Public Windows into Private Credit: BDC Disclosures and Same-Borrower Valuation Discipline

Diogo Sampaio Lima

July 2, 2026

Abstract

Private credit has grown rapidly, yet outsiders rarely observe how individual private loans are valued. I use public filings by Business Development Companies (BDCs)—which hold private-credit-like loans and disclose position-level cost and fair value in their schedules of investments—to compare the values different BDCs report for the same borrower at the same reporting date.

Reported marks differ across holders of the same borrower, and the differences are systematic. A BDC whose own portfolio is deteriorating marks a shared borrower lower than other holders do at the same date: in the pre-2026 homogeneous-debt sample, a 10 percentage point increase in the markdown on the rest of the portfolio predicts a 1.26 percentage point larger markdown on the shared borrower, identified from changes. The direction is the surprise. The bank fair-value literature finds that pressured institutions delay loss recognition to protect capital; pressured BDCs instead mark shared borrowers more conservatively. The relation concentrates in junior and structured debt, is absent in equally volatile first-lien positions, and carries no detectable power to forecast the borrower's future deterioration, placing it in lender-side judgment rather than superior borrower information.

Reported private-credit marks are thus not pure borrower signals: two lenders holding the same loan report systematically different values for reasons tied to the holder, not the borrower. A source-filing audit confirms that cost and fair-value extraction from public filings is highly reliable, so public BDC disclosures provide a scalable instrument for measuring private-credit valuation at the position level.

JEL codes: G23, G24, G32, G12, M41.

Keywords: private credit; business development companies; fair value; valuation; mark dispersion; private debt; financial stability.

1 Introduction

Private credit has become a major source of financing for middle-market firms, but the values assigned to individual private-credit loans remain largely hidden. Unlike public bonds or broadly syndicated loans, many private-credit positions trade infrequently. Their reported values are therefore produced through fair-value processes that combine borrower information, models, third-party inputs, and manager judgment. This makes it difficult for outside investors, researchers, and regulators to know whether reported marks reflect borrower fundamentals, stale valuations, or the valuation practices of the institution reporting the asset.

These marks matter. They determine reported net asset values, affect leverage and borrowing-base constraints, enter performance fees and investor reporting, and influence when credit problems become visible. If private-credit marks adjust slowly or differ across holders, credit deterioration may be harder to monitor. If marks move with lender-level pressure, reported valuations reveal borrower risk together with the condition and valuation process of the intermediary holding the loan.

Business Development Companies provide a rare public window into this market. BDCs are investment vehicles that lend to and invest in middle-market firms. Many of their positions look economically similar to private-credit loans held by private funds. Unlike most private-credit funds, however, BDCs file public schedules of investments. These schedules disclose each position's borrower name, investment description, cost, fair value, and often principal, coupon, maturity, or lien information.

Throughout the paper, a "mark" is the value reported for a position in a BDC filing. The main measure is fair value divided by cost, or equivalently one minus this ratio as a markdown. A position marked at 100 is reported at cost. A position marked at 90 is reported 10 percent below cost. Cost is not always the same as principal or par, so the paper also checks cleaned fair-value-to-principal measures when principal is reliable.

I use these disclosures to compare how different BDCs value the same borrower at the same reporting date. For example, if three BDCs hold positions in Firm X at the end of 2024Q2, I can compare the values they report for those positions. Because the borrower and date are the same, common borrower news is held fixed. The remaining differences show how reported valuations vary across holders.

This comparison is useful but imperfect. The same borrower can issue different claims: first-lien loans, second-lien loans, revolvers, delayed-draw facilities, notes, preferred equity, or warrants. These claims can have different values for legitimate reasons. The paper therefore does not treat "same borrower" as automatically meaning "same security." I first document broad same-borrower dispersion, then restrict the comparison to debt-like and more comparable security groups. A source-filing audit shows that cost and fair-value extraction are highly reliable, but fully automated same-security matching is too noisy to carry the main identification claim.

The paper asks two questions. First, how much do reported marks differ across BDCs that hold the same borrower at the same date? Second, are these differences related to BDC-level pressure?

Pressure is measured using the markdown on the rest of the BDC portfolio. For a given borrower, I exclude that borrower from the BDC portfolio and calculate how much the remaining portfolio is marked down. This leave-one-borrower-out construction avoids mechanically using the borrower's own mark to explain itself.

The main result is that changes in BDC portfolio pressure predict changes in reported marks. In the pre-2026 homogeneous-debt sample, a 10 percentage point increase in leave-one-borrower-out portfolio markdowns predicts a 1.26 percentage point increase in the markdown reported on the shared borrower. The result is strongest in changes, which reduces the concern that the estimate simply reflects persistent BDC marking style. The direction runs against the bank fair-value literature, where pressured institutions delay loss recognition to protect capital: pressured BDCs instead mark shared borrowers more conservatively. The effect is also concentrated where valuation is most discretionary: the relation is economically large in junior and structured debt but absent in first-lien senior secured positions, whose marks are pinned near par—a gradient that points to discretionary valuation rather than a mechanical repricing of all claims. Consistent with discretion rather than superior information, cross-holder disagreement does not detectably forecast the borrower's future deterioration.

The contribution is a public-data method for studying private-credit valuations. Existing private-credit research is constrained by opacity: researchers usually observe aggregate fund returns, public BDC outcomes, or proprietary valuation inputs. This paper uses public BDC filings to build a position-level borrower-date valuation panel. The panel reveals same-borrower mark dispersion and shows that reported marks carry a lender-level component concentrated in discretionary, junior claims. The contribution is therefore not just a method but a fact about private-credit valuation: the same loan is marked differently across holders in a way that tracks each holder's portfolio condition, so reported marks and the net asset values built from them are not pure borrower signals.

Lower marks need not mean misvaluation: they may reflect better information, different loan economics, or more conservative inputs. The economic content of the result is therefore not misconduct but comparability. This is a within-borrower reduced-form relationship between lender-level portfolio deterioration and reported fair values, not a structural effect on true loan values.

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 explains the institutional setting and definitions. Section 4 develops the conceptual framework and testable predictions. Section 5 describes the data, sample construction, and validation. Section 6 documents mark dispersion. Section 7 gives a simple example and the empirical design. Section 8 presents the main internal-pressure results. Section 9 discusses robustness, denominator checks, and diagnostics, including a March 2026 stress-period extension. Section 10 presents market pressure as supporting evidence. Section 11 concludes and sets out scope and external validity.

2 Related Literature

The paper contributes first to work on private-credit growth, opacity, and financial stability. Reports by International Monetary Fund (2024), Financial Stability Board (2026), and Bank for International Settlements (2024), Federal Reserve staff notes, and surveys such as Block et al. (2023) emphasize rapid growth, limited loan-level transparency, valuation uncertainty, and bank-NBFI interconnections. These sources motivate the measurement problem but do not identify how lenders mark the same private borrower. The BDC position-level panel helps fill that gap for the public-reporting segment of private credit.

Much of the policy discussion around private credit is necessarily aggregate because the market is opaque. Aggregate evidence can show growth, leverage, bank links, and possible vulnerabilities, but it cannot show whether two holders mark the same borrower differently at the same date. The contribution here is to move one level down, from market size and aggregate risk to the position-level marks that make those risks visible or invisible in reported accounts.

The paper also relates to research on BDCs as private-credit intermediaries. BDCs are not the full market, but their public SOI disclosures make them uniquely useful for measurement. Chernenko, Ialenti, and Scharfstein (2025) use BDCs to study private-credit growth and balance-sheet structure. Suhonen (2024) studies direct-lending returns using BDC performance. This paper differs by using position-level same-borrower comparisons to study reported marks.

The BDC literature also clarifies external validity. BDCs are not a random sample of private-credit lenders; they are the part of the market where public position-level disclosures exist. That makes them both special and useful. The paper's claim is not that BDCs summarize every private-credit fund. The claim is that BDC disclosures allow a measurement exercise that is otherwise difficult to conduct at scale, and that the resulting same-borrower comparisons are informative about valuation discipline in an important public-reporting segment of private credit.

The closest valuation paper is Jang and Kim (2026). They study valuation intermediation and third-party pricing in private credit using proprietary appraisal data. This paper studies same-borrower cross-holder valuation differences using scalable public SEC BDC data. The two approaches are complementary: one studies the valuation-intermediary process; the other studies public cross-holder mark dispersion.

The distinction between proprietary appraisal data and public BDC filings matters for what each approach can deliver. Proprietary data can observe details that public filings miss, but they are hard to scale and hard for regulators or outside researchers to reproduce. Public filings are noisier, but they are observable, repeatable, and linked to actual reporting obligations. A public measurement system is therefore valuable even if it is less complete than proprietary valuation records.

The paper also connects to stale-pricing and private-asset valuation smoothing. Work on appraisal smoothing, private equity valuation, hedge-fund illiquidity, and private real estate shows that private asset marks may adjust slowly (Geltner, MacGregor, and Schwann, 2003; Getmansky, Lo, and Makarov, 2004; Cumming and Walz, 2010; Jenkinson, Sousa, and Stucke, 2013; Brown, Gredil, and Kaplan, 2019; Couts, 2023). The contribution here is cross-sectional rather than only time-series: public BDC filings allow contemporaneous comparisons across holders of the same borrower.

