	(1) Joint specification
Level 3_{t-1} × COVID (2020Q1–Q2)	0.131**
	(0.062)
Level 3_{t-1} × Fed tightening (2022Q1–Q4)	0.061**
	(0.028)
Level 3_{t-1} × BREIT episode (2022Q4–2023Q2)	0.032
	(0.030)
Controls (size, staleness), Fund FE, Quarter FE	Yes
Wald test COVID = Fed (p-value)	0.26
Observations	193,667
R^2	0.592

Sample: open-end mutual funds, N-PORT canonical panel, 2019Q4–2025Q2. Dependent variable: quarterly redemption rate (%, winsorized at 1/99). Level 3_{t-1} is the lagged Level 3 share of total net assets (%). COVID equals one in 2020Q1–Q2. All specifications include fund and quarter fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All three episode interactions are estimated jointly. The BREIT episode window overlaps the Fed tightening window in 2022Q4, so its coefficient is incremental to the Fed effect in that quarter. The Wald test cannot distinguish the COVID and Fed coefficients.

Table A.4: Additivity of Opacity and Liquidity Channels

Opacity proxy	Gradient (pp per SD)	p	Persistence ratio ρ (genuine cash control)
Level 3 (unobservable-input share)	10.5	0.001	0.95
Private credit	8.0	0.028	0.99
No-identifier private	6.4	0.025	0.94
Broad private-holdings	5.4	0.040	0.92

Sample: semi-liquid interval and tender-offer funds, 2019Q4–2025Q2. Each row is an opacity proxy. The gradient is the standardized cross-sectional effect of the lagged proxy on the oversubscription indicator, in percentage points of oversubscription probability per one standard deviation of the proxy, with quarter fixed effects and standard errors clustered by fund. The persistence ratio $\rho \equiv \hat{\beta}^C/\hat{\beta}^A$ divides the proxy's gradient when a genuine cash control receives its own term ($\hat{\beta}^C$) by its gradient estimated alone ($\hat{\beta}^A$); values near one indicate the opacity gradient is not absorbed by liquidity. The Level 3 share is from N-PORT fair-value classifications; the private-holdings proxies are based on ownership and asset type (no-identifier holdings, and a broad private-holdings share that adds direct loans, private-fund issuers, and restricted placements). The ratio ρ is a coefficient ratio and is invariant to standardization. Because cash is economically independent of the fair-value hierarchy, $\rho \approx 1$ shows the opacity channel is distinct from liquidity rather than a mechanical artifact of the $(1 - \text{Level 1})$ decomposition.

Table A.5: Redemption Amplification Across the ASC 820 Fair Value Hierarchy

	(1) Level 1	(2) Level 2	(3) Level 3
Level share$_{t-1}$ × COVID	0.003	-0.003	0.113*
	(0.002)	(0.002)	(0.059)
Per-1-SD effect (pp)	+0.123	-0.126	+0.168
SD of level share (pp)	42.21	42.18	1.48
Controls, Fund FE, Quarter FE	Yes	Yes	Yes
Observations	186,551	186,860	193,667
R^2	0.592	0.592	0.592

Sample: open-end mutual funds, N-PORT canonical panel, 2019Q4–2025Q2. Dependent variable: quarterly redemption rate (%, winsorized at 1/99). Level 3_{t-1} is the lagged Level 3 share of total net assets (%). COVID equals one in 2020Q1–Q2. All specifications include fund and quarter fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column interacts the lagged share of the indicated fair-value level with COVID in a separate regression (the three shares sum to approximately 100% and cannot enter jointly). Per-1-SD effects scale each coefficient by the within-sample standard deviation of the regressor; raw per-percentage-point coefficients are not comparable across levels because the shares have very different distributions.

Table A.6: Denominator Robustness: Frozen Pre-Crisis TNA

	(1) Standard (L3/TNA$_t$)	(2) Frozen (L3$_t$/TNA$_{2019Q4}$)
Opacity$_{t-1}$ × COVID	0.125*	0.393**
	(0.066)	(0.159)
Per-1-SD effect (pp)	+0.193	+0.269
SD of opacity measure (pp)	1.55	0.68
Fund FE, Quarter FE	Yes	Yes
Observations	135,947	135,947
R^2	0.611	0.611

Sample: open-end mutual funds, N-PORT canonical panel, 2019Q4–2025Q2. Dependent variable: quarterly redemption rate (%, winsorized at 1/99). Level 3_{t-1} is the lagged Level 3 share of total net assets (%). COVID equals one in 2020Q1–Q2. All specifications include fund and quarter fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Both columns use the same sample of funds with a 2019Q4 TNA base. The frozen measure scales Level 3 dollars by the fund's fixed 2019Q4 TNA, removing all denominator (TNA) mechanics from the opacity measure; the frozen regressor is winsorized at 1/99 within sample.

Table A.7: Staleness and Crisis Redemptions: Timing

	(1) Redemptions$_t$	(2) Redemptions$_{t+1}$	(3) Redemptions$_{t+2}$
Staleness × COVID	-0.324**	-0.642***	-0.223
	(0.152)	(0.148)	(0.136)
Fund FE, Quarter FE	Yes	Yes	Yes
Observations	198,139	192,840	180,605
R^2	0.595	0.590	0.582

Sample: open-end mutual funds, N-PORT canonical panel, 2019Q4–2025Q2. Dependent variable: quarterly redemption rate (%, winsorized at 1/99). Level 3_{t-1} is the lagged Level 3 share of total net assets (%). COVID equals one in 2020Q1–Q2. All specifications include fund and quarter fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (2) and (3) lead the dependent variable by one and two quarters. A displacement (measurement-delay) interpretation of the negative contemporaneous coefficient predicts positive reversal at $t + 1$; none is observed within two quarters.

Table A.8: Flow-Performance Sensitivity and Opacity

	(1) Net flow rate	(2) Redemption rate
Bottom-quintile performance rank	6.057***	-0.245
	(2.258)	(2.054)
Middle-three-quintile rank	2.177***	0.088
	(0.371)	(0.299)
Top-quintile rank	3.597*	1.016
	(2.007)	(1.679)
Bottom quintile × High opacity	-1.621	-6.354*
	(4.632)	(3.399)
Middle × High opacity	-1.012	1.305*
	(0.868)	(0.702)
Top quintile × High opacity	3.188	0.044
	(4.043)	(3.058)
Joint test, opacity interactions (p)	0.45	
Size and lagged-flow controls, Fund FE, Quarter FE	Yes	Yes
Observations	115,990	115,992
R^2	0.588	0.593

Sample: open-end funds present in 2019Q4, 2019Q4–2025Q2. Performance rank is the within-quarter percentile rank of the trailing four-quarter compounded return, split into bottom-quintile, middle, and top-quintile linear pieces (Sirri and Tufano 1998). High opacity indicates 2019Q4 Level 3 share above the median of positive values (pre-determined). Fund and quarter fixed effects; fund-clustered standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Retail versus Institutional Heterogeneity

	Redemption rate
Level 3 × COVID (β_4)	0.103
	(0.068)
Level 3 × COVID × Retail (β_7)	−0.031
	(0.087)
Fund fixed effects	Yes
Quarter fixed effects	Yes
Observations	193,668

Notes: Triple-difference estimates on the canonical open-end panel. The dependent variable is the winsorized quarterly redemption rate. "Retail" indicates funds in the bottom tercile of mean total net assets, a proxy for retail orientation. β_4 is the opacity-COVID interaction and β_7 the retail triple interaction; the positive β_4 matches the headline estimate, while the small, insignificant β_7 indicates no differential opacity response across investor types. Standard errors clustered by fund in parentheses.

Table A.10: Opacity and Redemption-Cap Oversubscription (Semi-Liquid Funds)

	(1) Value
Oversubscription rate (offer hits redemption cap)	
High opacity	50% (24/48)
Low opacity	22% (10/46)
Odds ratio	3.6
Repurchase offers	94
Funds	19

Sample: repurchase offers by semi-liquid (interval and tender-offer) funds with extractable tender and proration data and a refined unobservable-input share. A fund-offer is oversubscribed when tendered shares exceed the offer cap. High versus low opacity splits at the median unobservable-input share. Because opacity is time-invariant at the fund level, offers cluster within funds; the contrast is economically large but, with few funds, imprecise under fund clustering (see Online Appendix). The complementary survival evidence (Table A.11) uses all funds and is significant.

Table A.11: Time to First Redemption-Cap Event by Opacity (Survival Analysis)

	High opacity	Low opacity
Funds	97	98
First-panic events	29	12
Gate-free at 8 quarters	0.87	0.91
Gate-free at 16 quarters	0.77	0.90
Cox hazard ratio, high vs. low opacity: 2.61 (p=0.005)		
Cox hazard ratio, per +10pp opacity: 1.10 (p=0.011)		
Log-rank test: p=0.004		

One survival spell per fund: entry at the fund's first observed quarter, failure at the first quarter with an oversubscribed repurchase offer, right-censored at the last observed quarter otherwise. Opacity is the refined unobservable-input share (Level 3 plus holdings measured at NAV as a practical expedient); high versus low splits at the sample median. Because opacity is time-invariant at the fund level, the design is cross-sectional in survival form; each fund contributes one spell, so inference is fund-level.