This cross-sectional feature separates the paper from pure return-smoothing tests. Stale-pricing papers often infer smoothing from serial correlation or delayed responses in aggregate returns. Here, the paper compares marks reported at the same date for the same borrower. The design does not require observing a transaction price. It asks whether public reports show contemporaneous cross-holder variation and whether that variation is related to lender-level pressure.

The mechanism connects the paper to the literature on discretion in fair-value accounting and loss recognition. Studies of bank financial reporting show that reported values reflect not only fundamentals but also managerial discretion that varies with the reporting institution's condition: discretion in loan-loss provisioning shapes the discipline of bank risk-taking (Bushman and Williams, 2012), write-downs are recognized with uneven timeliness across institutions (Vyas, 2011), valuations respond to bank incentives during stress (Huizinga and Laeven, 2012), and the role of fair-value rules in propagating distress remains debated (Laux and Leuz, 2010; Beatty and Liao, 2014). This paper brings that question to private credit, and the direction of the result is itself informative. The bank evidence typically finds that institutions under pressure use discretion to *delay* loss recognition and protect regulatory capital. BDC marks move the opposite way: more pressured BDCs report *lower* values on shared borrowers, and the effect is concentrated in junior and structured claims whose values are model- and judgment-based. Under Investment Company Act valuation requirements, board and auditor oversight, and public-market scrutiny, the discretionary component of private-credit marks appears to move toward conservatism rather than forbearance.

Finally, the paper follows the architecture of empirical macro-finance measurement and identification papers. Gilchrist and Zakrajsek (2012) construct a bottom-up credit-spread object before using it for macro analysis; Khwaja and Mian (2008) and Jimenez et al. (2012) use within-borrower variation to separate lender supply from borrower demand; and Krishnamurthy and Vissing-Jorgensen (2011) use channel-specific asset-price predictions. This paper brings that three-part architecture to private-credit valuation using public BDC filings: it builds a measurement object from granular data before studying its economic content, uses borrower-date comparisons to absorb borrower fundamentals, and maps competing mechanisms to distinct predictions rather than reading a single coefficient in isolation.

3 Institutional Background

BDCs are closed-end investment companies that finance small and middle-market firms. In practice, modern BDCs are an observable public-reporting segment of private credit. They invest primarily in privately originated debt, including first-lien senior secured loans, second-lien loans, unitranche facilities, revolvers, delayed-draw loans, and notes. They also hold preferred equity, common equity, warrants, and other structured claims.

BDCs are not ordinary private funds. They file periodic SEC reports and disclose detailed schedules of investments. These schedules list investment names, position labels, cost, fair value, principal, coupon or interest terms, maturity, lien category, non-accrual and PIK language where reported, industry labels, and filing metadata. The disclosures are public, repeated, and position-level, which makes them useful for measurement. They are also messy. Borrower names are not standardized across filers, security labels combine borrower and tranche information, and fields differ across BDCs and over time.

3.1 Why BDCs Are Informative

The BDC setting is useful because it combines three features that rarely appear together in private credit. First, BDCs hold the type of middle-market loans that make private credit economically important. Second, they report holdings at the position level rather than only at the portfolio level. Third, multiple BDCs often hold claims on the same borrower at the same reporting date. The intersection of these features creates a within-borrower comparison that is not available in most private-credit datasets.

The comparison is economically natural. Many middle-market borrowers borrow from lender groups rather than from a single lender. Several BDCs may therefore appear in public filings with claims on the same borrower during the same quarter. If all holders valued the exposure identically, the public filings would mostly confirm that BDC marks are mechanically pinned down by borrower identity. If marks differ across holders, the filings reveal a dimension of private-credit valuation that is usually hidden: the cross-holder component of reported fair value.

The setting also has a useful timing structure. BDCs report periodically, and the schedule of investments is tied to a reporting date. This allows the paper to compare marks at the same borrower-date rather than comparing one holder's mark before an event with another holder's mark after the event. The timing is still imperfect because filing dates differ and valuation committees may use information available before the filing, but the report-date comparison is substantially tighter than comparing annual fund returns or aggregate NAVs.

The SEC BDC Data Sets provide the starting point for the empirical work. The SEC describes these data as extracted from XBRL-tagged BDC periodic reports and other BDC disclosures. The datasets include schedule-of-investments reports, detailed financial datasets, and summary non-financial datasets in flat-file format. The SEC also cautions that these data are automatically extracted from registrant submissions and are not a substitute for full filings. This caveat is central for the paper: public BDC data are unusually valuable, but parser and classification validation are part of the research design.

3.2 What the Filings Reveal and What They Do Not

The schedule of investments reveals enough to build the empirical object, but not enough to remove every interpretation problem. It typically reports the borrower or investment name, security description, cost, fair value, maturity, interest rate information, and sometimes lien, PIK, non-accrual, or affiliation language. These fields allow the construction of borrower-date cells and security-comparability restrictions. They also allow the paper to distinguish debt-like positions from equity-like positions and to remove many observations where a fair-value-to-principal ratio is not economically meaningful.

The filings do not reveal the full private credit agreement, covenant package, amendment history, lender economics, or all borrower information available to the holder. Nor do they reveal whether two positions with similar labels have exactly the same economics. A BDC may hold a delayed-draw commitment while another holds a funded term loan; one may hold a first-out tranche while another holds a last-out piece; one may have participated in an amendment that changed economics differently across lenders. These limitations motivate the homogeneous-debt and security-cell restrictions. They also explain why the paper interprets strict cells as stronger evidence only where the filing-level validation audit supports them.

Fair value is the key accounting object. A BDC reports fair value for each position under valuation procedures that can combine market quotes, valuation models, borrower information, adviser inputs, third-party valuation support, board oversight, and auditor review. For many private-credit loans, observable secondary-market transactions are sparse. Reported marks are therefore accounting outcomes, not clean market prices.

3.3 Valuation as an Economic Outcome

Because many positions do not trade frequently, reported fair value is not only a measurement input; it is an economic outcome of the valuation process. That outcome affects reported NAV, leverage capacity, fee bases, investor communication, and the apparent timing of credit deterioration. A mark close to cost can mean that the borrower is healthy, but it can also mean that the position is hard to price, that adverse information is only partly reflected, or that valuation inputs differ across holders. This is why the empirical object is the reported mark itself.

For this reason, the paper uses markdown-to-cost as the main outcome. Cost is not a perfect economic denominator, because positions can be purchased above or below par and costs can differ across holders. But cost is the consistently available denominator in the public BDC data, and it maps directly into the reported cost and fair-value fields that BDCs disclose. Principal is a natural alternative for loans, but it is missing for many rows and unreliable for commitments, revolvers, hybrid instruments, and rows with scale mismatches. The principal-based tests are therefore used as robustness checks after explicit cleaning, not as the main outcome.

Marks can differ across holders for legitimate reasons. Two BDCs may hold different facilities of the same borrower, enter at different prices, report different cost bases, have different amendments or funding status, receive different borrower information, or apply different valuation processes. This is why the paper does not interpret mark dispersion mechanically as misvaluation. The research design first measures dispersion, then asks whether the component related to BDC-level pressure survives progressively tighter comparability restrictions.

BDCs are useful but not representative of all private credit. They face public reporting, BDC regulation, and in many cases public-market scrutiny. Private credit funds outside the

BDC sector may differ in leverage, investor base, liquidity structure, reporting frequency, and valuation governance. The paper's external-validity claim is therefore limited. BDCs are a transparent laboratory for studying valuation in private credit, not a complete map of the private credit universe.

The limitation cuts the other way for identification: the paper does not need BDCs to represent every private-credit vehicle. It needs BDC filings to provide repeated public observations of comparable private-credit claims. The same-borrower design uses those observations to study valuation discipline in the part of the market where outside researchers can see position-level data.

Three requirements follow from this setting. The public filings must be turned into a credible borrower-date panel; security heterogeneity must be measured and reduced, even if not eliminated; and the remaining variation must be shown not to be a parser artifact. The data and validation sections take up each in turn.

4 Conceptual Framework and Testable Predictions

The object of interest is the reported mark that BDC b assigns to a position in borrower i, security or tranche s, at report date t. A simple decomposition is:

$$\text{Mark}_{i,b,s,t} = F_{i,t} + S_{s,t} + V_{b,t} + \epsilon_{i,b,s,t}.$$

The term $F_{i,t}$ captures borrower-date fundamentals. Borrower-date fixed effects absorb this component. The term $S_{s,t}$ captures security-specific features: seniority, lien, coupon type, revolver status, maturity, funding status, and other tranche economics. Homogeneous-debt, first-lien, flexible-security, and strict-security cells reduce variation in this component. The term $V_{b,t}$ captures BDC-level valuation behavior, pressure, or governance. Internal pressure and market pressure proxy for this component. The residual contains remaining measurement error, unmatched tranche differences, and filing noise.

This decomposition clarifies what the empirical design can and cannot identify. Borrower-date fixed effects make the comparison close to the within-firm logic in banking papers: they compare different intermediaries exposed to the same borrower at the same date. But same borrower does not always mean same security. Security-cell restrictions therefore tighten the comparison along economically relevant dimensions. Pressure changes strip out persistent BDC-level style.

The decomposition organizes the competing readings of a pressure–mark relation around the design lever that discriminates among them; Table 1 collects the mechanisms, their predictions, and the corresponding tests.

Two readings imply that the relation should survive tighter comparisons. Under *valuation discipline*, a rise in BDC pressure draws closer scrutiny from boards, auditors, valuation committees, and public investors, so changes in pressure predict lower marks for shared borrowers. This requires only that the valuation process carries discretion and that scrutiny moves with the intermediary's condition, not intentional misstatement. Because that discretion is concentrated in the capital structure, the channel makes a sharp cross-sectional prediction: the relation should be largest in junior and structured claims and absent in first-lien debt, whose marks are anchored by collateral and observable prices. Under *superior information*, a pressured BDC marks lower because it holds claims with worse private information about the borrower. This is the leading alternative, and borrower-date fixed effects do not dispose of it when holders own different parts of the capital structure. But it makes a forecast that discipline does not: if disagreement encoded borrower information, the most pessimistic mark would predict the borrower's future deterioration. Table 9 shows no detectable forecasting content.

Three further readings concern persistence rather than borrower fundamentals. Portfolio composition makes pressure a proxy for sector, vintage, or sponsor structure; valuation governance makes it a proxy for stable differences in adviser structure, board oversight, and audit process; market monitoring makes it investor skepticism priced into listed BDC stock and NAV discounts. All three predict persistent, partly anticipated pressure *levels* rather than contemporaneous co-movement, which is why the weight falls on the change specification and market pressure enters as supporting evidence. The change specification asks the narrower question: when the rest of a BDC's portfolio deteriorates from one period to the next, does its reported mark on a shared borrower deteriorate relative to other holders at the same date?

The sign of the relation shapes the language. More pressured BDCs mark shared borrowers lower, a direction consistent with discipline, scrutiny, or lender-specific information rather than an incentive to hide losses; the paper therefore speaks of valuation discretion and discipline, not manipulation, and its tests measure systematic mark differences rather than assign blame. The mechanisms can coexist—a pressured BDC may hold worse information and also face greater scrutiny—so the strategy does not assign the coefficient to one channel. It disciplines interpretation by asking which predictions survive successively tighter comparisons.

5 Data, Sample Construction, and Validation

The analysis uses SEC BDC structured datasets and EDGAR filing links from 2022Q4 through 2025Q3. March 2026 is available but excluded from headline inference because the

Table 1: Mechanisms, Predictions, and Tests

Mechanism	Prediction	Test
Valuation discipline	Mark changes move with changes in BDC portfolio pressure.	Internal-pressure change regressions.
Valuation discretion (seniority gradient)	The relation concentrates in junior and structured debt and is absent in first-lien debt, where marks are pinned near par.	Seniority-class change regressions.
Information	Effects weaken as comparability tightens; and cross-holder disagreement forecasts the borrower's future deterioration.	Security-cell estimates; forecasting tests of dispersion and the pessimistic mark on future markdown and nonaccrual.
Portfolio composition	Level pressure is persistent and future pressure predicts current marks.	Lag-current-future and change specifications.
Governance or valuation style	Persistent BDC differences explain part of mark variation.	BDC fixed effects and dynamic diagnostics.
Market monitoring	Public BDC stress predicts lower marks in broad cells.	Stock-pressure, NAV-discount, and residualized-pressure tests.
Denominator measurement	Cost-based and principal-based marks differ when principal is unreliable.	Principal taxonomy and cleaned FV/principal tests.

SEC schema changes. The raw files include schedule-of-investments data, detailed financial data, summary non-financial data, and filing identifiers. The parser extracts borrower labels, BDC identifiers, report dates, filing dates, position labels, cost, fair value, principal amount, interest-rate fields, maturity, lien category, and filing URLs.

5.1 Raw Data and Sample Funnel

The raw data are not a ready-made econometric panel. They are filing-level records that must be converted into position observations, matched to reporting periods, assigned to BDC families, and filtered to current-period holdings with usable cost and fair-value fields. The sample funnel therefore does two jobs. It reports the number of observations lost at each step, and it tells the reader which empirical comparison remains after each restriction.

Table 2: Sample Construction

Sample step	Observations
Raw SEC SOI rows, 2022Q4-2025Q3	1,257,431
Candidate position rows	1,148,545
Validated SOI position rows	374,993
Clean pre-2026 regression panel	374,993
Cross-BDC borrower-date identifying rows	168,137
Borrower-BDC-period identifying rows	106,429
Homogeneous-debt borrower-date cells	29,457
Headline homogeneous-debt observations	84,191
Broad public-BDC market-pressure observations	31,289

This table shows how SEC schedule-of-investments rows become the main analysis samples. No regression estimates are reported.

Table 3: Regression-Sample Attrition

Step	Observations	Dropped since prior step
Headline homogeneous-debt rows	84,191	0
Require prior-quarter BDC-borrower mark	65,671	18,520
Require prior-quarter leave-one-out pressure	65,671	0
Require nonmissing borrower-date and BDC fixed effects	65,671	0
Final homogeneous-debt change regression rows	65,671	0

This panel reconciles the headline homogeneous-debt sample with the first-difference regression sample. Counts are generated from the analysis panel.

Table 2 shows the sample funnel. The parser begins with 1.26 million raw SEC SOI rows from 2022Q4 through 2025Q3. Current-period and validation filters reduce the sample to 374,993 clean position rows. The same-borrower designs require borrower-date cells with at least two BDC holders; this produces 168,137 cross-BDC borrower-date identifying rows. The headline homogeneous-debt sample contains 84,191 observations. Table 3 shows why the first-difference regression uses 65,671 observations: the drop is driven by the requirement that the same BDC-borrower pair has a prior-quarter mark.

The funnel also clarifies the distinction between three samples that otherwise look similar. The broad borrower-date sample is the best sample for measuring how much cross-holder dispersion appears in public filings. The homogeneous debt sample is the preferred sample for the main internal-pressure result because it removes equity-like positions and obvious instrument differences. The strict-security sample is the most comparable sample, but it is most dependent on correct parsing of security attributes and therefore requires filing-level validation before it can carry the full evidentiary burden.

Borrower matching is the central data-construction challenge. The parser normalizes legal suffixes, punctuation, capitalization, and common instrument descriptors to construct borrower match keys. This process creates a scalable panel but cannot eliminate all errors. False positives can arise when different borrowers share a sponsor or legal stem. False negatives can arise when the same borrower appears under different legal names. The validation sample is therefore not optional; it is part of the identification strategy.

5.2 Borrower Matching

Borrower matching follows the logic of an empirical credit-registry paper, but with public filing text rather than administrative identifiers. The parser standardizes capitalization, strips legal suffixes, removes common instrument phrases, and builds a borrower key intended to capture the operating borrower rather than the security label. The procedure is conservative enough to create many repeated borrower-date cells, but it necessarily trades off false matches against missed matches. The paper therefore treats borrower matching as a measured source of uncertainty. The same-borrower tables and validation sample are meant to expose this uncertainty rather than hide it.

Security classification is the second challenge. The classifier uses position labels, principal amount, interest-rate fields, lien category, PIK and non-accrual flags, and text markers to classify debt-like positions, equity-like positions, seniority, rate type, and maturity buckets. The homogeneous-debt sample requires mostly debt-like rows and excludes equity-like rows. Flexible security cells condition on borrower-date, debt class, and seniority. Strict security cells add rate type and maturity bucket.

Online Appendix Table A.1 reports the coverage of these classifications. The broad borrower-date analysis panel—the 106,429 borrower-BDC-period rows that follow from the 168,137 cross-BDC identifying rows in the funnel once positions are collapsed to the borrower-BDC-period level—contains 37,787 comparison cells, 8,006 borrowers, and 164 BDCs. The homogeneous-debt restriction keeps 84,191 rows. The first-lien-only restriction is much smaller, with 42,972 rows, which is why it is a robustness sample rather than the baseline. Flexible and strict-security cells retain broad coverage while absorbing increasingly detailed security attributes.

The online appendix reports the same tradeoff visually. More comparable cells reduce the chance that the result is driven by obvious instrument differences, but they also create a stronger need for validation of the underlying security labels. The same tradeoff runs through the paper: the broad sample buys measurement, the narrow samples buy credibility.

The main outcome is markdown to cost, defined as one minus fair value divided by cost. Cost is widely available, but it is not always par. Principal is economically attractive because many loans are originated near par, but it is frequently missing or economically non-comparable. The principal-denominator audit reviews 373,363 rows—the 374,993 validated rows minus 1,630 positions where a principal-based outcome is undefined—and identifies 75,182 problem rows, mostly missing principal. Other categories include near-zero principal, likely scale mismatch, partially funded or unfunded commitments, delayed-draw facilities, revolvers, equity or hybrid securities, and restructuring or PIK/nonaccrual cases.

Online Appendix Table A.2 explains why the cost-based mark is the main dependent variable. FV/cost is available for the full analysis sample and its distribution has plausible tails. The headline specifications use the raw markdown for positions with positive cost and do not winsorize the dependent variable; winsorizing FV/cost at the 1st and 99th percentiles is reported among the robustness checks and leaves the result unchanged. Raw FV/principal, by contrast, is not usable as a main outcome: the mean and standard deviation are dominated by a small number of denominator failures even though the median is economically reasonable. Data issues of this kind can produce spurious results when they are not separated from the economics.