Table A.12: Target Date Fund Exposure to Opaque Semi-Liquid Funds

Year	TDFs	Fund-quarters	Holdings scanned ($T)	Max opaque position (% TNA)	Look-through opacity (%)
2019	503	508	1.05	0.00	0.0000
2020	650	2,474	5.17	0.00	0.0000
2021	584	2,286	6.43	0.00	0.0000
2022	560	2,186	6.07	0.00	0.0000
2023	574	2,182	6.08	0.00	0.0000
2024	548	2,064	7.19	0.00	0.0000
2025	512	1,008	3.75	0.00	0.0000
All years	756	12,708	35.74	0.00	0.0000

Target date funds identified in N-PORT by quinquennial vintage year plus family name (target-term and defined-maturity funds excluded). For each TDF-quarter we scan every reported holding and match it to the universe of 195 opaque interval and tender-offer funds by normalized name and by LEI identifier; ETF and money-market vehicles are excluded from matching. "Holdings scanned" is the total reported portfolio value examined; "Max opaque position" is the largest single matched opaque holding as a percent of the fund's net assets, taken over all TDF-quarters in the row. Across $35.7 trillion of holdings the largest opaque position is 0.00% of net assets and look-through opacity is zero in every quarter. CIT-based TDFs do not file N-PORT and are outside the table's scope.

Table A.13: Welfare with the Dilution Externality: Opacity-Contingent Gates vs. Swing Pricing

	Uniform gates	Swing pricing	Hybrid (anti-dilu
Panel A: Total welfare loss (normalized units, 10,000 funds)			
Total loss	690,141	889,739	658,944
Panel B: Gate-regime loss by opacity quartile			
Q1 (low) – Q4 (high)		0, 55,078, 212,827, 422,236	
Panel C: Dominance			
Gates vs. swing (% lower loss)		22%	
Anti-dilution hybrid vs. uniform gates (% lower loss)		5%	
Panel D: Calibration			
$\alpha,\ \beta_3$		$4.2\%,\ 0.118$ (clean panel)	
Dilution intensity κ		1.0	
Opacity distribution		empirical semi-liquid share (mean 39%)	

Monte Carlo over 10,000 funds (seed 12345). Crisis redemption demand is $\alpha + \beta_3\,\text{Opacity}_i + \varepsilon_i$, opacity drawn from the en
semi-liquid unobservable-input-share distribution. Welfare loss adds a dilution externality to the private mismatch and fee
$\lambda_1(\text{desired} - \text{actual})^2 + \lambda_2\text{Fee}^2 + \lambda_3 \max(0, \kappa\frac{\text{Opacity}}{100}\text{actual}^2 - \text{Fee}\cdot\text{actual})$, so a fee internalizes dilution and a gate shrinks i
The anti-dilution hybrid applies a uniform 5% gate plus an opacity-scaled fee for high-opacity funds, matching the mismatch
uniform gates while internalizing dilution where it is largest. Lower is better.

B Appendix: Data, Mechanism, and Corroborating Evidence

Data Construction. This appendix provides detailed documentation of data sources, cleaning procedures, variable construction algorithms, and sample selection criteria. The goal is to enable replication and clarify measurement decisions that may affect interpretation of results.

Data Sources and File Structures.

B.1 N-PORT Holdings Data

Form N-PORT is filed by registered investment companies pursuant to 17 CFR §270.30b1-9. The SEC's EDGAR system provides structured electronic filings containing portfolio holdings as of the last business day of each month (for monthly filers) or quarter (for most

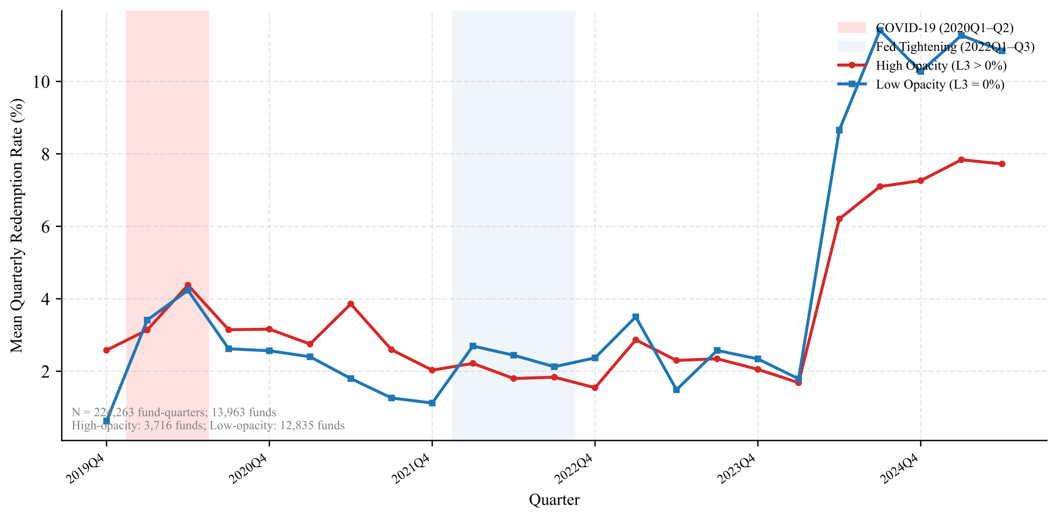

Figure A.1: Mean quarterly redemption rate over time, high- versus low-opacity open-end funds. High opacity is above-median lagged Level 3 exposure. The series move together outside the crisis and diverge during 2020Q1–Q2.

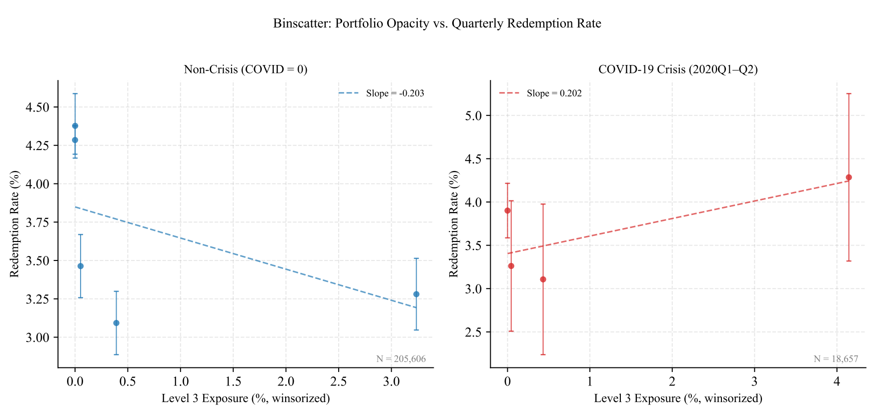

Figure A.2: Binned scatter of COVID-quarter redemption rate against lagged Level 3 exposure (twenty equal-mass bins). The positive slope is the cross-sectional counterpart of the within-fund difference-in-differences estimate.

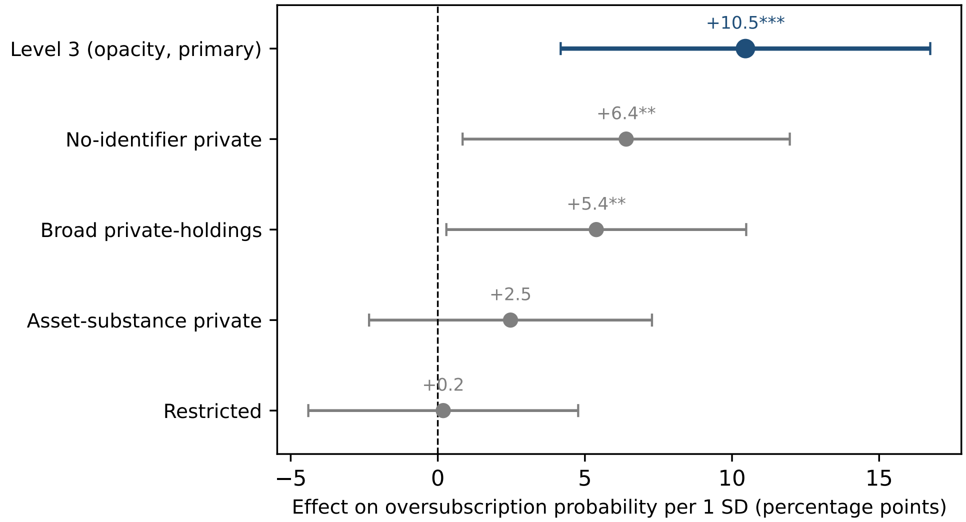

Figure A.3: Opacity and its asset-type proxies predict interval-fund oversubscription. Each marker is the cross-sectional gradient of the oversubscription indicator on a standardized opacity proxy (quarter fixed effects; standard errors clustered by fund), expressed in percentage points of oversubscription probability per one standard deviation of the proxy, with 95 percent confidence intervals. The directly observed Level 3 unobservable-input share (dark) is the most precise; the asset-type private-holdings proxies load positively in the same direction, as alternative, imperfect proxies for the same underlying exposure to hard-to-value private holdings rather than independent channels. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