Internal pressure is the leave-one-borrower-out BDC portfolio markdown. For each borrower-BDC-period observation, the pressure measure excludes that borrower from the BDC portfolio. This avoids mechanically using the dependent variable in the pressure measure. Internal pressure changes are quarter-to- quarter changes in this measure. Market pressure is negative BDC stock performance over the prior 90 days for public BDCs. NAV discount and residualized market pressure are secondary diagnostics.

5.3 Validation Strategy

Online Appendix Table A.3 reports summary statistics. Most positions are marked close to cost, but the tails are large enough to matter. The 5th percentile of FV/cost is 0.765. The median leave-one-borrower-out portfolio markdown is small but positive. This distribution is typical of private-credit marks: many positions sit near par, while credit deterioration appears in the tails.

Validation has two layers. First, the paper validates extraction by opening the source SEC filings behind the validation sample and searching the filing text for borrower tokens, security-label tokens, and the reported cost, fair value, and principal values. Second, it evaluates whether rows are usable for borrower-date and strict-security comparisons. The strict-security validation sample contains 4,000 rows stratified by strict cells, large mark changes, large internal-pressure changes, first-lien cells, public-BDC pressure cells, high-dispersion cells, March 2026 cells, principal outliers, extreme FV/cost rows, and suspicious borrower matches.

The validation thresholds are demanding because the strict-cell design carries the paper's strongest comparability claim. The borrower match must be accurate enough that borrower-date fixed effects are meaningful. The cost and fair-value fields must be accurate enough that the dependent variable is not a parser artifact. Debt and lien classifications must be accurate enough that homogeneous-debt and first-lien restrictions have economic content. The strict-security cell must be usable often enough that the remaining regression sample is not a selected set of unusually clean filings. These thresholds convert data construction from a background task into an explicit credibility condition. The filing audit passes the numeric extraction thresholds but fails the borrower-match, debt/equity, and strict-cell usability thresholds. The split cuts both ways: public BDC filings are strong enough to support a scalable mark-measurement contribution, while the strictest same-security design remains too noisy to carry the main evidence.

The validation strategy is staged because different claims require different evidence. Filing-snippet checks can validate that the parser extracted cost and fair value from the filing text. They cannot by themselves prove that a strict security cell is economically identical across holders. The paper's preferred claim is therefore the more conservative one: internal-pressure changes predict reported values for shared borrowers in homogeneous-debt and flexible-security comparisons, with strict cells used as supporting robustness rather than headline evidence.

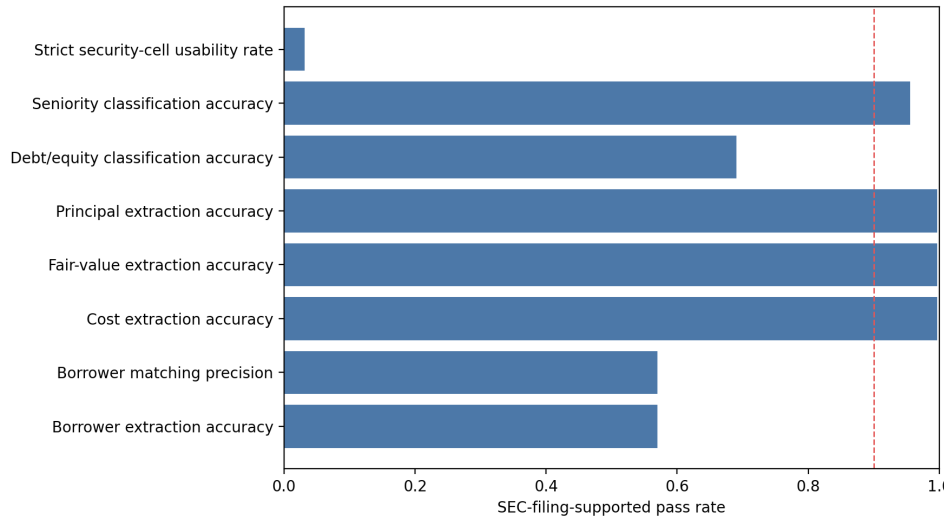

Figure 1: Source-filing validation audit. Cost and fair-value extraction are near-perfect in the audit sample, while borrower matching and fully automated strict-security usability are too noisy to support strict-security cells as the headline design.

6 Mark Dispersion in Public BDC Filings

Before studying whether BDC pressure predicts marks, I first ask whether same-borrower mark dispersion is empirically meaningful and measurable in public filings. This is the measurement object of the paper. If BDC marks were mechanically identical across holders, pressure regressions would have little economic content. The data show otherwise.

The dispersion exercise also clarifies the object being measured. The paper is not trying to infer an unobserved transaction price for every private loan. Instead, it measures cross-holder differences in reported fair value for borrowers that appear in more than one public BDC filing. That object is important even when none of the marks is a transaction price, because reported fair values are the numbers that enter NAV, performance measurement, and public disclosure.

Table 4: Rule-Screened Same-Borrower Example

Borrower	Report date	BDC	Adviser family	Seniority	Cost ($m)	Fair value ($m)	FV/cost	Internal pressure
Asurion LLC	2023-03-31	OWL ROCK CAPITAL CORP	Blue Owl	first_lien	7.95	6.54	0.823	0.008
Asurion LLC	2023-03-31	ARES STRATEGIC INCOME FUND	Ares	first_lien	1.51	1.43	0.945	0.003
Asurion LLC	2023-03-31	HANCOCK PARK CORPORATE INCOME, INC.	Hancock Park	first_lien	1.69	1.67	0.992	0.054
Asurion LLC	2023-03-31	OWL ROCK CORE INCOME CORP.	Blue Owl	first_lien	13.93	14.32	1.028	0.003

The same-borrower examples are illustrative and rule-screened. They are not hand-adjudicated same-security evidence and do not imply that any reported mark is incorrect. The empirical analysis uses the full sample.

Table 4 reports a rule-screened same-borrower example. The example is selected from first-lien debt-like borrower-date cells with multiple BDC holders. It illustrates the comparison used in the paper, but it is not hand-adjudicated same-security evidence and is not evidence of misreporting.

The examples serve the same purpose as institutional examples in benchmark empirical papers: they make the identifying variation concrete. The reader can see that the unit of comparison is not an abstract BDC average but a borrower that appears in multiple public portfolios at the same reporting date. The examples also reveal why the paper must be cautious. Even when the borrower and lien language look similar, the filings may not disclose every contractual feature that determines value. The examples therefore motivate both the measurement contribution and the validation requirement.

Table 5: Mark Dispersion Across BDCs Holding the Same Borrower

Sample	Cells	Borrowers	Mean holders	Median SD (pp)	P75 SD (pp)
Borrower-date	37,787	8,006	2.82	0.30	1.24
Narrow same-security diagnostic	34,073	6,373	2.69	0.16	0.76

Dispersion is measured in percentage points of markdown, where markdown equals one minus fair value divided by cost.

Table 5 summarizes dispersion. Borrower-date cells have a median markdown standard deviation of 0.30 percentage points and a 75th percentile of 1.24 percentage points. Strict-security cells have a median of 0.16 percentage points and a 75th percentile of 0.76 percentage points. Security heterogeneity explains part of dispersion, but a measurable cross-holder component remains after tightening the cell definition.

The tail of the distribution matters more than the median. A median dispersion near zero is consistent with many private-credit positions being marked close to cost in normal times. The economically relevant question is whether the public data reveal cells in which holders disagree by enough to matter for NAV and monitoring. The 75th percentile and the same-borrower examples show that they do, which is why I treat dispersion as a measurement fact rather than a side result.

Online Appendix Table A.4 shows that dispersion varies with the number of BDC holders in the cell. Borrower-date cells with two BDC holders have a median standard deviation of 0.22 percentage points, while cells with six to ten holders have a median of 0.78 percentage points. This increase partly reflects more opportunities to observe extreme marks, but it also shows why multi-holder borrowers are valuable for measuring valuation dispersion.

The holder count shapes what each cell can show. A two-holder cell can reveal disagreement, but not whether one mark is unusual relative to a broader cross-holder distribution. Cells with more holders provide a richer comparison set and make it easier to detect whether dispersion reflects a single outlier or a broader disagreement. At the same time, high-holder cells are not random; they likely involve larger, more widely financed borrowers. The paper therefore uses holder count descriptively and does not treat it as a source of exogenous variation.

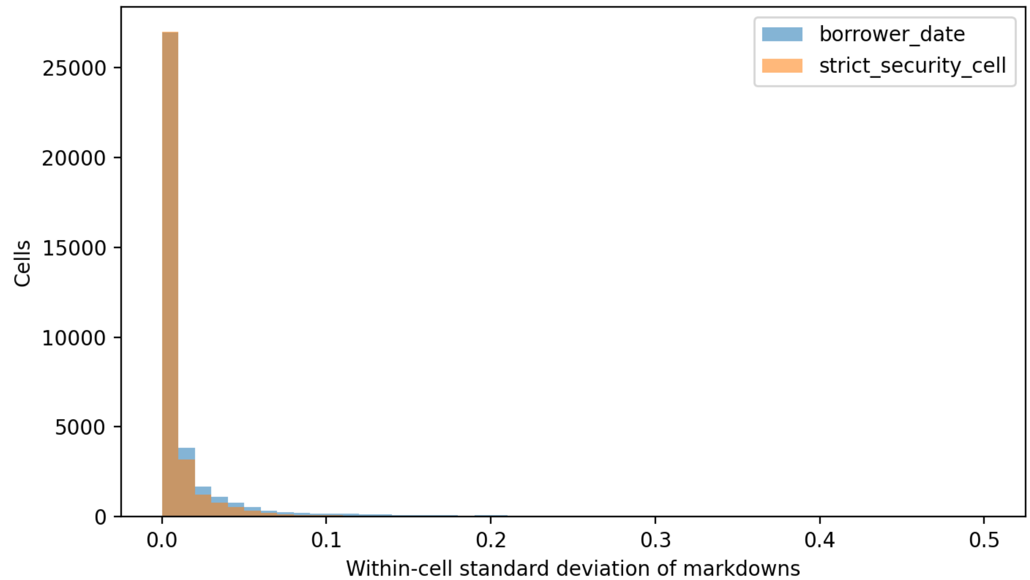

Figure 2: Same-borrower mark dispersion is concentrated near zero but has an economically visible upper tail. Dispersion falls in narrow same-security diagnostic cells, showing that security heterogeneity matters but does not account for all cross-holder differences.