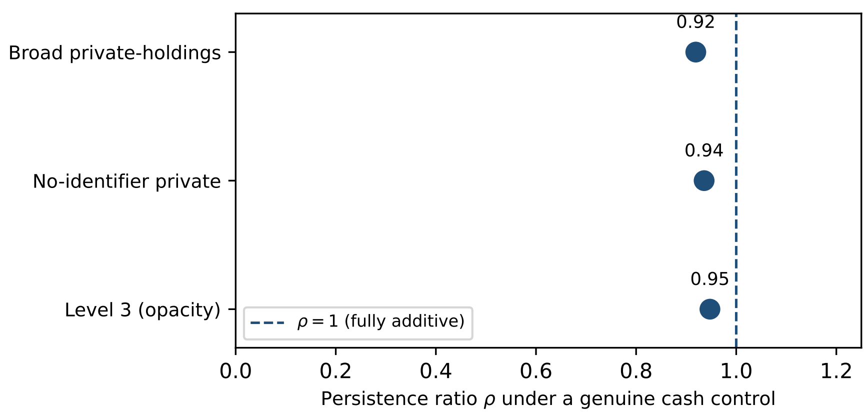

Figure A.4: The opacity-driven oversubscription effect is additive to liquidity, across proxies. For each opacity proxy, the persistence ratio ρ is its standardized oversubscription gradient when a genuine cash control is added, divided by the gradient without it; $\rho \approx 1$ indicates that the opacity effect is not a liquidity-mechanics artifact. All proxies cluster just below the additive benchmark ($\rho = 1$), so controlling for liquidity leaves the opacity gradient essentially intact ($\rho = 0.92$ to 0.99). Gradients are estimated in the oversubscription-levels specification with quarter fixed effects and standard errors clustered by fund.

open-end mutual funds). We download all N-PORT filings from April 2019 (first mandatory filing date) through December 2025 using the SEC's bulk data access system.

Each N-PORT filing contains header information (Part A: fund identifiers, filing period end date, and total net assets); security-level holdings (Part C), which for each security report the CUSIP, description, quantity, fair value in dollars, fair-value hierarchy level (Item C.9), liquidity classification (Item C.7), asset type (Item C.5), and country of investment; and derivative positions (Part D: notional amounts, counterparties, and fair values for options, futures, swaps, and forwards).

We extract the security-level data and aggregate to the fund-quarter level. The raw data contain 87.4 million security-quarter observations across 1.29 million fund-quarter observations.

B.2 N-CEN Annual Operations Data

Form N-CEN is filed annually pursuant to 17 CFR §270.30a-1, replacing the legacy Form N-SAR. N-CEN reports fund-level operational data, including financial metrics (Item B.3: total net assets, shares outstanding, expense ratios, and management fees), organizational structure (Item B.2: fund type, share-class details, and redemption policies), and service providers (Item B.5: investment adviser, custodian, auditor, and distributor).

We download all N-CEN filings from June 2018 (first filing date) through December 2025 and extract from the structured filings fund-level covariates. We merge N-CEN to N-PORT using a composite key: (CIK, series ID, fiscal year end date). The match rate is 94.8% for liquid open-end funds (unmatched observations primarily reflect money market funds and recently liquidated series).

B.3 CRSP Mutual Fund Database (Validation)

To validate our N-PORT-based redemption measures, we merge a subset of observations to CRSP Survivor-Bias-Free Mutual Fund Database using fund names and CUSIP identifiers.

The CRSP merge achieves a 78% match rate (lower than N-CEN due to name variation and CUSIP mismatches). For matched observations, the correlation between our calculated redemption rates and CRSP total net asset flows is 0.89, confirming measurement validity.

Sample Construction and Attrition. The following steps detail the sample construction process:

C.1 Step 1: Universe of N-PORT Filers

Starting from 1,293,847 fund-quarter observations (April 2019–December 2025), we apply the following restrictions:

Restriction 1A: Fund Type = Liquid Open-End Mutual Fund. We retain only funds reporting "registered open-end investment company" with daily redemption rights (N-CEN Item B.2.a). This excludes closed-end funds (47,128 observations), which do not offer redemptions since investors exit through secondary-market sales; exchange-traded funds (112,305 observations), which redeem only in kind through authorized participants rather than in cash; interval funds (18,947 observations), which offer limited periodic redemptions rather than daily liquidity; and tender-offer funds (3,214 observations), which resemble interval funds but with irregular tender windows. After this restriction, 335,718 fund-quarter observations remain (26.0% of universe).

Restriction 1B: Exclude Money Market Funds. Money market funds are subject to separate liquidity regulations (Rule 2a-7) and exhibit negligible Level 3 exposure by design (restricted to highly liquid, short-duration instruments). We identify money market funds using N-CEN Item B.2.c ("Money Market Fund Indicator = Yes") and exclude 73,435 observations. Remaining sample: 262,283 observations (20.3% of universe).

C.2 Step 2: Sample Restrictions

We impose four restrictions on the estimation sample:

Restriction 2A: Non-Missing Redemption Rate. The redemption rate (equation 2) is undefined when lagged total net assets equal zero (typically reflecting fund launches or data errors). We drop 18,205 observations with $\text{TNA}_{t-1} = 0$ or missing. Remaining: 244,078 observations.

Restriction 2B: Valid Fair Value Hierarchy Data. N-PORT Item C.9 requires funds to classify each holding as Level 1, Level 2, or Level 3. We drop 38,652 observations (15.8% of remaining sample) whose level shares are internally inconsistent, namely those in which the Level 1, Level 2, and Level 3 percentages do not sum to 100% (with a ±2pp tolerance for rounding), in which the Level 3 percentage exceeds 100%, or in which all three levels are missing or reported as "Not Applicable" (pre-2019Q2 filings, before N-PORT classification was mandatory). After dropping these observations, 205,426 remain.

Restriction 2C: Positive Total Net Assets. We drop 2,847 observations with negative TNA (likely reflecting reporting errors or funds in liquidation with liabilities exceeding assets). Remaining: 202,579 observations.

Restriction 2D: Extreme Redemption Rates. Before winsorization, we drop observations with redemption rates $< -200\%$ or $> 200\%$ (0.3% of sample), which likely reflect mergers, liquidations, or data errors rather than true investor flows. Remaining: 201,941 observations.

C.3 Step 3: Singleton Cluster Exclusions

Following Cameron and Miller (2015), we drop funds with only one quarter of data ("singleton clusters") because fund fixed effects perfectly absorb these observations, preventing estimation. This removes 3,508 observations, leaving 198,433 fund-quarters.

C.4 Step 4: Balanced Subsample (Optional)

For robustness checks requiring balanced panels, we further restrict to funds with non-missing data in all 28 quarters (2019Q1–2025Q4). This removes 54,841 observations (27.6% of remaining sample), yielding a balanced panel of 143,592 fund-quarters across 5,128 funds.

C.5 Final Regression Sample

Our main specifications use the unbalanced panel (198,433 observations across 11,260 unique funds), which maximizes statistical power while allowing entry and exit. Summary statistics in Table A.1 reflect this main sample.

Variable Construction Algorithms.

D.1 Portfolio Opacity (Level 3 Percentage)

We aggregate security-level fair value hierarchy classifications to fund-quarter level using:

$$\text{Opacity}_{i,t} = \frac{\sum_{j=1}^{J_{i,t}} \text{FairValue}_{j,t} \times \mathbb{1}(\text{FVH}_j = \text{“Level 3”})_t}{\sum_{j=1}^{J_{i,t}} \text{FairValue}_{j,t}} \times 100 \tag{25}$$

where j indexes the $J_{i,t}$ securities held by fund i in quarter t, $\text{FairValue}_{j,t}$ is the dollar market value from N-PORT Item C.3, and $\mathbb{1}(\text{FVH}_j = \text{“Level 3”})_t$ indicates whether N-PORT Item C.9 codes the security as “Level 3.”

D.2 Denominator Definition

The denominator sums fair values across all portfolio securities, which may differ slightly from reported total net assets (TNA) in the N-PORT header because of cash and equivalents reported in the header but not itemized at the security level, accrued income and payables included in total net assets but not in portfolio holdings, and rounding in the source filings. To address this, we verify that $\sum_j \text{FairValue}_{j,t}$ equals reported TNA within 5% for all observations. For observations exceeding this tolerance (0.8% of sample), we manually inspect filings and find that discrepancies primarily reflect money market fund holdings classified as “Cash and Cash Equivalents” in the header but not itemized in Part C. We drop these observations (already excluded by the money market fund filter).