7 Simple Example and Empirical Design

7.1 A Simple Example

Consider a borrower, Firm X, that appears in the filings of three BDCs in the same quarter. BDC A reports its position at 100 cents on cost, BDC B reports 96, and BDC C reports 93. The borrower-date comparison asks whether these differences are related to the condition of the BDCs reporting them. If BDC C's broader portfolio has also deteriorated more than the portfolios of BDCs A and B, this suggests that reported marks contain a lender-level component. The empirical analysis repeats this comparison across thousands of borrower-date cells and then asks whether the pattern survives when the comparison is restricted to debt-like and more comparable securities.

7.2 Empirical Design

The empirical difficulty is to separate lender-level pressure from borrower fundamentals, security heterogeneity, and persistent BDC valuation style. A BDC may report a lower mark because the borrower is weaker, because the BDC holds a junior or differently structured position, because the BDC has different information, or because pressure affects valuation behavior. The design addresses these threats sequentially.

The structure is analogous to within-borrower credit-supply designs, but the interpretation is different. In loan-supply papers, the borrower-date fixed effect separates borrower demand from lender balance-sheet shocks. Here, the borrower-date fixed effect separates borrower fundamentals from lender-level marking variation. The remaining problem is not loan demand; it is whether the holders being compared own economically comparable securities and whether BDC pressure captures persistent valuation style. The security-cell and dynamic tests are designed for those two problems.

Table 6: Identification Threats and Design Responses

Threat	Why it matters	Design response
Borrower fundamentals	Weak borrowers should have lower marks for all holders.	Borrower-date fixed effects.
Security heterogeneity	Same borrower can issue economically different claims.	Homogeneous debt, first-lien, flexible, and strict cells.
Persistent BDC style	Some BDCs may mark conservatively in every period.	BDC fixed effects, changes, lead/lag diagnostics.
Portfolio composition	Pressure may proxy for sector, vintage, or sponsor exposure.	Leave-one-borrower-out pressure and dynamic tests.
Cost-basis differences	FV/cost can differ when entry prices differ.	Cleaned FV/principal and denominator taxonomy.
Parser or match error	False borrower or security matches can create false dispersion.	Filing snippets and validation sample.
Market-pressure endogeneity	Stock prices can anticipate future marks.	Treat market pressure as supporting evidence.

Table 6 summarizes the logic. The empirical strategy is not a single fixed-effect regression. It is a sequence of tests that asks whether the same economic relation survives after the most obvious alternative explanations are disciplined. This is why the paper reports borrower-date, homogeneous-debt, flexible-security, and strict-security results side by side, and why it interprets pressure changes more strongly than pressure levels.

The level specification is:

$$Markdown_{i,b,t} = \alpha_{i,t} + \gamma_b + \beta Pressure^{-i}_{b,t} + \varepsilon_{i,b,t}.$$

The dependent variable is one minus fair value divided by cost for borrower i, BDC b, and report date t. Borrower-date fixed effects $\alpha_{i,t}$ absorb borrower-date fundamentals. BDC fixed effects γ_b absorb persistent BDC differences. The coefficient β measures whether BDCs with more leave-one-borrower-out pressure mark the same borrower lower at the same date.

The main specification uses changes:

$$\Delta Markdown_{i,b,t} = \alpha_{i,t} + \gamma_b + \beta \Delta Pressure^{-i}_{b,t} + \varepsilon_{i,b,t}.$$

This comparison asks whether deterioration in the BDC's other portfolio positions predicts deterioration in its mark on a shared borrower. The identifying assumption is that, conditional on borrower-date and BDC fixed effects, changes in leave-one-borrower-out pressure are not driven by unobserved changes in the specific security being valued.

The security-cell specification tightens the comparison:

$$\Delta Markdown_{i,b,s,t} = \alpha_{i,s,t} + \gamma_b + \beta \Delta Pressure^{-i}_{b,t} + \varepsilon_{i,b,s,t}.$$

Here the fixed effect absorbs borrower, report date, and security-cell features. Flexible and strict cells reduce tranche heterogeneity. These estimates are more credible but rely more heavily on correct classification, which is why the validation audit determines the claim hierarchy.

The market-pressure specification replaces internal pressure with public BDC market stress:

$$Markdown_{i,b,t} = \alpha_{i,t} + \gamma_b + \beta MarketPressure_{b,t} + \varepsilon_{i,b,t}.$$

Market pressure is externally priced but not a clean shock. It reflects expected future marks, risk premia, leverage, liquidity, and market sentiment. The paper therefore treats market pressure as broad stress evidence, not the main evidence.

Standard errors are clustered at the BDC level where the identifying variation is lender-level pressure. This is the conservative clustering dimension for the main regressions because the pressure measure varies at the BDC-period level and may be correlated across borrowers within the same BDC. The borrower-date fixed effects absorb common borrower news and market conditions at the reporting date, while BDC fixed effects absorb persistent cross-BDC differences in reporting style, portfolio mix, and valuation governance.

The identifying assumption for the change specification is narrower than for the level specification. It does not require that high-pressure and low-pressure BDCs are otherwise comparable in levels. It requires that changes in the rest of a BDC's portfolio markdown, after borrower-date and BDC fixed effects, are not entirely driven by unobserved changes in the exact security being valued. This is why the leave-one-borrower-out construction and security-cell tests are central. The pressure measure excludes the borrower, and the security-cell tests ask whether the result survives among more comparable claims.

8 Main Results: Internal-Pressure Changes and Same-Borrower Marks

Table 7: Main Result: Portfolio-Pressure Changes and Reported Marks

Row	Borrower-date	Homogeneous debt	Flexible security	Strict security diagnostic	Cross-adviser families	Family collapsed
Pressure-change estimate	0.054	0.126***	0.055**	0.062**	0.130***	0.077***
Standard error	(0.056)	(0.031)	(0.026)	(0.027)	(0.044)	(0.026)
Observations	83,366	65,671	73,849	72,980	32,776	26,298
Cells	29,715	23,018	27,675	27,258	–	–
BDCs / families	152	129	128	121	84	84
Role	Broad measurement	Main estimate	Main robustness	Diagnostic only	Family robustness	Family robustness

The homogeneous-debt column is the headline estimate. The strict-security column is diagnostic because automated strict matching is too noisy to carry the main design. Standard errors are in parentheses.

Table 7 reports the core results. The main estimate is the internal-pressure change coefficient. In the homogeneous-debt sample, the coefficient is 0.1264 with a standard error of 0.0314. A 10 percentage point increase in leave-one-borrower-out portfolio markdowns predicts a 1.26 percentage point increase in the markdown reported on a shared borrower.

This magnitude is economically visible in a market where many positions are reported near cost. Quarterly changes in portfolio pressure are themselves small—a standard deviation of about 0.6 percentage points—so a one-standard-deviation increase predicts roughly a 0.08 percentage point larger markdown on the shared borrower. That is about a quarter of the median same-borrower mark dispersion documented in Section 6, and it accumulates as portfolio stress persists across quarters. The ten-percentage-point figure is a linear scaling convention for comparability, not a typical one-quarter move; it does not imply that every 10 percentage point movement in portfolio pressure mechanically changes all marks by 1.26 percentage points. It says that, within the same borrower and report date, the BDC whose other portfolio positions deteriorate more also reports a lower mark on the shared borrower.

The magnitude is also meaningful at the portfolio level. A BDC with a large portfolio of near-par loans does not need every position to move by many percentage points for NAV to change materially. A one percentage point lower fair value on a broad set of debt positions can matter for leverage ratios, reported performance, and investor perception. The coefficient should not be read as a mechanical NAV calculation, because the regression is within shared borrowers and not a full portfolio simulation. The estimate is a within-borrower reduced-form relationship between lender-level portfolio deterioration and reported fair values, not a structural effect on true loan values; within that scope, it shows that the pressure-related component of marks is large enough to be economically visible in reported valuation data.

The level result is positive but lower in the claim hierarchy. The internal pressure level coefficient is 0.0745 in the homogeneous-debt sample, but it weakens after lagged and future pressure enter together. This pattern is consistent with persistent BDC valuation style or slow-moving portfolio stress. The paper therefore treats level estimates as descriptive and places the main weight on changes.

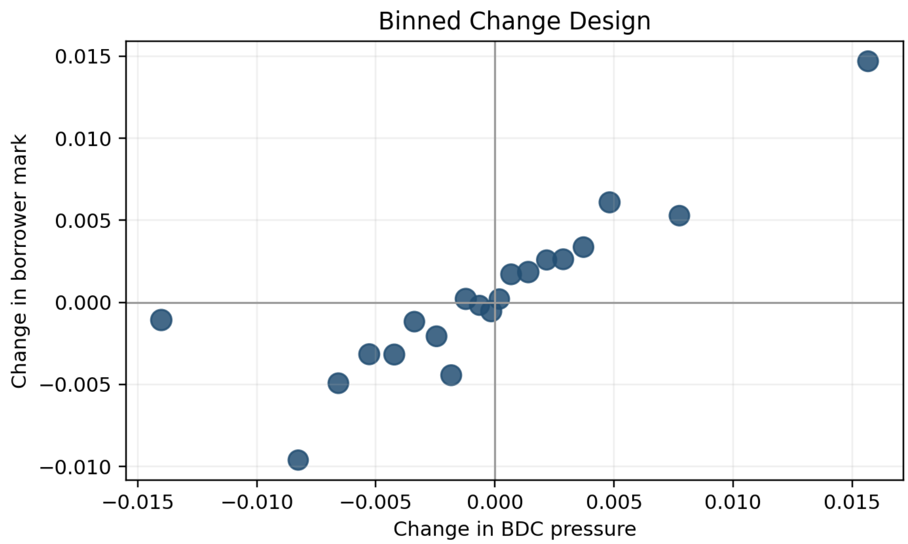

Figure 3: BDCs whose other portfolio positions deteriorate more also report lower values for shared borrowers. The binned relationship visualizes the main reduced-form result in changes.

Online Appendix Table A.5 reports the full security-comparability tests. If the internal-change result only reflected different tranches, it should disappear when the comparison holds fixed more security characteristics. The coefficient falls but remains positive and then stabilizes. The flexible-security estimate, which holds debt class and seniority fixed within the borrower-date, is 0.0546; the strict-security estimate, which additionally matches on rate type and maturity, is 0.0616. The two are statistically indistinguishable, so the relation settles at roughly 0.05 to 0.06 once seniority is controlled rather than continuing to decline. The flexible cell is the conservative, seniority-controlled anchor; the strict cell is reported as a consistent but match-limited diagnostic, because the validation audit shows that fully automated strict matching is noisy.

The reduction in magnitude is also informative. If borrower-date estimates were entirely driven by BDCs holding different instruments of the same borrower, the coefficient should collapse under tighter security cells. If the strict-cell estimate were identical to the borrower-date estimate, the security restrictions would not be doing much work. The observed pattern is in between: comparability matters, but it does not eliminate the relation between internal pressure changes and reported marks.

8.1 The Effect Concentrates in Discretionary Debt

These estimates raise a sharper question: *where* in the capital structure does the pressure–mark relation live? If reported marks move with lender pressure because valuation involves discretion, the relation should be strongest where discretion is largest—junior and structured claims whose fair values rely on models and judgment—and weakest in first-lien senior secured debt, whose marks are pinned near par by collateral and observable comparables. Table 8 tests this directly, estimating the internal-pressure change specification within seniority classes.

Table 8: Seniority Gradient: the Pressure–Mark Relation Concentrates in Discretionary Debt

	First lien	Second lien	Subordinated / mezzanine	All non-first-lien debt	Homogeneous debt (headline)
Pressure-change estimate	-0.012	0.123**	0.396	0.143***	0.126***
Standard error	(0.030)	(0.052)	(0.271)	(0.040)	(0.031)
SD of Δmark	0.041	0.040	0.035	0.039	0.040
Observations	32,386	1,787	794	35,685	65,671
BDCs	88	39	32	99	129

Internal-pressure change specification (first differences, borrower-date and BDC absorbed, standard errors clustered by BDC). Each column restricts to a seniority class. The relation is absent in first-lien senior secured debt and concentrated in junior and structured debt, even though first-lien marks are no less volatile (SD of Δmark). Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The contrast is sharp along one margin in particular. In first-lien senior secured debt the relation is essentially zero (-0.012, not significant). It is positive and economically large across non-first-lien debt (0.143) and positive in the second-lien cell (0.123), while the most subordinated cell is too thinly populated to estimate precisely. (The seniority columns are drawn from the full multi-holder debt sample and do not partition the headline cell.) The first-lien null is not an artifact of static marks: the standard deviation of mark changes is essentially identical across seniority classes, around 0.04. First-lien marks move as much as junior marks; they simply do not move *with lender pressure*. It is the clearest evidence against a mechanical repricing story: if pressure simply moved all of a holder's marks, first-lien marks, which are equally volatile, would move too. Instead the relation is confined to claims whose reported values rest on models and judgment rather than on collateral and observable prices. The same composition logic explains why the security-cell estimates in Table 7 (flexible 0.055, strict 0.062) are smaller than the homogeneous-debt estimate and serve as the conservative, seniority-controlled anchor: the cell fixed effects partly absorb the seniority composition that carries the effect.

8.2 Discretion or Superior Information?

The remaining interpretive question is whether the lender-level component reflects *discretion* in valuation or *superior borrower information* held by some lenders. The two are observationally similar within a borrower-date cell, but they make opposite predictions about the future. If cross-holder differences encode private information about the borrower, then disagreement today—and especially a low mark from the most pessimistic holder—should forecast the borrower's subsequent deterioration. If the differences instead reflect lender-side discretion, they should carry no such borrower-level signal.

Table 9 runs this test on a borrower-date panel with calendar-date fixed effects and standard errors clustered by borrower. Neither same-borrower mark dispersion nor the most-pessimistic-holder gap predicts the borrower's future change in markdown or its future nonaccrual share; all four estimates are statistically indistinguishable from zero, and the markdown estimates are, if anything, the wrong sign for an information story. The nonaccrual forecasts are tightly bounded, with 95 percent confidence intervals within roughly ± 0.01, while the markdown forecasts are less precise but wrong-signed, so the data rule out all but small positive forecasting effects. Cross-holder mark differences thus do not detectably anticipate realized borrower outcomes. This weighs against a superior-information interpretation and is consistent with a lender-condition component in judgment-based marks: the differences track the holder's situation, not the borrower's future.

Two readings of that lender component remain. A shared-borrower mark may move because portfolio stress invites closer scrutiny, or because the holder applies a common shift in valuation inputs across its judgment-based positions. The borrower-date comparison nets out market-wide input changes, so the relevant version is holder-specific, and the data here cannot fully separate the two. Both, however, are lender-level rather than borrower-level, which is what the comparability reading requires: the same loan is valued differently across holders for reasons tied to the holder, not the borrower.

The same table explains why market pressure is not the headline result. Public market pressure is positive in broad borrower-date public-BDC cells, but it weakens or turns negative in homogeneous-debt and strict security-cell samples. The main security-comparability evidence is therefore internal-pressure changes.

The same result disciplines the mechanism interpretation. If the main result were only a broad public-market repricing story, the stock-pressure variables would be expected to

Table 9: Do Cross-Holder Mark Differences Forecast Borrower Deterioration?

Predictor (today)	Outcome	Estimate	SE	Obs.
Mark dispersion	Δ markdown (t+1)	-0.065	(0.052)	23,063
Most-pessimistic gap	Δ markdown (t+1)	-0.093	(0.062)	23,063
Mark dispersion	Nonaccrual share (t+1)	0.002	(0.004)	23,063
Most-pessimistic gap	Nonaccrual share (t+1)	0.002	(0.005)	23,063

Borrower-date panel of homogeneous-debt cells with at least two BDC holders; calendar-date fixed effects; standard errors clustered by borrower. "Most-pessimistic gap" is the highest holder markdown minus the borrower-date mean. Nonaccrual regressions control for the current nonaccrual share; markdown regressions control for the current mean markdown. No predictor is statistically distinguishable from zero. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

remain strong in the same tight security cells. Instead, the strongest robust pattern is internal pressure changes. This points the paper toward lender-level valuation discipline and portfolio pressure rather than a pure public-market monitoring channel.

9 Robustness, Denominator Checks, and Diagnostics

The most serious threat is persistent valuation style. Dynamic diagnostics show that internal-pressure levels weaken after lag and future controls, while internal-pressure changes remain strong.

9.1 Persistent BDC Style and Dynamic Diagnostics

This threat is not minor. A BDC that usually marks conservatively may also have a higher average portfolio markdown. A BDC concentrated in a troubled sector may look pressured for several quarters. A persistent valuation process can therefore create a positive level coefficient even without a contemporaneous pressure channel. The dynamic tests are included to discipline that interpretation.

Online Appendix Table A.6 reports the dynamic specifications. The lag-current-future internal-pressure level specification weakens materially, whereas the internal-change specification remains economically large. This evidence rules out a simple interpretation of pressure levels as transitory shocks.

The future-pressure terms are not nuisance controls: they show that BDC pressure is persistent and that part of the level relation reflects stable BDC-level differences. This is why the change specification, which asks whether changes in the rest of the portfolio move with changes in the reported value of a shared borrower, carries the main weight.