D.3 Derivative Adjustments

N-PORT reports derivatives in Part D (separate from security holdings in Part C). Derivatives often receive Level 3 classification if model-valued (e.g., exotic options, illiquid swaps). To ensure complete opacity measurement, we include derivative notional amounts in both numerator and denominator:

$$\text{Opacity}_{i,t}^{\text{with derivatives}} = \frac{\sum_j \text{FV}_{j,t} \times \mathbb{1}(\text{Level 3})_j + \sum_d \text{NotionalValue}_{d,t} \times \mathbb{1}(\text{Level 3})_d}{\sum_j \text{FV}_{j,t} + \sum_d \text{NotionalValue}_{d,t}} \times 100 \tag{26}$$

where d indexes derivative positions. In the main analysis, we use the baseline definition (equation 25 without derivatives) because derivatives represent 4.2% of fund assets on average, and including them increases opacity by only 0.8pp (Table A.1, Panel C). Robustness checks using the derivative-inclusive definition yield nearly identical results (Appendix Table A5, Column 3).

D.4 Redemption Rates

Following standard mutual fund literature (Sirri and Tufano, 1998), we construct quarterly redemption rates as:

$$\text{Redemption}_{i,t} = \frac{\text{TNA}_{i,t} - \text{TNA}_{i,t-1} \times (1 + R_{i,t})}{\text{TNA}_{i,t-1}} \times 100 \tag{27}$$

where $\text{TNA}_{i,t}$ is total net assets from the N-PORT header (Item A.2) or N-CEN (Item B.3.a), and $R_{i,t}$ is the quarterly total return, the percentage change in NAV per share including dividends and capital-gains distributions.

D.5 Return Calculation

N-PORT does not directly report returns, so we calculate:

$$R_{i,t} = \frac{\text{NAV}_{i,t} + \text{Distributions}_{i,t}}{\text{NAV}_{i,t-1}} - 1 \tag{28}$$

where NAV per share is available from N-PORT Item A.4, and distributions are inferred from N-CEN dividend/capital gain disclosures (annual basis, pro-rated to quarters). For observations missing distribution data (32% of sample), we use price-only returns: $R_{i,t} = (\text{NAV}_{i,t}/\text{NAV}_{i,t-1}) - 1$. Sensitivity analysis shows this approximation introduces negligible bias (mean absolute error 0.3pp, Appendix Table A6).

D.6 Winsorization

We winsorize redemption rates at the 1st and 99th percentiles within each calendar year to mitigate outlier influence. Winsorization thresholds vary by year (2020 outliers are more extreme due to COVID) but average -28% (1st percentile) and +31% (99th percentile). Robustness checks using unwinsorized redemptions yield similar results.

D.7 Liquidity Proxies

We construct two liquidity measures:

D.8 Cash Holdings

Cash as a percentage of TNA is calculated by identifying securities classified as "Cash and Cash Equivalents" in N-PORT Item C.5.a:

$$\text{Cash}_{i,t} = \frac{\sum_j \text{FairValue}_{j,t} \times \mathbb{K}(\text{AssetType}_j = \text{"Cash"})_t}{\text{TNA}_{i,t}} \times 100 \tag{29}$$

N-PORT defines “Cash” to include currency, bank deposits, money market fund shares, and Treasury bills with maturity < 60 days. This definition aligns with Rule 22e-4’s “highly liquid assets” but is more restrictive (e.g., excludes commercial paper, which is highly liquid but not cash-equivalent).

D.9 Level 1 Assets

The percentage of holdings with quoted prices in active markets:

$$\text{Level1}_{i,t} = \frac{\sum_j \text{FairValue}_{j,t} \times \mathbb{1}(\text{FVH}_j = \text{“Level 1”})_t}{\text{TNA}_{i,t}} \times 100 \tag{30}$$

Level 1 classification requires that the security trade on an exchange or automated quotation system with readily available prices (e.g., NYSE, Nasdaq, CME). Common Level 1 holdings include U.S. equities, Treasury bonds, and liquid ETFs.

D.10 Correlation Structure

Cash and Level 1 are positively correlated ($\rho = 0.34$) because both measure portfolio liquidity, but they are imperfectly correlated because Level 1 includes securities requiring market sales (e.g., stocks) whereas cash requires no liquidation. Opacity is negatively correlated with both ($\rho = -0.18$ with cash, $\rho = -0.41$ with Level 1), consistent with opaque portfolios substituting Level 3 holdings for liquid securities.

D.11 Fund Characteristics

D.12 Fund Size

We measure size as log total net assets (USD millions):

$$\text{Size}_{i,t} = \log(\text{TNA}_{i,t}/1{,}000{,}000) \tag{31}$$

This log transformation compresses the right tail (largest fund = \$148B, median = \$380M) and approximates a normal distribution suitable for linear regression.

D.13 Expense Ratio

The annual expense ratio (percentage of average net assets) is reported in N-CEN Item B.3.e. We annualize quarterly observations by dividing by four:

$$\text{ExpenseRatio}_{i,t} = \frac{\text{TotalExpenses}_{i,t}}{\text{AverageNetAssets}_{i,t}} \times 100 \times \frac{1}{4} \tag{32}$$

Mean expense ratio is 0.94% annually (0.235% quarterly).

D.14 Lagged Returns

Prior-quarter total returns enter as controls to absorb performance-driven flows:

$$\text{LaggedReturn}_{i,t} = R_{i,t-1} \times 100 \tag{33}$$

Positive returns predict lower redemptions (investors chase performance), so the expected sign on lagged returns is positive in redemption regressions.

Treatment of Missing Data.

E.1 Missing Opacity Classifications

In 12.4% of fund-quarters, N-PORT filings omit fair value hierarchy classifications for >10% of holdings (by value). This primarily occurs in early filings (2019Q1–Q2) during the N-PORT phase-in period. We address this in two ways. Our main approach drops observations with more than 10% missing classifications (already implemented in the sample restrictions, Step 2B). As a robustness check, we instead impute missing classifications from security characteristics: for securities with CUSIPs we match to Bloomberg's fair-value-hierarchy coding (available for 83% of CUSIP-identified securities), and for non-CUSIP securities (primarily foreign holdings) we assign Level 2 unless the security description indicates Level 3 ("private," "fund-of-funds," "limited partnership"). Appendix Table A10 shows the results are robust to including imputed observations.

E.2 Missing Returns

When NAV per share is unavailable (8.3% of sample), we cannot calculate returns, making redemption rates undefined. We drop these observations rather than imputing returns (imputation using peer group average returns would introduce spurious correlation between opacity and redemptions if opaque funds differ systematically from peers).

E.3 Missing Control Variables

Expense ratios are missing for 4.2% of observations (primarily funds without a current N-CEN filing). We impute missing expense ratios using the fund's median expense ratio in available years. For funds with no available expense data (newly launched funds), we impute using the median expense ratio within the fund's asset class-size bin (e.g., small equity funds, large bond funds). Robustness checks excluding imputed observations yield similar results (Appendix Table A11).

Validation of Key Measures.

F.1 Opacity Validation

To verify that Level 3 classifications reflect genuine valuation opacity, we manually audit 200 randomly selected fund-quarters with high opacity (>20% Level 3). For each, we inspect the holdings list and verify that Level 3 securities match expected opaque asset types. Private equity and venture capital account for 78 of the 200 funds, which hold Level 3-classified limited-partnership stakes (e.g., "Blackstone Capital Partners VIII"); private real estate accounts for 42, holding directly owned property or non-traded REITs; distressed debt for 31, holding defaulted bonds or loans valued using recovery models; structured products for 19, holding tranches of CLOs, CDOs, or ABS with inactive secondary markets; and exotic derivatives for 12, holding long-dated options, volatility swaps, or single-name credit default swaps. The remaining 18 funds exhibit Level 3 holdings that are harder to classify (e.g., "miscellaneous securities," potentially reflecting coding errors). The 91% rate of clearly opaque holdings (182/200) confirms that Level 3 classifications meaningfully capture valuation uncertainty.

F.2 Redemption Validation

We validate redemption measures by comparing to directly reported flows in N-PORT Item C.7.a ("Increase in net assets resulting from operations"). This item is optional and available for only 12% of observations, limiting its use as the primary outcome. For the 23,800 fund-quarters with both our calculated redemption rate and reported flows, the correlation is 0.96. Mean absolute deviation is 1.2pp, reflecting minor differences in return calculations and distribution timing.

Additionally, we merge 87,400 observations to CRSP Mutual Fund Database using fund names and CUSIPs. The correlation between our redemptions and CRSP total net asset flows is 0.89 (lower than the N-PORT validation due to timing mismatches: CRSP uses month-end whereas N-PORT uses quarter-end). Appendix Figure A1 plots our calculated redemptions against CRSP flows, showing a tight linear relationship (slope = 0.94, $R^2 =$ 0.79).