A direct placebo confirms that the change design is not picking up a pre-trend. Figure 4 adds leads and lags of portfolio-pressure changes to the headline change specification. The contemporaneous coefficient is positive and significant, while both the lagged and the *lead* pressure-change coefficients are economically negligible and statistically indistinguishable from zero (lead estimate 0.021, $p = 0.45$). The absence of a lead effect is what matters here: future deterioration in the rest of the portfolio does not predict current changes in the shared-borrower mark. That is what separates the change design from the level specification, where future pressure looks predictive only because pressure levels are persistent and partly anticipated. The change estimate therefore reflects contemporaneous co-movement, not an anticipatory or mechanical trend. Adding the leads and lags leaves the contemporaneous coefficient essentially unchanged (0.082 without the controls versus 0.086 with them on the same sample), so the difference from the full-sample headline of 0.126 is attributable to the three-consecutive-quarter requirement, not to the dynamic controls absorbing the effect.

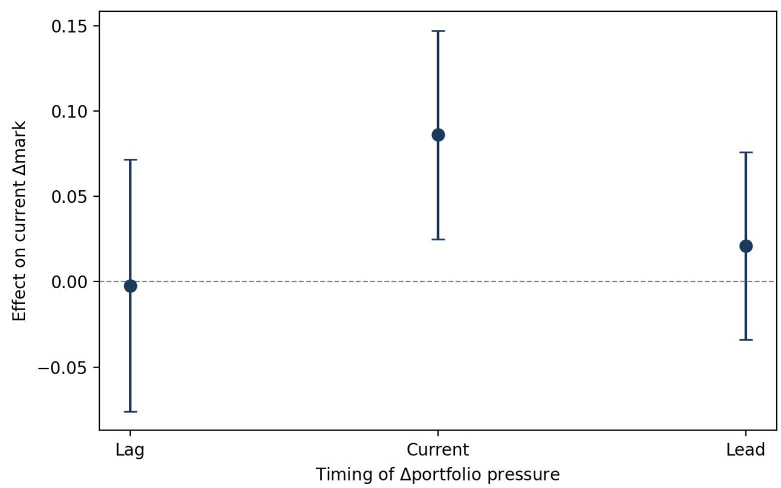

Figure 4: Lead–lag placebo for the change design. Only the contemporaneous portfolio-pressure change moves the shared-borrower mark; the lead coefficient is zero, so future portfolio deterioration does not predict current mark changes. Bands are 95 percent confidence intervals with standard errors clustered by BDC.

The first-difference sample is also not selected. Positions retained in the change specification, which require a prior-quarter mark, are well balanced against the positions dropped for lacking one: the standardized difference in the mark level is 0.05 and in portfolio pressure 0.01, both far below conventional balance thresholds. The change result is therefore not an artifact of which positions survive the prior-quarter requirement.

The result is also not confined to a single period. Splitting the homogeneous-debt change sample at the start of 2024 gives the same sign in both halves, with the relation strongest in the more recent, higher-stress window (0.139, $p < 0.01$, 2024–2025) and positive but smaller and imprecise in the earlier window (0.067, 2022–2023). This is the pattern the discretion mechanism implies: portfolio pressure has more scope to move reported marks when portfolios are under greater stress. The relation is not an artifact of any single quarter.

Online Appendix Table A.7 reports additional robustness and diagnostic tests. The level result is stable when the design requires at least three or five BDC holders, excludes the largest BDCs, excludes the largest borrowers, and winsorizes the dependent variable. The shuffled-pressure placebo is close to zero, which supports the idea that the result is not generated by arbitrary date-level variation. The future-pressure and fake-cell diagnostics are significant in the *level* design; the change design, by contrast, passes its own placebo (the lead–lag test below). These level diagnostics are why the paper does not read pressure levels as clean shocks. They show that persistent BDC-level marking patterns are real and must be built into the interpretation.

A second concern is the denominator itself. Cost is widely available but is not always par. Raw FV/principal is contaminated by missing principal, near-zero principal, scale mismatches, partially funded commitments, delayed-draw facilities, revolvers, hybrid instruments, and restructuring cases. The cleaned-principal analysis removes clear data-error categories.

9.2 Denominator Measurement

Online Appendix Table A.8 shows that cleaned FV/principal supports internal-pressure changes. The clean FV/principal internal-change coefficient is 0.1392, and the 1/99 winsorized cleaned-principal coefficient is 0.1409. Both estimates are close to the cost-based internal-change coefficient. Market-pressure principal-denominator estimates are small and imprecise.

Online Appendix Table A.9 reports the taxonomy of principal-denominator problems. Most problem rows have missing principal, but the economically important cases are not limited to missingness. Near-zero principal, scale mismatches, unfunded commitments, delayed-draw facilities, and hybrid securities can make FV/principal explode even when the fair-value field itself was correctly parsed. This is why raw principal-denominator estimates are not used as a main outcome.

The denominator evidence supports a conservative conclusion. The internal change result survives when principal can be cleaned enough to create an economically meaningful denominator. The market-pressure result does not show the same robustness in principal-denominator tests. This again reinforces the paper's evidence hierarchy: internal-pressure changes are the main result; market pressure is supporting evidence.

The remaining threat is parser and matching error. The source-filing validation audit opens the SEC filings behind the 4,000-row validation sample and checks whether borrower tokens, security-label terms, and reported numerical fields appear in the filing text. This turns validation into an empirical object of the paper rather than a background coding exercise.

9.3 Parser, Matching, and Same-Security Validation

The validation audit has two messages. First, the numerical fields are reliable: cost and fair value are found in the source filings in almost all audited rows. This supports the measurement contribution. Second, fully automated same-security matching is not reliable enough: borrower matching and security classification fail too often in the strict validation sample. This means the paper should not make strict-security cells the main identification design.

The audit therefore strengthens one part of the paper and narrows another. It strengthens the claim that public BDC filings can be used to measure reported cost and fair value at scale. It narrows the claim that fully automated strict same-security cells can serve as the main design. For that reason, the main evidence in the paper comes from homogeneous-debt and flexible-security comparisons. Strict-security cells remain useful diagnostics, but not the main source of identification.

Table 10: Source-Filing Validation Audit

Audit item	Pass rate	Interpretation
Cost extraction	99.7%	Numeric mark field reliable
Fair-value extraction	99.7%	Numeric mark field reliable
Principal extraction conditional on availability	99.7%	Reliable when present
Borrower matching	57.0%	Too noisy for strict matching
Debt/equity classification	69.0%	Needs caution
Seniority conditional on debt-like rows	95.6%	Good conditional on debt-like rows
Strict-security usability	3.2%	Diagnostic only

The audit validates the measurement layer but limits the strict same-security design.

9.4 Adviser-Family Dependence and Inference

Several BDC vehicles can share an adviser family, valuation infrastructure, portfolio managers, valuation vendors, and reporting conventions. Treating each vehicle as a fully independent valuation decision could therefore overstate the effective number of independent marks. To address this issue, I construct manually reviewed adviser-family identifiers from BDC names, tickers, and platform names and rerun the internal-pressure change specification under alternative dependence assumptions.

Table 11 reports inference robustness. The internal-pressure change result remains positive when standard errors are clustered by BDC, adviser family, borrower, and two-

Table 11: Inference Robustness

Specification	Estimate	SE	Observations	Clusters
Baseline BDC clustering	0.126***	0.031	65,671	129
Adviser-family clustering	0.126***	0.035	65,671	87
Borrower clustering	0.126***	0.031	65,671	4991
Two-way BDC and borrower clustering	0.126***	0.038	65,671	129
Two-way adviser-family and borrower clustering	0.126***	0.038	65,671	87

The internal-pressure change result remains positive under alternative clustering choices.

way combinations of BDC or adviser family with borrower. The standard errors increase under two-way clustering, as expected, but the coefficient remains economically large and statistically distinguishable from zero. Because the identifying variation is lender-level and the number of BDC clusters is modest, I also compute a Rademacher wild cluster bootstrap by BDC with 399 replications. The bootstrap p-value for the homogeneous-debt change estimate is below 0.001 (bootstrap point estimate 0.107), so few-cluster inference does not overturn the result.

Online Appendix Table A.10 reports additional adviser-family diagnostics. The estimate remains positive in cells with at least two adviser families, after collapsing observations to adviser-family-by-cell marks, and after excluding large families one at a time. These results reduce the concern that the main pattern is driven by one large adviser platform and address dependence across vehicles managed by the same platform.

9.5 Stress-Period Extension: March 2026

March 2026 tests whether the measurement system extends to a stress-period release with a changed SEC schema. The old parser recovers no validated March-2026 rows, while the fallback mapping recovers 26,917 validated rows. Automated filing-text checks validate nearly all sampled cost and fair-value fields. March 2026 is therefore no longer a pure parser failure.

The detailed March 2026 schema and result tables are reported in the online appendix.

Online Appendix Table A.15 reports provisional March results. March-only internal-level estimates are imprecise, and the public-BDC March-release subsample has fewer than 10 BDCs for market-pressure regressions. The paper therefore excludes March 2026 from headline results and presents it as a stress-period extension.

The March evidence nonetheless shows that the measurement system can be extended after a schema break, while illustrating why stress-period evidence should not be pooled mechanically with the headline sample. A stress-period release changes both the economic environment and the data pipeline. The conservative approach is to document March 2026 separately, validate the parser change, and then decide whether the extended sample can be pooled with the pre-2026 panel.