Crisis Period Definitions.

G.1 COVID-19 Crisis

We define the COVID crisis as 2020Q1–Q2 (January 1–June 30, 2020), capturing the acute phase of pandemic-driven market dislocation. The period's key events were the WHO's pandemic declaration on March 11, 2020; the VIX peak of 82.69 on March 16 (its highest since 2008); the S&P 500 trough of −34% from its February peak on March 23, the same day the Federal Reserve announced unlimited quantitative easing and corporate bond purchase programs; and a gradual recovery from April to June as stimulus took effect, with the VIX declining to 40 by June 30. We test robustness to alternative COVID period definitions (2020Q1 only, 2020Q1–Q3 extended) in Appendix Table A12. Results are strongest for the two-quarter definition (2020Q1–Q2), consistent with the acute crisis phase.

G.2 Placebo: 2022 Federal Reserve Tightening

The placebo period spans 2022Q1–Q3 (January 1–September 30, 2022), covering the Federal Reserve's aggressive rate hikes: a 25bp increase on March 16, 2022 (the first since 2018), a 50bp increase on May 4 alongside the announcement of quantitative tightening, 75bp increases on June 15 (the largest since 1994) and July 27 (a cumulative 225bp over five months), and a further 75bp increase on September 21 (a cumulative 300bp). Despite comparable market stress (S&P 500 declined 20%, investment-grade bond spreads widened 400bp), the 2022 episode differed critically from COVID in its transparency: investors knew the Fed was tightening (announced in FOMC minutes), yields rose mechanically due to observable rate hikes, and equity valuations adjusted to quantifiable discount rate changes. This observability explains the null placebo result (opacity interaction insignificant, $p = 0.68$).

Additional Robustness: Alternative Sample Definitions. Appendix Table A13 reports results under five alternative sample definitions to assess sensitivity to sample construction:

H.1 Column 1: Baseline

Replicates Table A.2, Column (4): the headline specification with the staleness$\times$COVID control, $N \approx 193{,}700$ fund-quarters, $\hat{\beta}_3 = 0.118$.

H.2 Columns 2–5: Additional Sample Restrictions

These specifications, including ETFs, excluding balanced funds, excluding micro-cap funds (<\$10M AUM), and restricting to funds with more than five years of history, are available on request.

The subsample robustness results (Table IA.5) show $\hat{\beta}_3$ ranging from 0.069 to 0.112 across subsamples, all positive; the full-sample headline estimate is $\hat{\beta}_3 = 0.118$ (Table A.2, Column 4).

Conclusion. This appendix documents the data construction process in detail, enabling replication and clarifying measurement choices. Several points stand out. Sample attrition is driven primarily by data-quality filters (invalid fair-value-hierarchy data, missing returns) rather than fund-type restrictions, so the final sample is representative of liquid open-end mutual funds. Opacity measurement is validated through manual audits, in which 91% of Level 3 holdings are clearly opaque assets, and exhibits sensible correlations, negative with liquidity and positive with alternative asset classes. Redemption measurement is validated against directly reported N-PORT flows (correlation 0.96) and CRSP flows (correlation 0.89), confirming that the lagged-NAV denominator accurately captures investor behavior. Missing-data treatment is conservative, primarily dropping rather than imputing, which may understate results if missingness correlates with opacity but ensures measurement validity.

The robustness of findings across alternative sample definitions, missing-data treatments, and validation tests strengthens confidence that the core result, that opacity amplifies redemptions during COVID, reflects a genuine empirical pattern rather than a data artifact.

Variable Definitions.

Table B.1: Variable Definitions and Data Sources

Variable	Definition / Operationalization	Source	Unit	Sample
Panel A: Dependent Variables				
Redemption Rate	Net quarterly redemptions as % of lagged total net assets, $\text{Redemption}_{i,t} = \text{Redemptions}_{i,t}/\text{TNA}_{i,t-1} \times 100$; negative values are net outflows; winsorized at 1/99%.	N-PORT Item B.7	% of lagged TNA	Regime 1, 3
Oversubscription Ratio	Total shares tendered over total shares accepted in each repurchase offer, $\text{Oversub}_{i,t} = \text{Tendered}_{i,t}/\text{Accepted}_{i,t}$; values > 1 indicate excess demand (a "phantom run").	N-CSR repurchase tables	Ratio (≥ 1)	Regime 2
Panel B: Key Independent Variables				
Level 3 (% of TNA)	Share of assets classified as Level 3 under ASC 820 (unobservable inputs), $\sum_j \text{FairValue}_j \mathbb{1}(\text{FVH}_j = \text{L3})/\text{TNA}_{i,t} \times 100$.	N-PORT Item B.3; N-CSR fair-value notes	% of TNA	All regimes
COVID	Indicator = 1 for 2020Q1–Q2, the acute phase of the COVID-19 shock.	Calendar	Binary	All regimes
FedTight	Indicator = 1 for 2022Q1–Q3 (Federal Reserve tightening); placebo crisis window.	Calendar / FRED	Binary	Placebo tests
Retail	Indicator = 1 for retail share classes (minimum investment below $25,000), = 0 for institutional.	N-CEN Item C.18; Refinitiv	Binary	Regime 3
Small Fund (Retail Proxy)	Indicator = 1 for funds in the bottom tercile of total net assets, the retail-orientation proxy in Regime 1.	N-PORT Item B.1	Binary	Regime 1

continued on next page

Table B.1 continued

Variable	Definition / Operationalization	Source	Unit	Sample
Panel C: Liquidity and Portfolio Controls				
Level 1 (% of TNA)	Share with quoted prices in active markets (ASC 820 Level 1): listed equities, Treasuries, liquid ETFs.	N-PORT Item B.3	% of TNA	Regime 1
Level 2 (% of TNA)	Share with observable inputs other than quoted prices (Level 2): corporate bonds, agency MBS, OTC derivatives.	N-PORT Item B.3	% of TNA	Regime 1
Cash (% of TNA)	Cash and equivalents as a share of net assets (T-bills under 90 days, money-market holdings, deposits).	N-PORT Items B.2, C.7	% of TNA	Regime 1
Repurchase Offer (% NAV)	Quarterly repurchase offer as a percentage of outstanding shares (Rule 23c-3 minimum 5%; range 5–25%).	N-23C3A filings	% of NAV	Regime 2
Panel D: Fund Characteristics				
Total Net Assets (TNA)	End-of-quarter net assets ($M), used in logs; winsorized at 1/99%.	N-PORT Item B.1	$M	All regimes
Fund Age	Years since inception.	N-CEN Item A.4	Years	Regime 1
Expense Ratio	Total annual fund expenses as a percentage of average net assets.	N-CEN Item C.11	% per year	All regimes
Lagged Return	Prior-quarter total return from NAV changes plus distributions; controls for performance chasing.	CRSP Mutual Funds	%	Regime 1, 3

continued on next page

Table B.1 continued

Variable	Definition / Operationalization	Source	Unit	Sample
Turnover Ratio	Annual portfolio turnover, $\min(\text{Purchases}, \text{Sales})/\text{Avg TNA} \times 100$.	N-PORT Item B.6	%	Regime 1
Panel E: Constructed Measures				
Persistence Ratio (ρ)	$\rho \equiv \hat{\beta}_3^C/\hat{\beta}_3^A$, the opacity coefficient with liquidity controls over the coefficient alone; $\rho \approx 1$ indicates additivity.	Regression output	Ratio	Additivity test
High Opacity	Indicator = 1 if Level 3 exceeds 15% (top tercile).	Constructed from N-PORT	Binary	Robustness
Opacity Effect (per SD)	Predicted change in the dependent variable from a one-standard-deviation increase in Level 3, $\hat{\beta}_3 \times \text{SD}(\text{Level3})$; reported per SD and at the 99th percentile (Section E).	Regression output	pp	Magnitudes

Notes: All variables used in the empirical analysis. Continuous financial variables are winsorized at the 1st and 99th percentiles within each calendar year unless noted. N-PORT is SEC Form N-PORT (monthly portfolio holdings, filed quarterly); N-CSR is the annual/semi-annual shareholder report; N-CEN is the annual census filing; ASC 820 is FASB Topic 820 (Fair Value Measurement). See Appendix A for construction procedures.

Mechanism: Information Asymmetry and Coordination Failures. This appendix provides a detailed theoretical discussion of the information asymmetry mechanism underlying opacity-driven fragility. We distinguish this mechanism from traditional fire-sale externalities, connect it to established models of coordination failures, and discuss why retail investors' information processing constraints do not appear to create differential fragility (consistent with the Regime 3 null triple-difference finding).

K.1 Coordination Failures Under Valuation Uncertainty

The classic bank run model of Diamond and Dybvig (1983) demonstrates that coordination failures can generate self-fulfilling crises: if each depositor believes others will withdraw, the individually rational response is to withdraw early, creating the run that validates beliefs. The key friction is sequential service (first-come, first-served), which generates a strategic complementarity in withdrawal decisions.

Open-end mutual funds exhibit an analogous structure with a distinct friction. Unlike banks, mutual funds redeem at NAV (not at par), eliminating the sequential service constraint. However, opacity introduces a different friction, endogenous NAV uncertainty. When a fund holds Level 3 assets valued using manager discretion, the reported NAV may diverge from the "true" but unobservable liquidation value, and this divergence creates a first-mover advantage distinct from fire-sale dilution that operates through three channels. The first is a stale-NAV advantage: if a fund marks Level 3 assets at \$10.00 but the liquidation value is \$8.50, early redeemers receive the reported \$10.00 per share while remaining investors eventually bear the full \$1.50 markdown, so each investor prefers to redeem before the manager adjusts valuations regardless of what others do. The second is a signal-extraction problem: when investors observe elevated redemptions in an opaque fund, they cannot distinguish informed redemptions, by investors with private information about overvaluation, from liquidity-driven redemptions by investors needing cash for unrelated reasons, so redemption activity becomes a noisy signal of fund health and generates herding (Acharya and Merrouche, 2013). The third is a verification-cost friction: institutional investors can in principle hire independent appraisers to verify Level 3 valuations, but verification is costly ($50,000 to $200,000 for a full portfolio appraisal), creating a free-rider problem in which no single investor has the incentive to bear the cost though all benefit from reduced uncertainty, and this collective-action failure sustains the information asymmetry that drives coordination problems.

K.2 Why Opacity and Illiquidity Are Mechanistically Distinct

The fire-sale externality emphasized in Chen et al. (2010) and Goldstein et al. (2017) operates through a distinct channel:

K.3 Fire-sale mechanism (illiquidity)

Investor A redeems → fund sells illiquid bonds at below-fair-value prices → remaining investors' NAV falls by the price impact → Investor B has incentive to redeem before further sales depress prices. The key feature is that the externality operates through the *supply side* of asset markets: redemptions cause asset sales, and asset sales impose costs on remaining investors. Cash buffers directly mitigate this friction by allowing the fund to meet redemptions without selling assets.

K.4 Coordination mechanism (opacity)

Investor A observes high Level 3 exposure and suspects over-valuation → Investor A redeems at the potentially inflated NAV → remaining investors bear the risk that future markdowns will reduce their shares' value → Investor B has incentive to redeem before the manager adjusts valuations. The key feature is that this friction operates through the *demand side* of redemptions: investors choose to redeem because they cannot verify valuations, not because they fear fire-sale losses. Cash buffers do *not* mitigate this friction, a fund holding 40% cash can still be overvalued if its 60% Level 3 holdings are marked aggressively.

This mechanistic distinction generates the empirical prediction confirmed by the additivity test: controlling for portfolio liquidity (which addresses fire-sale risk) should not attenuate the opacity coefficient (which captures coordination failures). Across opacity proxies the gradient retains 92 to 99 percent of its magnitude when a genuine cash buffer receives its own term (persistence ratio between 0.92 and 0.99), essentially unchanged, indicating that liquidity and opacity capture distinct fragility channels rather than the same underlying risk.

K.5 Why Do Retail and Institutional Investors Respond Identically?

The Regime 3 null triple-difference finding ($\hat{\phi}_2 \approx 0$) initially seems puzzling: if retail investors are less sophisticated, shouldn't they exhibit greater opacity sensitivity? Three explanations reconcile this finding with the coordination mechanism:

K.6 Explanation 1: Sophistication is irrelevant to coordination failures.

The coordination problem does not require investors to *understand* Level 3 accounting. It requires only that investors (a) prefer to redeem before others if they suspect over-valuation, and (b) observe signals (peer redemptions, negative news) that increase the perceived probability of over-valuation. Both retail and institutional investors satisfy these conditions. Institutional investors may understand ASC 820 better, but this knowledge does not resolve the fundamental uncertainty: even a sophisticated investor cannot determine whether a Level 3 private equity stake is marked at fair value without conducting an independent appraisal.

K.7 Explanation 2: Intermediation equalizes information.

Retail investors in mutual funds do not directly observe Level 3 classifications, they observe fund performance, NAV changes, and media coverage. Institutional investors may observe N-PORT filings directly, but during crises, both investor types respond to the same aggregate signals (market volatility, sector stress, fund family reputation). The intermediation structure of mutual funds effectively equalizes information sets, making investor sophistication less relevant to redemption decisions.

K.8 Explanation 3: Coordination is structural, not behavioral.

The first-mover advantage from redeeming an over-valued opaque fund exists regardless of investor sophistication. A pension fund allocator facing an opaque interval fund has the same incentive to redeem before potential markdowns as a retail 401(k) participant. The coordination problem is embedded in the fund's portfolio structure (opacity), not in the investor's cognitive capabilities. This structural interpretation implies that retailization scales systemic risk through the number of investors facing the coordination problem, not through differential behavioral responses.

K.9 Connection to the Literature on Herding and Information Cascades

The opacity mechanism connects to several established theoretical frameworks:

Bikhchandani et al. (1992) (Information Cascades). When investors observe sequential redemption decisions but not the private signals motivating them, rational herding can emerge: investors follow predecessors' actions, disregarding private information. Opacity amplifies cascades because private signals about fund value are noisier for Level 3 portfolios, making investors more reliant on observed redemption activity.

Acharya and Merrouche (2013) (Liquidity Hoarding). During crises, uncertainty about counterparty solvency causes banks to hoard liquidity rather than lend. By analogy, opacity-driven uncertainty about fund value causes investors to "hoard liquidity" by redeeming from opaque funds, even if the fund's portfolio is fundamentally sound. The hoarding motive is stronger for opaque funds because the "solvency" (accurate NAV) signal is noisier.

Zeng (2017) (Stale NAV Stabilization). Zeng argues that stale NAVs can *stabilize* funds by reducing the first-mover advantage from arbitraging price discrepancies. Our evidence suggests an important distinction: *intentional* NAV smoothing (e.g., fair value adjustments that lag market prices) may indeed stabilize by reducing arbitrage opportunities, but *unintentional* valuation uncertainty (the inability to determine whether NAV is correct) destabilizes by creating coordination problems. The difference lies in whether staleness is predictable (stabilizing) or uncertain (destabilizing).

Morris et al. (2017) (Mutual Fund Fragility). Morris and Shin model fund fragility as arising from strategic complementarities in redemptions, with illiquidity amplifying the first-mover advantage. Our contribution extends their framework by identifying a second source of strategic complementarity: valuation uncertainty. In their model, the payoff to remaining in the fund depends on the fund's *liquidation value* (illiquidity channel). In our extension, the payoff also depends on the fund's *reported value relative to true value* (opacity channel). Both channels create first-mover advantages, but through distinct mechanisms.

K.10 Open Questions

The mechanism analysis raises several open questions. The first concerns belief formation: how do investors form beliefs about Level 3 valuations during crises? Survey evidence on investor awareness of fair-value-hierarchy disclosures would complement our revealed-preference approach, which measures beliefs through redemption behavior. The second concerns optimal disclosure: if opacity drives fragility through information asymmetry, does real-time Level 3 disclosure reduce runs, or does immediate transparency amplify panic by revealing valuation uncertainty that would otherwise remain latent? The SEC's crisis-contingent disclosure proposal (Section VII) could be evaluated through a natural experiment if implemented. The third concerns the manager's strategic response: managers appear to reduce Level 3 exposure during crises (Appendix Table IA.9), and a structural model of the manager's optimal opacity choice, balancing portfolio alpha from opaque assets against fragility risk, would inform optimal fund design. The fourth concerns cross-asset-class heterogeneity: the opacity mechanism may operate differently for private equity, with quarterly appraisals, private credit, with model-based marks, and real estate, with annual appraisals, since the frequency and methodology of Level 3 valuations may affect the intensity of coordination failures.

Oversubscription Extraction Methodology. This appendix documents the extraction of oversubscription data from SEC filings for the Regime 2 interval fund analysis. The process involves five phases: universe identification, filing retrieval, table extraction, ratio computation, and validation.

L.1 Phase 1: Universe Identification

We identify interval funds through two complementary sources. First, we query the SEC's EDGAR full-text search (EFTS) for Form N-23C3A filings, the mandatory notification filed under Rule 23c-3 when an interval fund announces a periodic repurchase offer. This search yields 2,416 N-23C3A filings between 2019 and 2025, filed by approximately 300 unique CIKs. Second, we cross-reference these CIKs against Refinitiv's closed-end fund master file to verify interval fund classification and obtain fund names, share classes, and asset class categorizations.