10 Market Pressure as Supporting Evidence

Because public prices are externally determined, I also ask whether BDC stock-market stress predicts lower reported marks. In broad borrower-date listed-BDC cells, the answer is yes. The broad market-pressure panel has 31,289 listed-BDC borrower-date rows, while Table 12 uses the narrower regression sample after requiring the specific 90-day stock-pressure window, nonmissing return data, report-date alignment, and fixed-effect support. Table 13 reports this attrition. A 10 percentage point worse 90-day BDC stock return predicts roughly a 0.52 percentage point lower reported value for a shared borrower.

Table 12: BDC Stock-Market Stress as Supporting Evidence

Measure	Estimate	SE	Observations	Cells	BDCs
90-day stock pressure	0.052**	0.021	5,426	2,537	42
90-day stock pressure, change	0.038***	0.013	3,616	1,724	36
90-day residualized pressure	0.024	0.015	6,847	3,193	42
90-day abnormal pressure	0.052**	0.021	5,426	2,537	42
NAV discount, 30d before report	0.025	0.015	6,612	3,093	42

Market pressure is strongest in broad cells and weakens in tighter security comparisons, so it is supporting evidence rather than the main identification result.

Table 13: Market-Pressure Sample Attrition

Step	Observations	Dropped since prior step
Broad listed-BDC market-pressure rows	31,289	0
Narrow 90-day stock-pressure regression sample	5,426	25,863

This panel reconciles broad listed-BDC market-pressure rows with the narrower 90-day stock-pressure supporting-evidence regression sample.

Table 12 reports these results. The 90-day stock-pressure measure is the strongest market signal. Residualized market pressure and NAV-discount measures have the expected sign but are imprecise. Because market pressure weakens in strict-security cells, the correct interpretation is broad BDC stress rather than a strict same-security mark adjustment channel.

Online Appendix Table A.11 shows that the 90-day window is the strongest timing window. The 30-day and 60-day windows are weak, while the 180-day window is positive but less precise. This pattern is consistent with the timing of BDC reporting and market information.

Public equity investors may anticipate valuation deterioration before the filing date, but the mapping from daily stock prices to quarterly reported marks is necessarily noisy.

Online Appendix Tables A.12 and A.13 report two additional market-pressure diagnostics. Residualized stock pressure removes a common BDC market component, and NAV discounts compare market value to reported net asset value. Both measures generally have the expected sign, but they are not precise enough to support a separate headline result, which is why the broad market-pressure evidence supplements rather than replaces the internal-pressure change result at the center of the paper.

Online Appendix Table A.14 reports future-pressure diagnostics for market pressure. Future market pressure is sometimes predictive, which is exactly what one should expect if public BDC stock prices partly anticipate future valuation news. This is another reason to avoid a strong structural interpretation. Public-market pressure is informative about stress, but it is not a clean instrument for valuation discipline.

11 Conclusion

Private-credit valuations are hard to observe, but public BDC filings make part of the market visible. This paper uses those filings to build a borrower-date panel of reported values and to compare marks across BDCs holding the same borrower at the same reporting date. The first result is descriptive: reported values for shared borrowers differ across holders. The second result is empirical: deterioration in the rest of a BDC's portfolio predicts changes in the value it reports for shared borrowers. This relationship is confined to junior and structured claims and is absent in equally volatile first-lien debt, and cross-holder disagreement does not detectably forecast the borrower's future deterioration, together pointing to a lender-level rather than information-based source of the differences.

The paper's contribution is twofold: a scalable public-data measurement of private-credit valuation, and reduced-form evidence that reported marks carry a lender-level component. Lower marks need not mean misvaluation—they can reflect information, security differences, valuation governance, or conservative judgment—but the comparability implication holds regardless of which channel dominates: the same loan is valued differently across holders depending on each holder's portfolio condition. Reported private-credit marks and the net asset values built from them therefore cannot be read as pure borrower signals.

The validation audit is central to this interpretation. It shows that cost and fair-value extraction from public filings is highly reliable, which supports the measurement contribution. It also shows that fully automated same-security matching is too noisy to be the main design. A credible reading of the evidence therefore puts the weight on homogeneous-debt and flexible-security comparisons, uses strict-security cells as diagnostics, treats market pressure as supporting evidence, and keeps March 2026 as provisional appendix evidence.

11.1 Scope and external validity

The results describe the public-reporting segment of private credit and should be read at that scope. BDCs are the part of the market where position-level marks are observable; reported fair values are accounting marks rather than transaction prices, and the lender-level component documented here is a reduced-form relationship rather than a structural effect on true loan values. The same features make the setting usable: BDC filings provide repeated, public, position-level observations of comparable private-credit claims, and the same-borrower design uses that structure to hold borrower fundamentals fixed while comparing holders.

Public filings can make private-credit valuation less opaque. They do not reveal true transaction prices or the entire private-credit market, but they let outside researchers and regulators observe reported marks at the borrower level and study how those marks differ across holders.

References

Altavilla, Carlo, Marco Pagano, and Saverio Simonelli. 2017. "Bank Exposures and Sovereign Stress Transmission." *Review of Finance*.

Avalos, Fernando, Sebastian Doerr, and Gabor Pinter. 2025. "The Global Drivers of Private Credit." BIS Quarterly Review.

Beatty, Anne, and Scott Liao. 2014. "Financial Accounting in the Banking Industry: A Review of the Empirical Literature." *Journal of Accounting and Economics*.

Berrospide, Jose, Fang Cai, Siddhartha Lewis-Hayre, and Filip Zikes. 2025. "Bank Lending to Private Credit: Size, Characteristics, and Financial Stability Implications." FEDS Notes.

Bank for International Settlements. 2024. *Annual Economic Report*.

Block, Joern, Young Soo Jang, Steven N. Kaplan, and Anna Schulze. 2023. "A Survey of Private Debt Funds." NBER Working Paper No. 30868.

Brown, Gregory W., Oleg R. Gredil, and Steven N. Kaplan. 2019. "Do Private Equity Funds Manipulate Reported Returns?" *Journal of Financial Economics*.

Bushman, Robert M., and Christopher D. Williams. 2012. "Accounting Discretion, Loan Loss Provisioning, and Discipline of Banks' Risk-Taking." *Journal of Accounting and Economics*.

Chernenko, Sergey, Robert Ialenti, and David Scharfstein. 2025. "Bank Capital and the Growth of Private Credit." Working paper.

Couts, Spencer J. 2023. "Valuation Estimates and Return Smoothing: Evidence from Private Equity Real Estate Assets." Working paper.

Cumming, Douglas, and Uwe Walz. 2010. "Private Equity Returns and Disclosure Around the World." *Journal of International Business Studies*.

Erel, Isil, Thomas Flanagan, and Michael Weisbach. 2024. "Risk-Adjusting the Returns to Private Debt Funds." NBER Working Paper No. 32278.

Financial Stability Board. 2026. "Report on Vulnerabilities in Private Credit."

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*.

Geltner, David, Bryan MacGregor, and Gregory Schwann. 2003. "Appraisal Smoothing and Price Discovery in Real Estate Markets." *Urban Studies*.

Getmansky, Mila, Andrew W. Lo, and Igor Makarov. 2004. "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns." *Journal of Financial Economics*.

Gilchrist, Simon, and Egon Zakrajsek. 2012. *American Economic Review*.

Hanson, Samuel G., and Jeremy C. Stein. 2015. "Monetary Policy and Long-Term Real Rates." *Journal of Financial Economics*.

Huizinga, Harry, and Luc Laeven. 2012. "Bank Valuation and Accounting Discretion During a Financial Crisis." *Journal of Financial Economics*.

Investment Company Institute. 2026. "Valuation Governance Considerations for Private Credit Assets in Regulated Funds."

International Monetary Fund. 2024. "The Rise and Risks of Private Credit." Global Financial Stability Report, Chapter 2.

Jang, Young Soo, and Ginha Kim. 2026. "Information Asymmetry and Valuation Intermediation in Private Credit." SSRN Working Paper.

Jenkinson, Tim, Miguel Sousa, and Ruediger Stucke. 2013. "How Fair Are the Valuations of Private Equity Funds?" *Review of Accounting Studies*.

Jimenez, Gabriel, Steven Ongena, Jose-Luis Peydro, and Jesus Saurina. 2012. "Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications." *American Economic Review*.

Jorda, Oscar, Moritz Schularick, and Alan M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking*.

Kaplan, Steven N., and Antoinette Schoar. 2005. "Private Equity Performance: Returns, Persistence, and Capital Flows." *Journal of Finance*.

Khwaja, Asim Ijaz, and Atif Mian. 2008. "Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market." *American Economic Review*.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates." *Brookings Papers on Economic Activity*.

Laux, Christian, and Christian Leuz. 2010. "Did Fair-Value Accounting Contribute to the Financial Crisis?" *Journal of Economic Perspectives*.

Mian, Atif, Amir Sufi, and Emil Verner. 2017. "Household Debt and Business Cycles Worldwide." *Quarterly Journal of Economics*.

U.S. Securities and Exchange Commission. 2026. "Business Development Company (BDC) Data Sets."

Suhonen, Antti. 2024. "Direct Lending Returns." *Financial Analysts Journal*.

Vyas, Dushyantkumar. 2011. "The Timeliness of Accounting Write-Downs by U.S. Financial Institutions During the Financial Crisis of 2007–2008." *Journal of Accounting Research*.

Zou, Jason. 2026. "Private Credit Markets: Theory, Evidence, and Emerging Frontiers." Working paper.