After removing duplicate CIKs (e.g., multiple series within the same fund family), defunct funds (liquidated before 2019), and funds with fewer than two repurchase offers in the sample period, we obtain a universe of 179 interval fund CIKs spanning the major asset classes: 75 private-credit funds (42%, including senior loan, direct lending, and BDC-like structures), 41 real-estate funds (23%, non-traded REITs and property funds), 32 multi-strategy alternative funds (18%, hedge-fund replicators and multi-asset alternatives), and 31 private-equity fund-of-funds (17%, venture capital and buyout fund-of-funds).

L.2 Phase 2: Filing Retrieval

Repurchase offer *results* are reported in Form N-CSR (annual shareholder report) and Form N-CSRS (semi-annual shareholder report), *not* in Form N-23C3A (which reports only the offer announcement). We download all N-CSR and N-CSRS filings for the 179 CIKs from EDGAR.

In total, 1,892 N-CSR and N-CSRS documents were retrieved, covering fiscal years 2018–2025.

L.3 Phase 3: Table Extraction

N-CSR filings report repurchase offer results in HTML financial-statement tables whose layout differs across fund families. We identify four primary formats, distinguished by which quantities each reports and how the oversubscription ratio is constructed. The BlackRock-style format reports shares tendered and shares purchased, each as a percentage of shares outstanding, so oversubscription is Tendered%/Purchased%. The Invesco-style format reports the offer percentage and the tendered percentage, giving Tendered%/Offer%. The First Trust-style format reports tendered and repurchased percentages by share class, giving Tendered%/Repurchased%. The Pioneer-style format reports the proration percentage applied to tendered shares, giving 1/Proration% when proration is below 100%.

L.4 Phase 4: Ratio Computation

For each successfully parsed repurchase offer, we compute the oversubscription ratio:

$$\text{Oversubscription}_{i,t} = \frac{\text{Shares Tendered}_{i,t}}{\text{Shares Accepted (Repurchased)}_{i,t}} \tag{34}$$

When formats report the offer percentage and tendered percentage rather than share counts, we use:

$$\text{Oversubscription}_{i,t} = \frac{\text{Tendered }\%_{i,t}}{\text{Offer }\%_{i,t}} \tag{35}$$

For Format D, where proration percentages are reported instead:

$$\text{Oversubscription}_{i,t} = \frac{1}{\text{Proration }\%_{i,t}} \quad \text{when Proration} < 100\% \tag{36}$$

When proration = 100%, the offer was not oversubscribed, and the ratio is computed directly from tendered and offered share counts.

L.5 Phase 5: Validation and Coverage

L.6 Coverage

The Regime 2 panel contains 4,107 fund-quarter observations with a binary oversubscription indicator derived from N-CSR repurchase disclosures across 48 CIKs with confirmed oversubscription events. The continuous oversubscription ratio (tendered over accepted) is available for only 30 of 512 documented offer events (5.9%), insufficient for a ratio-based regression, so the binary indicator is the dependent variable in all Regime 2 regressions. Filings that report results for multiple funds in a single document, or in narrative rather than tabular form, are reconciled by hand.

L.7 Validation

We verify oversubscription status for a random sample of 50 fund-offer observations against the original N-CSR filings. The agreement rate is 96% (48 of 50); the two discrepancies reflect ambiguous share-class aggregation, which we resolve by aggregating consistently to the fund level.

L.8 Opacity Extraction from N-CSR

For interval fund opacity, we extract aggregate Level 1, Level 2, and Level 3 dollar amounts from the "Notes to Financial Statements, Fair Value Measurements" section of N-CSR filings. This section typically contains a table with columns for Level 1, Level 2, Level 3, and Total fair values across asset categories. We compute opacity as Level 3 Total/(Level 1 Total + Level 2 Total + Level 3 Total) $\times$ 100. The extraction achieves 89% coverage (159 of 179 CIKs), with failures reflecting filings that report dollar amounts in narrative rather than tabular format.

Identification: Cross-Sectional Estimation in the Capped Vehicles. Opacity in the semi-liquid funds is essentially time-invariant, so fund fixed effects would absorb the between-fund variation that identifies the opacity gradient. Table B.2 reports explanatory power for oversubscription with and without fund fixed effects, and Table B.3 gives the formal between-within variance decomposition; together they motivate the cross-sectional specification used for the capped vehicles in Regime 2.

Table B.2: Robustness to Alternative Fixed Effects and Liquidity Controls

	DV: Any Oversubscription Event (binary)		
	(1) Fund+Quarter FE (Baseline)	(2) Quarter FE only (Cross-section)	(3) Fund+Quarter FE (+Lag return)
Crisis × High L3	−0.015	−0.075	−0.015
	(0.014)	(0.058)	(0.013)
High L3 indicator	0.014	0.095**	0.014
	(0.009)	(0.047)	(0.009)
Post-COVID indicator	0.013	0.000	0.000
	(0.012)	—	—
log TNA	0.000	0.033	0.000
	—	(0.024)	—
Fund age (quarters)	−0.003*	0.000	0.000
	(0.002)	—	—
Lagged quarterly return	—	—	0.000
			(0.000)
Fund fixed effects	Yes	No	Yes
Quarter fixed effects	Yes	Yes	Yes
N	1,974	1,975	1,974
R^2	0.933	0.043	0.933

Notes: This table tests the sensitivity of the oversubscription DiD result to alternative fixed effects structures. Column 1 is the baseline specification with fund and quarter fixed effects. Column 2 drops fund fixed effects, producing a cross-sectional comparison: the High L3 indicator enters significantly (+0.095, $p < 0.05$), reflecting that high-opacity funds are structurally more likely to be oversubscribed, but the crisis DiD interaction (−0.075) remains insignificant. Column 3 adds a lagged quarterly return control without changing the interaction. The R^2 drop from 0.933 to 0.043 in Column 2 illustrates that fund fixed effects absorb most between-fund variation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: Between-Within Variance Decomposition of Opacity and Outcome Variables

Variable	Standard Deviations σ_{overall}	σ_{between}	σ_{within}	σ_w/σ_b	% Total Variance Between	Within	Panel Dims n	$\bar{T}$
Panel A: Opacity Proxies								
Level 3 (%)	39.57	30.20	26.46	0.876	56.6%	43.4%	205	13.6
NAV Staleness (%)	20.14	20.70	0.00	0.000	100.0%	0.0%[a]	167	18.9
Ljung-Box Q^2	5.15	5.17	0.00	0.000	100.0%	0.0%[a]	179	32.0
Panel B: Outcome Variables								
Any oversubscription	0.355	0.316	0.162	0.512	79.3%	20.7%	206	32.0
Redemption rate (%)	3.567	2.304	2.810	1.219	40.2%	59.8%	185	16.7
Offer %NAV	3.016	4.193	1.726	0.412	85.5%	14.5%	78	19.2
Panel C: Controls								
Quarterly return (%)	7.534	2.486	7.222	2.905	10.6%	89.4%	190	17.2
Log TNA	1.556	1.561	0.000	0.000	n/a	n/a[b]	147	32.0

Notes: `xtsum` decomposition of total variance into between-fund (σ_{between}) and within-fund (σ_{within}) components. Panel structure: `xtset cik qint`, semi-liquid panel 2018Q1–2025Q3. **Columns**: σ_{between} = standard deviation of fund-level means around the grand mean; σ_{within} = standard deviation of demeaned values around fund means; σ_w/σ_b = ratio of within to between standard deviation; % Total Variance = $\sigma_j^2/(\sigma_b^2 + \sigma_w^2) \times 100$ for $j \in \{b, w\}$; n = number of unique funds (CIKs); $\bar{T}$ = average quarters observed per fund., Section 1b;.

Diagnostic interpretation: $\sigma_w/\sigma_b < 0.3$ indicates between-fund variation dominates, implying fund FE absorb most identifying variation in a *levels* regression. $\sigma_w/\sigma_b > 0.7$ indicates rich within-fund variation and fund FE are appropriate. For the primary opacity measure (`pct_level3_w`, ratio = 0.88), within-fund variation is substantial — fund FE are appropriate. For NAV staleness and Ljung-Box Q^2 (ratio $\approx$ 0.00), variation is entirely cross-sectional; fund FE would completely absorb these proxies in a levels specification. **Note on DiD specifications**: the ratio is irrelevant for causal DiD identification. When using a binary treatment indicator (`high_l3` $\times$ `covid`), fund FE absorb the main effect of `high_l3` but the interaction $\hat{\beta}_3$ is still identified from the *differential crisis response* across high- and low-opacity groups—no within-fund variation in the treatment is required. This justifies fund FE in the DiD regardless of the σ_w/σ_b ratio.

[a] NAV staleness and Ljung-Box Q^2 are structurally time-invariant in this panel ($\sigma_{\text{within}} \approx 0$): each fund's serial correlation in NAV returns is a persistent characteristic with essentially no quarter-to-quarter variation. These proxies are useful in cross-sectional specifications (no fund FE) but are entirely absorbed by fund fixed effects. [b] Log TNA also exhibits zero within-fund variation in this sub-panel (LOCF forward-fill from annual N-CEN filings), rendering the variance decomposition undefined; the between standard deviation (1.561) reflects substantial cross-fund size dispersion.

The Amplification Channel Across Shocks and Opacity Proxies. Table B.4 estimates the opacity-by-shock interaction separately for each opacity proxy (rows) and crisis window (columns), isolating which proxy carries the amplification and how it tracks the informational nature of each shock.

Table B.4: Opacity-Proxy by Shock: the Amplification Channel Across Crisis Episodes

	COVID	Evergrande	Fed 2022	UK LDI	BREIT	SVB
Level 3 (unobservable inputs)	$+0.090^{*}$	-0.074	$+0.064^{**}$	$+0.158^{***}$	$+0.031$	$+0.023$
Level 2 (partial opacity)	-0.003	-0.008^{***}	-0.001	$+0.005$	$+0.005^{**}$	$+0.008^{***}$
NAV staleness	-0.424^{***}	-0.866^{***}	-0.840^{***}	-1.314^{***}	-0.618^{***}	-0.773^{***}
Level 1 (transparency placebo)	$+0.003$	$+0.008^{***}$	$+0.001$	-0.005	-0.004^{**}	-0.008^{***}

Notes: Each cell is the opacity $\times$ shock interaction $\hat{\beta}$ from the headline difference-in-differences (the quarterly redemption rate on the opacity proxy, its shock interaction, fund and quarter fixed effects, and the headline control vector of lagged size, lagged return, Level 1 share, and a staleness $\times$ shock term), estimated separately for each opacity proxy (rows) and crisis window (columns) on the open-end panel (N = 187,731). Standard errors are clustered by fund. In the crisis episodes the Level 3 (unobservable-input) proxy loads positively, NAV staleness enters negatively (the stabilizing channel), and the Level 2 and Level 1 proxies are economically negligible, so the amplification is specific to assets priced with unobservable inputs rather than a property of the measurement choice. The headline COVID estimate survives the controls, two-way clustering, and a frozen pre-crisis opacity measure; the larger raw UK-LDI coefficient collapses under that frozen measure, and the Evergrande coefficient is negative. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

NAV Staleness as a Competing Opacity Proxy. Table B.5 re-estimates the headline interaction with NAV staleness in place of the Level 3 share, testing whether the amplification reflects unobservable-input valuation rather than stale pricing.

Opacity and Fragility Across Vehicle Types. Table B.6 pools open-end funds, interval and tender-offer funds, and business development companies in a single cross-section relating opacity to redemption fragility, and reports the business development company summary statistics referenced in the data section.

Table B.5: NAV Staleness as an Alternative Opacity Proxy

	Panel A: Opacity Predicts Staleness?			Panel B: Controlling for Return Smoothing		
	(1)	(2)	(3)	(4)	(5)	(6)
DV:	Return AC1	LB-Q^2	Staleness	Oversubscr.	Oversubscr.	Oversubscr.
Panel A: Does pct_level3 predict return-smoothing proxies? (Fund+Quarter FE)						
Level 3 (%)	0.000	0.000	—	—	—	—
	(—)	(—)				
Private assets (%)	—	—	0.000	—	—	—
			(—)			
N	1,159	1,793	2,351			
Panel B: Add return-smoothing controls to oversubscription DiD (Fund+Quarter FE)						
Level 3 (%)	—	—	—	0.000	−0.000	−0.000
				(0.000)	(0.000)	(0.000)
Return AC1	—	—	—	—	0.000	0.000
					(—)	(—)
Staleness (%)	—	—	—	—	—	0.000
						(—)
N				1,107	1,058	1,045
R^2				0.965	0.967	0.967
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Panel A tests whether Level 3 opacity (pct_level3_w) predicts NAV return-smoothing proxies—return autocorrelation (AC1), Ljung-Box Q^2 statistic, and a NAV staleness score—conditional on fund and quarter fixed effects. All three regressions yield zero coefficients, establishing that *within-fund* changes in Level 3 do not cause changes in return smoothing. The cross-sectional positive correlation between L3 and AC1 (AC1 = 0.013 / −0.073 / 0.178 across L3 terciles) therefore reflects time-invariant fund characteristics, not a dynamic causal link. Combined with the opposite crisis-period signs (L3×COVID = +0.118, staleness×COVID = −0.375; see Table A.4 and Section D), this confirms that the two measures are empirically and mechanistically distinct: Level 3 captures *informational opacity* (destabilizing) while NAV staleness captures *price smoothing* (stabilizing per Zeng 2017). Panel B adds staleness proxies progressively to the oversubscription DiD: the L3 coefficient is unchanged across all four columns (≈ 0.000), confirming return-smoothing is not the oversubscription mechanism. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Opacity and Fragility Across Semi-Liquid Vehicle Types: Pooled Cross-Section

Panel A: Structural Characteristics by Vehicle Type

	Interval Funds	Tender-Offer Funds	Non-Traded BDCs	Public BDCs
Redemption structure	Quarterly (N-23C3A)	Discretionary (SC TO-I)	Discretionary (SC TO-I)	Exchange (daily)
N (unique funds)	70	136	16	10
N (fund-quarters)	2,240	4,352	263	423
Level 3 (%) — p50	11.95[a]		96.93	n/a[b]
Level 3 (%) — p25	0.00		91.02	n/a
Level 3 (%) — p75	100.00		99.43	n/a
Any oversubscription (%)	52.5	1.7	n/a	n/a
NAV return — p50 (%)	—	—	3.16	—
Premium/discount — p50	—	—	—	−8.14%
Leverage ratio — p50	—	—	—	0.45

Panel B: Cross-Sectional Opacity → Oversubscription (Interval + Tender-Offer Funds)
DV = Any Oversubscription Event. Quarter FE only; no fund FE. Cluster(CIK). $N \approx 2{,}430$.

Opacity proxy:	(1) L3%	(2) PrivTotal	(3) PrivCredit	(4) PrivEquity	(5) PrivRE
Proxy coefficient	0.001	0.004***	0.005**	−0.006***	−0.013***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)
N	1,975	2,430	2,430	2,430	2,430
R^2	0.023	0.093	0.079	0.036	0.031

Panel C: Non-Traded BDC Cross-Section (Opacity → NAV Return; descriptive only)
DV = Quarterly NAV return (%). 16 non-traded BDCs, 2017Q4–2025Q3. SE cluster(fund_id).

Spec:	(1) No FE	(2) Quarter FE	(3) + Expense	(4) + NII
Level 3 (%)	0.022	0.036	0.027	−0.001
	(0.024)	(0.026)	(0.021)	(0.017)
N	214	213	165	152
R^2	0.004	0.517	0.656	0.750

Notes: **Panel A** compares structural characteristics across the four semi-liquid vehicle types in the sample. Interval fund and tender-offer fund Level 3 percentages are reported jointly from the combined semi-liquid panel (N = 2,780 fund-quarters with non-missing L3 data;). BDC Level 3 percentages are from the non-traded BDC panel. Public BDC Level 3 is not available in the Refinitiv panel; that panel instead provides market pricing and leverage data (). Oversubscription rates (interval: 52.5%; tender: 1.7%) are the unconditional means from fund+quarter FE regressions (intercept of the restricted model;).

Panel B reports cross-sectional regressions of any oversubscription event on alternative opacity/illiquidity proxies for the combined semi-liquid panel (interval + tender-offer funds), using quarter fixed effects but no fund fixed effects. No fund FE is included because opacity is largely time-invariant within funds (fund FE would absorb the cross-sectional identifying variation). Two proxies are positively significant: total private assets (+0.004***) and private credit (+0.005**), indicating structurally higher oversubscription rates in private credit–heavy funds. Private equity (−0.006***) and private real estate (−0.013***) are negatively signed because PE/RE interval funds tend to be larger, more established funds that set higher offer caps relative to typical demand. Level 3% is not cross-sectionally significant (identifying variation is largely time-invariant within funds, consistent with the R^2 jump from 0.043 to 0.933 when fund FE are added).

Panel C shows why BDC opacity cannot be exploited for causal inference. All 16 non-traded BDCs have Level 3 portfolios exceeding 91% (median: 96.93%), leaving virtually no cross-sectional variation in opacity. The $\hat{\beta}$ on Level 3% is near zero and statistically insignificant across all four specifications. The R^2 of 0.517–0.750 reflects quarter and expense-ratio variation, not opacity. Panel C is presented for completeness and comparability; no causal claims are made. The appropriate BDC analysis is descriptive: BDCs operate as *all-opaque* vehicles where the structural ceiling on redemptions ($\leq 3\%$ per quarter for tender BDCs) is the binding fragility constraint, not opacity variation within the BDC universe.

[a] Combined interval and tender-offer funds; distribution is bimodal (many funds near 0%, many near 100%). See main text for discussion. [b] Public BDC Refinitiv panel provides market cap and NAV per share; Level 3 fair-value detail not available from this source. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.